

Minerva Access is the Institutional Repository of The University of Melbourne

Author/s:

Lampe, B;de Fontenay, C;Nugent, J;Jomini, P

Title:

Climbing the Jobs Ladder Slower: Young People in a Weak Labour Market

Date:

2022-03-01

Citation:

Lampe, B., de Fontenay, C., Nugent, J. & Jomini, P. (2022). Climbing the Jobs Ladder Slower: Young People in a Weak Labour Market. Australian Economic Review, 55 (1), pp.40-70. <https://doi.org/10.1111/1467-8462.12444>.

Persistent Link:

<https://hdl.handle.net/11343/302907>

Climbing the jobs ladder slower: Young people in a weak labour market

Bryn Lampe*, Catherine de Fontenay, Jessica Nugent, Patrick Jomini

Government of Australia Productivity Commission, Melbourne, Victoria, Australia

*Correspondence Bryn Lampe, Government of Australia Productivity Commission, Melbourne, Victoria, Australia. Email: bclampe@gmail.com

Abstract

The downturn in Australia's labour market from 2008-2018 was reflected in dimensions other than unemployment. We use the Household Income and Labour Dynamics in Australia dataset to explore one key dimension — occupational choice. Using an occupational ranking (AUSEI06) based on earning potential and education level, we show young graduates (aged 20-35) found work in lower-scored occupations after 2010. We examine potential heterogeneity of the occupational outcomes using Markov transitions between quartiles of the occupational distribution. We find the likelihood of transitioning to better quartiles worsened over the 2001-2018 period, suggesting poor initial outcomes can have long-term effects on one's occupation.

1 Introduction

The Global Financial Crisis (GFC), followed by the slowing of the mining boom, ushered in a weak labour market in Australia. These macroeconomic shocks weakened labour demand, while higher participation rates and increases in educational attainment greatly increased supply, particularly for high-skilled jobs (Productivity Commission 2020a).

Notably, however, the weak labour market after 2008 was not marked by a high unemployment rate. The unemployment rate did jump two percentage points in 2008-2009, but it then remained steady and did not return to its higher pre-2001 levels until the COVID-19 crisis. This was in contrast to past downturns, such as the 1991 recession, during which unemployment rose much higher. Evidence suggests that the

This is the author manuscript accepted for publication and undergone full peer review but has not been through the copyediting, typesetting, pagination and proofreading process, which may lead to differences between this version and the [Version of Record](#). Please cite this article as [doi: 10.1111/1467-8462.12444](https://doi.org/10.1111/1467-8462.12444).

labour market adjusted through mechanisms other than unemployment (PC 2020a). For example, full-time employment declined and part-time employment increased among workers aged 15-24 (PC 2020a).¹ Another adjustment mechanism is wage rates: young people aged 20-34 experienced nearly zero wage rate growth from 2008 to 2018 (PC 2020a), in part because of lower starting wages.

This paper focuses on one key adjustment mechanism, the choice of occupation. We examine the occupations of young people after they graduate, using data from the Household Income and Labour Dynamics in Australia (HILDA) survey from 2001 to 2018 and occupational scores developed by McMillan, Beavis and Jones (2009).² As a first step, we show that occupational scores for recent graduates declined from 2008 to 2018, when controlling for demographic characteristics and education. Occupational scores remained lower even four or six years after graduation.

But this decline in average occupational scores masks significant heterogeneity in outcomes. Both university graduates and vocational education graduates are distributed across a wide range of occupational scores. The decline in occupational scores has not been uniform – rather, some young workers remain in very high-scored occupations, while there have been more ‘unlucky’ people than in earlier years, obtaining work in occupations with much lower occupational scores. A natural follow-up question is whether these unlucky young workers are gradually working their way back to their desired occupations over the next few years. Has mobility from the lower-scored occupations increased now that more young people are ‘unlucky’? If workers can catch up, then the effects of entering a weak labour market would not be permanent. If, instead, the effects are permanent, this could lead to a phenomenon called ‘scarring’.

Scarring occurs when workers who experience a negative shock in the labour market suffer negative long-term consequences. This negative shock could be a spell of unemployment (Arulampalam, Gregg and Gregory 2001; Knights, Harris and Loundes 2002), or it could be a spell of underemployment or of low-wage work (Buddelmeyer, Lee and Wooden 2010; Chalmers and Hill 2007; Fok, Scutella and Wilkins 2015; Fouarge and Muffels 2009; Mosthaf 2014).³ More recent work on scarring has

¹ The decline took place among young people who were not studying full-time or part-time, suggesting that the decline in hours was involuntary (PC 2020a).

² These occupational scores are a socioeconomic index that connects a person’s education with their earning potential. The index provides researchers with a ‘means of assigning sociologically meaningful occupational status scores to data coded in accordance with the official occupational classifications of the Australian Bureau of Statistics (ABS)’ (McMillan, Beavis and Jones 2009). The index allows us to connect a person’s occupation with a quantifiable continuous measure.

³ ‘Wage scarring’ occurs if a long spell of unemployment leads to lower wage rates in the future (Arulampalam 2001). The earnings of graduates who enter the labour market in a recession are lower than those of other cohorts (Altonji, Kahn and Speer 2016; Oreopoulos, von Wachter and Heisz 2012).

considered the ‘quality’ of jobs. It is possible that holding a job that does not fully use one’s skills can raise the probability of unemployment or reduce future wage rate growth (Guvnen et al. 2020; Mavromaras, Sloane and Wei 2015; Naidoo, Packard and Auwalin 2015).⁴ And a negative shock today can lead to holding lower-scored occupations in the future. For instance, Kahn (2010) found that in the United States, cohorts that graduated when unemployment was high were in lower-scored occupations on average, and that this effect had not disappeared after 10 years.⁵

Regressions are ill-suited to answer whether transitions from lower-scored occupations have become more likely over time. As our main exercise, we undertake discrete-time Markov chain analysis of transitions between quartiles of the distribution of occupational scores, controlling for demographic characteristics using a matching procedure. This is the first paper to apply Markov analysis to occupational scores. We examine transition rates for young Australians who graduated between 2001 and 2014.

We see no evidence that transitioning to better outcomes became more likely after 2008, suggesting that poor initial outcomes are likely to have long-term effects on one’s occupation. If anything, when we examine young graduates between 2013 and 2014, people who found work in lower quartiles of the score distribution were more likely to remain in those quartiles compared with earlier cohorts. Thus, young people who obtained a lower-scored job in the post-2008 weak labour market are likely to remain on a low-scored job trajectory.

There is no evidence that employers are now more willing to consider an applicant from a lower-scored occupation, taking into account that some high-skill workers may have been unlucky. This may be because the labour market is still weak, and there are many high-skill applicants among recent graduates. Also, it could be that occupational mobility will recover when the labour market finally improves (which would now be well after the COVID-19 crisis). But the fact that the weak labour market lasted for at least a decade means that many young workers may face scarring.

⁴ Other authors who look at occupational scores include: Sicherman (1991) — who started a literature that looked at the occupational mobility of over-educated and over-skilled people — and Ralston et al. (2016) — who looked at whether persistent unemployment leads to occupational scarring. Unlike our paper, these papers are not concerned with scarring effects that might be specific to downturns.

⁵ Although Kahn (2010) shows that the effects of a downturn on occupational scores can endure, it is not clear whether the negative effect was from longer job tenure or from getting caught on a low-skilled job trajectory: Kahn also found that workers who graduated in a downturn were less likely to change jobs. It is important to examine subsequent transitions because the primary mechanism through which workers increase their wages is moving to better-paying employers (Berge 2018, Deutscher 2019). For young workers, these moves are particularly important because young workers have the largest wage rate gains and losses associated with moves up and down the occupational ladder (Forsythe 2019).

The rest of the paper proceeds as follows: Section 2 presents the data used and descriptive statistics. Section 3 outlines the methodology we use throughout the paper, including the regressions of occupational status and the Markov Analysis of occupational status, including details of the matching technique. Section 4 presents the results of the regressions and of the transitions between quartiles of the occupation distribution to show that the transitions did not improve. Section 5 conducts a robustness check of the results.

2 Data

The Household Income and Labour Dynamics in Australia (HILDA) is a high-quality panel dataset beginning in 2001; the latest year available is 2019. Over 17 000 individuals are surveyed in each year. The panel tracks members of the original dataset, and those who join their households, thus recent immigrants are under-sampled; an additional sample was added in 2011 to help correct this. Table 1 lists sample statistics for the young people in the HILDA dataset, according to the year in which they finished their studies.

We use the AUSE106 scale, a 0-100 scale of occupational scores based on the ABS occupation classifications in 2006 (McMillan, Beavis and Jones 2009).⁶ Broadly, the occupational score is a function of the average education required for an occupation and the average earnings in that occupation — for example, neurosurgery has a high score. Table 2 provides some examples of occupational scores.

We use the 2006 scores to analyse the entire 2001–2018 sample period. An occupation has the same score throughout the sample period, even if the average level of education or the earnings in an occupation changed over the period. Keeping the scores constant allows us to verify whether there was a change in the occupations that young people were able to find work in.

The resulting scores can be thought of as being ordered along an occupational ladder (in a rough sense). It is worth bearing in mind that individual job seekers have their own preferences and, of course, the occupational score (or earning potential) is not the only important characteristic of a job. However, if we see a significant decline in occupational scores for young workers over time, it is likely to reflect an involuntary worsening of their employment outcomes.

⁶ The occupational scale replicates for Australia the methodology behind the widely-used International Socioeconomic Index, as described in Ganzeboom et al. (1992). Essentially, using 2006 Census data on earnings and education, the algorithm attempts to explain earnings as a function of education and occupation, minimising the direct effect of education on earnings and maximising the indirect effect of education on earnings through its impact on occupation. The rank ordering of occupations is derived from these results.

Figure 1 plots the empirical probability density functions (PDFs) of occupational scores in 2001 and 2018. The most noticeable feature of the occupation distributions is that they are bimodal (i.e. with two peaks) for all age groups and education levels. One implication is that people who are not successful in obtaining a job in the right-hand peak may find themselves in the left-hand peak, a long way down the occupation distribution. The left-hand peak is much higher for workers aged 20-24, possibly because some of them held low-wage jobs during their studies. Given how few people are in the middle of the distribution, one implication is that some of these young people are likely to move a long way up in the distribution over the course of their life. We will analyse where young people move in the occupation distribution in section 3.

When we look at how the distribution has changed over time, the overall direction of change seems positive. Workers aged 25-34 and 35-54 were more likely to be in higher-scored occupations in 2018 than in 2001. The picture is more nuanced for workers aged 20-24: they were more likely to be in mid-scored occupations than in low-scored occupations relative to 2001. A relatively small number of workers aged 20-24 were in high-scored occupations, but the number was even smaller by 2018.

Next we graph how occupational outcomes changed for a worker of a given education level. Average educational attainment increased steadily over the past 20 years (PC 2019), and the mix of occupations is shifting towards those with higher education requirements. However, the increased supply of graduates might worsen the prospects of the average graduate. In figure 2, we examine outcomes for graduates of different ages. Graduates include people with a vocational education qualification (VET)⁷, identified as ‘sub-bachelor’, and people with a bachelor or post-graduate qualification.

Once we control for education, a gloomier picture emerges — the distribution of occupations for graduates under 35 worsened (figure 2). People aged 25-34 with bachelor degrees are substantially worse off in 2018 than in 2001, particularly in their likelihood of securing a high-scored occupation. Outcomes have slightly worsened for other groups with bachelor’s degrees. Outcomes for people aged 20-24 have slightly deteriorated, but mostly at the very low-score end of the distribution, while outcomes are also worse for people aged 35-54 in the middle and low-score end of the distribution.

This is consistent with the growth in the supply of skilled people exceeding the growth in demand for those skills, resulting in increased competition for high-scored

⁷ VET courses are generally offered outside of the high school curriculum, but can include both high school and post-high school students. This study focuses on graduates aged 20-34, who are generally post-high school students.

occupations.⁸ A person looking for a high-scored occupation in 2018 may have faced a more competitive labour market, and therefore was more likely to end up in a lower-scored occupation than they would have in 2001. Note, however, that 35-54 year olds in high-scored occupations appear to have been shielded from competition, maybe by virtue of incumbency: there is little decline in their occupational scores.

When we turn to people with sub-bachelor qualifications, the probability of obtaining a high-scored occupation declined for all groups. And for people aged 20-24, the decline in average occupational score was substantial. It may be that as workers move down the ‘jobs ladder’, there was more competition in lower-scored occupations.

We graph the average occupational score of each graduation cohort from 2001 to 2015 (figure 3). While the results are noisy, they confirm that although educational attainment increased, average occupational score did not increase. As a result, young graduates towards the end of the period were accepting lower-scored occupations than earlier graduates with a comparable level of educational attainment.

It is true that graduates took longer to find their first graduate job after 2008 (Pennington and Stanford 2019, p. 58).⁹ This delay could look like a step ‘down the jobs ladder’ on a static occupational score distribution for young workers, if they held a casual job while searching for a graduate job. But figure 3 suggests that even after four years, occupational scores are still worse.

3 Methodology

We undertake two separate econometric exercises. First, to assess whether young people are in lower-scored occupations on average, we regress occupational score on demographic variables with year and cohort effects. Second, to assess whether mobility from lower quartiles has increased, we estimate Markov transitions between quartiles of the income distribution, by year and cohort.

⁸ Changes to which jobs require a qualification (‘credential creep’) can artificially change this distribution. Coelli and Wilkins (2009) showed that credential creep did occur in an earlier period (1981 to 2004) and that this may have suppressed estimates of women’s wage growth by 6 percentage points. The main changes to credentials outlined in Coelli and Wilkins (2009) were the re-classification of existing nursing and teaching qualifications to bachelor degrees. This could have the effect shown in the ‘Bachelor or above’ category in figure 2 (shifting the PDF to the left).

⁹ The number of bachelor’s degree graduates in full-time employment as a percentage of people available to work full-time hours steadily increased from the late 1990s until 2008 and then fell to its lowest rate in 17 years in 2014 (Pennington and Stanford 2019, p. 58).

3.1 Mincer equations in occupation

We begin by estimating an equation for the occupational score of graduates (both VET and university). This is similar to the classic Mincer equation, but rather than considering wage outcomes, the outcome of interest is occupational score.

$$y_{i,t} = \alpha_0 + \alpha_1 \lambda_t + \alpha_2 e_{i,t} + \alpha_3 e_{i,t}^2 + \mathbf{x}_{i,t}' \boldsymbol{\beta} + \varepsilon_{i,t} \quad (1)$$

$$y_{i,t} = \alpha_0 + \alpha_1 \lambda_t + \alpha_2 c_i + e_{i,t} \mathbf{c}_i \boldsymbol{\alpha}_3 + e_{i,t}^2 \mathbf{c}_i \boldsymbol{\alpha}_4 + \mathbf{x}_{i,t}' \boldsymbol{\beta} + \varepsilon_{i,t} \quad (2)$$

where $y_{i,t}$ is the occupational score of individual i at time t . The regressions include year dummies, λ_t (to control for economic conditions in that year), linear and quadratic terms in the years of experience of the worker, $e_{i,t}$ (measured as the number of years since graduation), and other individual characteristics, $\mathbf{x}_{i,t}$.

The second specification interacts experience with a cohort variable, c_i , a dummy variable that captures the year of graduation for individual i , measured in three-year cohorts. We restrict the sample to workers under the age of 35 in their graduation year and the first four years after graduation.¹⁰ This specification allows us to verify whether cohorts have a different path over time. If, for example, young people are taking longer to secure their first ‘real job’ after graduation, there will be a decline in initial occupational scores but the slope will be steeper.

Our third specification verifies that labour-market conditions are the main reason that different cohorts have different outcomes in the labour market. Following Kahn (2010), we include experience variables interacted with the regional unemployment rate $u_{i,g}$ when the person was graduation.¹¹

$$y_{i,t} = \alpha_0 + \alpha_1 \lambda_t + \alpha_2 e_{i,t} + \alpha_3 e_{i,t}^2 + \alpha_4 u_{i,g} + \alpha_5 e_{i,t} u_{i,g} + \mathbf{x}_i' \boldsymbol{\beta} + \varepsilon_{i,t} \quad (3)$$

We also control for the possibility that young people facing a weak labour market might delay their graduation by undertaking further study, using an instrumental variable (IV) approach, as in Kahn (2010). We instrument our experience variable — years since graduation — with years since reaching the age of 22 (with sensitivity tests for ages 21 and 23). We also instrument the regional unemployment rate at year of graduation with the regional unemployment rate at the age of 22 of the region the graduate lived in when they completed secondary school. These instruments control for

¹⁰ And then we estimate the same regression with six years of experience rather than four, which requires us to leave out the 2013-2015 cohort.

¹¹ Kahn (2010) used the mode age at graduation, which was 22 in their sample. In our sample the mode is 21.

the possibility that the graduate moved regions or delayed graduation in response to the weak labour market.

Note that these specifications do not control for any unobserved heterogeneity that might be affecting the results, contrary to many other labour market regressions such as Heckman estimation. For example, if graduates from well-off families become less likely over time to participate in the casual labour market while searching for a full-time job, occupational status will increase. Section 5 discusses the Rosenbaum tests we undertake to confirm that our results are not driven by unobserved heterogeneity, and we estimated Heckman equations in prior work to verify that there are no significant changes in participation decisions (PC 2020a).

All specifications cluster standard errors at the 1-digit ANZSCO occupational level.

3.2 Markov analysis of occupations

The Mincer regressions focus on changes in average occupational outcomes. However, an average masks considerable heterogeneity, as revealed in the bimodal distribution of occupations in section 2. In addition to understanding outcomes at any point in time, we want to understand how people move through the distribution over time. How much upward mobility is there in occupations? If more people start at low occupational scores in the later period, are more of them likely to subsequently move up in the distribution? We turn to analysing Markov transitions between quartiles in the occupational score distribution to understand how much mobility there is across occupations.

This paper uses a discrete-time Markov chain analysis (Craig and Sendi 2002) on an unmatched and then a matched sample. We separate young people into five cohorts: people who graduated from 2001 to 2003, 2004 to 2006, 2007 to 2009, 2010 to 2012 and 2013 to 2014, based on their last observed graduation date. We include both university and VET graduates.

Occupational transitions are observed for the four years after graduation — that is, someone who graduated in 2007 is observed between 2008 and 2011. These transitions are used to calculate a probability transition matrix, $\hat{M}$, where the probability of transitioning from state i to state j is the number of transitions from state i to state j divided by the number of people originally in state i .

Four years might seem too short a time to observe transitions post-graduation, but there is a trade-off between the number of cohorts that can be compared and the length of time for which each cohort can be observed. For instance, in order to include the 2013–2014 cohort, it is necessary to restrict transitions to four years after graduation.

For both the unmatched and matched samples, the distribution is cut into quartiles based on the unmatched distribution of occupational scores for young people who graduated in 2001–2003 and conditioned on two age groups: 20 to 24 and 25 to 34. The probability of transitions between these quartiles is estimated. The Markov chain

analysis, therefore, measures how quickly graduates moved up or down the occupational score quartiles from 2001-2003 for each age group.

This is a homogenous Markov chain, in that the state-to-state probabilities are assumed to be constant for all members of a cohort, over the four years after graduation. Thus the maximum likelihood estimate of the transition probability from one quartile to another is just the proportion of all observed transitions from that quartile to the quartile of interest, from this cohort.

Confidence intervals for these estimates are derived through bootstrapping, as in Craig and Sendi (2002): we simulate populations with the estimated transition probabilities and derive a 95% confidence interval for each transition estimate. Confidence intervals are reported in appendix B.

Matching methodology

One limitation of the Markov methodology is that we do not allow for transitions to differ according to an individual's characteristics. The assumption is maintained in order to have sufficient sample sizes to estimate transition probabilities for each cohort across all the quartiles. However, if transition probabilities actually differ according to demographic characteristics, then if subsequent cohorts have different demographic characteristics, the estimated transition probabilities will be different.

We resolve this issue by comparing samples that have the same demographic characteristics, using our matching process. We separate young people into one control group (people who graduated from 2001 to 2003) that we match four separate times, once each to four treatment groups (people who graduated from 2004 to 2006, 2007 to 2009, 2010 to 2012 and 2013 to 2014) based on their last observed graduation date.

We use genetic matching (Diamond and Sekhon 2013; Ho et al. 2011; Sekhon 2011) based on observables at year of graduation. Genetic matching is an iterative algorithm that minimises a generalised version of Mahalanobis distance.

The Mahalanobis distance between the X covariates for two individuals i and j is:

$$MD = \sqrt{(X_i - X_j)^T S^{-1} (X_i - X_j)}$$

where S is the sample covariance matrix of X , the characteristics of the treatment and control group together.

For each individual i , genetic matching pairs her with the individual j who minimises:

$$GMD = \sqrt{(Z_i - Z_j)^T S^{-1/2} W S^{-1/2} (Z_i - Z_j)}$$

where W is a positive definite weighting matrix of diagonal elements and $S^{-1/2}$ is the Cholesky decomposition of S^{-1} . The Z matrix can be simply the X matrix above, or it can be the X matrix augmented by propensity scores; thus the weighting matrix can allow for more or less of the weight to be attributed to propensity scores. A loss function is specified and genetic matching iteratively searches over weighting matrices and the minimised distance. In this instance, the loss function used was one of minimising the largest individual discrepancy, based on p-values from Kolmogorov-Smirnov tests and paired t-tests for all variables that are being matched on. Appendix B lists the characteristics that were used for matching (figures B.1 to B.4).

Using this matching methodology, we match each person in the treatment group to someone in the control group of 2001–2003 graduates. We match with replacement of control group observations and hold the treatment groups fixed (dropping treatment observations only if they cannot be appropriately matched).

Once each individual i in the treatment cohort has been matched to an individual k in the 2001–2003 cohort, we have the matching weights. For example, if person k was matched ω_k times, then every transition person k makes counts as ω_k transitions, not one transition.¹² We calculate a weighted transition matrix for each treatment group as well as for each control group, which we then use to calculate a matrix of transition probabilities, $\hat{M}$, for each treatment and control group.¹³ The weighted observed transitions will carry across so that the transitions probabilities are also weighted by the matching weights. The final step is to take the difference of the probability matrices of each treatment group and their respective control group: $\Delta\hat{M} = \hat{M}_T - \hat{M}_C$. Taking this difference ensures that:

1. each treatment group is compared to its corresponding control group and
2. the matching weights are incorporated appropriately (because only the control group is weighted by matching weights).

Figure 7 shows the difference in transition probabilities between the treatment and control groups, $\Delta\hat{M} = \hat{M}_T - \hat{M}_C$, where the transitions of the control group have been weighted by the matching weights.

¹² As an example, suppose rural/urban and male/female were the only characteristics of interest, and suppose that these characteristics are evenly distributed in the control (25% are rural and male, 25% are urban and male, and so on). Suppose the treatment group had 40% rural and male, and 20% of the other three groups, and the total population was the same size. Then rural and male observations in the control would have a weight of 1.6, and the other observations would have a weight of 0.8.

¹³ Each transition is also weighted by the HILDA responding person weights, so that they are representative of the Australian population more generally. Thus in practice the control group is also weighted by the HILDA weights.

The extent to which this procedure provides valid comparisons of the transitions between the treatment and control groups depends on whether the conditional independence assumption is satisfied and whether there is any omitted variable bias. If these assumptions are violated, then the changes in transitions are in part due to graduation cohort and in part due to differences in the unobserved characteristics. We minimise the risk of unobserved heterogeneity by matching on an extensive list of twenty-six demographic characteristics (Appendix B). The characteristics chosen are ones that are likely to be correlated with labour market outcomes, and also ones that proxy for parental wealth and education and language, to the extent possible, as parental background is the most likely unobserved variable. And in section 5 we undertake tests to check for the possibility of unobserved heterogeneity, using Rosenbaum bounds.

4. Results

Mincer regressions in occupation

Results for regressions (1) and (2) are reported in Table 3, column (1) and (2) respectively. The year dummies in regression (1) are negative but not significant. However, when we allow each cohort's outcomes to vary by experience (regression 2), the path of each graduate cohort follows the general pattern observed from the descriptive data in figure 3. Average cohort outcomes worsen from about 2010, as represented in Figure 4a.¹⁴ The decline in occupational scores is significant.

Changes in the labour market from year to year can mean that graduates in a particular graduate year are more likely to start in a lower-scored occupation after they graduate. Conditions in subsequent years can then also affect mobility in the first few years of a graduate's career. Figures 3 and 4 show both of these effects, with some graduate years starting at a lower average occupational score and some graduate years climbing more slowly. For instance, after controlling for observables, the 2013–2015 graduates had a lower average occupational score one year after graduation while the 2010–2012 graduates' experience function was almost flat (figure 4).

¹⁴ We also split these regressions into two samples, where one sample includes graduates with a bachelors degree or above and the other sample includes graduates with diplomas, vocational certificates I to IV and other degree types. In both samples, occupation outcomes at the end of the period are worse than at the start of the period. University graduates had better occupational outcomes in the first few cohorts, then a strong decline in outcomes in 2010–2012 and 2013–2015. Vocational graduates were in worse occupations from the 2004–2006 cohort onwards, but their situation has been slightly improving since that cohort — possibly because so many young people substituted towards university degrees. Tables available from authors on request.

If we extend the analysis to include more years of experience (Figure 4b), the last cohort in the sample (2010–2012) appears to experience the worst occupational score outcomes. Compared with the three previous cohorts, the 2010–2012 cohort begins at a low occupational score and the occupational score does not grow quickly with experience, which suggests that long-term outcomes may be declining.¹⁵ Note however that t-tests of difference between the average marginal effect of experience at each graduation cohort show no statistical difference from the 2001–2003 cohort.

Instrumental Variable results in occupation

Given that the first Mincer regression (equation 1) shows that the decline in occupational scores tended to only be statistically significant towards the end of the sample, we now turn to the IV regressions (equation 3). The baseline OLS regressions show that the unemployment rate at year of graduation has a small and statistically significant negative effect (table 4).

This negative effect would be likely decreasing occupational scores between 2011 and 2014, with the effect peaking for the 2013–2014 graduate cohort.¹⁶ This is because over this time the unemployment rate was increasing for both people aged 20–24 and 25–34 before decreasing from 2014 onward.

The results from equation (3) confirm that the worsening results for successive cohort are driven by worsening labour-market prospects. Regional unemployment at the age of 22 negatively affects occupational outcomes. As predicted by Kahn (2010), the effect is even larger if we instrument experience with years since the age of 22, although not all specifications show significance.

Kahn (2010) argues that the effect of the unemployment rate will be biased downwards by endogeneity, because some people can defer graduation until the unemployment rate is more favourable. This is confirmed in the IV results, where the effect of the unemployment rate is smaller in the OLS regressions than in the IV regressions (table 4). This endogeneity may explain why some of the Mincer regressions are not statistically significant.

¹⁵ Including a cubic experience term yields similar results of a decline for the 2010–2012 cohort (PC 2020).

¹⁶ Sensitivity tests of the matching procedure included calculations of the Average Treatment Effect for each graduate cohort both one and four years after graduation. These sensitivity tests show that the matched cohort of 2010–2012 graduates were in lower occupations both one and four years after graduation, with statistical significance at the 95 per cent level of confidence. The treatment effect for the 2013–2014 cohort one year after graduation was also statistically negative.

Markov estimation results

The Mincer results showed that recent cohorts of graduates have experienced a significant decline in their average occupational score, and that outcomes over the first six years continue to be poor. However, outcomes for the 2013–2015 cohort (figure 4a) are intriguing: while this cohort appears to begin at a much lower occupational score than earlier cohorts, their score rises very quickly. (The pattern is there in the raw data in figure 3 as well, although it is less pronounced.)¹⁷ Is it possible that later cohorts start with lower initial occupation scores, but those low scores are offset by higher mobility of the affected workers? To answer this question, we turn to Markov analysis to consider how young workers move through the occupation distribution.

Figure 5 presents raw transition probabilities in each cohort. Figure 6 presents comparative results for the control group (2001–2003) matched to the four subsequent cohorts, using matching weights to ensure that the characteristics of each cohort are the same. Appendix A presents the results with standard errors.

Figure 5 suggests that, over the sample period, about 75 per cent of young people remained within the same occupational quartile each year. This could reflect that quartiles are a relatively coarse measure of the distribution, or that job-to-job transitions were relatively infrequent over the period. However, using another dataset, Deutscher (2019, p. 5) estimated job switching rates of about 20 per cent per year for people in their early 20s in this period. In that context, our estimate of mobility is high: we estimate about a 25 per cent probability per year that a young person moves to an occupation in a different quartile. This suggests that quartiles are not too coarse to capture many job-to-job transitions.

We can assess how much mobility exists between quartiles, and whether that mobility has improved over time. The data suggest that mobility has not improved. We find that young people who graduated after 2007 were less likely to move out of the lowest quartile, and no more likely to move upward from the second quartile, when compared with graduates from 2001–2003. Concerningly, the lower quartiles did not show large upward movement as the economy recovered by 2017–2018. This could imply scarring for young graduates in 2010–2012 and 2013–2014. They may find themselves in

¹⁷ This is also seen in the regression results: the cohort effect for 2013–2015 graduates is statistically below the cohort effect for 2001–2003 (-5.88 with 95% confidence). However, the average marginal effect of graduating in 2013–2015 (which considers the average of sample so also summaries the subsequent trajectory of graduates' experience as well as the 'initial shift') is slightly smaller than the cohort effect and not statistically different from the 2001–2003 graduates (-4.67 with p-value = 0.17).

lower-scored occupations (which usually imply lower wage growth) over the long term.¹⁸

Figure 6 allows us to compare transition probabilities in each year with transitions in 2001–2003 for a matched cohort. The effects of the weak labour market after the GFC are noticeable. Relative to the period before the GFC, people aged 20–24 are now more likely to remain in the bottom quartile of the occupation distribution (2004–2006 and 2010–2014 graduates), and less likely to move to the second (2004–2006 and 2010–2014 graduates) or third quartile (2010–2014 graduates).¹⁹

Workers aged 25–34 are substantially more likely to remain in the second quartile (all graduates) (figure 6). They are also slightly more likely to move from the third quartile downwards (2004–2006 and 2010–2014 graduates); less likely to move from the second quartile to the third (2004–2012 graduates) or fourth (2007–2012 graduates); and less likely to move from the third to the fourth (2013–2014 graduates). Taken together, these results suggest that the recession not only made it more likely that young people would find a low-scored occupation, but more likely that they would remain in low-scored occupations.

Now we focus on transitions in the 2013–2014 cohort and compare them to earlier cohorts. The most striking feature of these Markov transitions is that not much improvement takes place for 20–24 year-olds in that last cohort. Transitions between quartiles look very similar or worse for 20–24 year-olds graduating in 2007–2009 and for 20–24 year-olds graduating in 2013–2014, even though the 2013–2014 graduates are transitioning during the period after 2015, in which the labour market showed signs of recovery.

For 25–34 year-olds, 2013–2014 graduates are less likely to remain in the second quartile than 2007–2009 graduates²⁰ — although both cohorts are significantly more likely to remain than their respective control group. However, they were more likely to transition downward from the third quartile and less likely to move upward, when compared with 2007–2009 graduates. And although there is a 15 percentage point decrease in the probability of remaining in the lowest quartile in 2013–2014 — many of those transitions are to the second quartile (figure 6).

¹⁸ Relatedly, Storesletten, Telmer and Yaron (2004) found that cohorts who live through more downturns can have larger dispersions in their labour income.

¹⁹ The cohorts reported in brackets are the cohorts for which the difference is statistically significant at the 95 per cent confidence level or better.

²⁰ In 2007–2009 the treatment group was 19 percentage points more likely to remain in quartile two, whereas in 2013–14 the treatment group was 13 percentage points more likely.

Over this whole period, transitions between quartiles are relatively uncommon, as one would expect; indeed the raw data suggests that average occupational score stabilises quickly after graduation (figure 5). But a decline in average occupational score for later cohorts and a bimodal distribution of occupation scores implies that there are more ‘unlucky’ workers than in earlier years. (The ‘unlucky’ are young workers whose occupational score is far lower than we would have predicted from their characteristics in earlier generations.) If employers are aware that a larger share of workers in lower-scored occupations are ‘unlucky’, they might be more willing to try out a worker who is currently in a lower-scored occupation, perhaps at a lower wage. So as the number of ‘unlucky’ young workers increases over time, one might observe some corresponding improvement in upward transitions.

Yet the data suggest that mobility has slightly worsened, and certainly there is no evidence of an improvement in mobility for 20-24 year-olds. And improvements for 25-34 year-olds moving from the second quartile may have been offset by worse transitions for those moving from the third quartile. This has serious implications: more young people looking for work after 2010 found themselves in lower-quartile jobs, either because they continued in a casual job they held during their studies, or because this was the best job they could find after graduation. There is little sign that they were able to move out of these lower quartiles in the following years. This may be because in a weak labour market, employers still face an excess supply of recent graduates and school leavers from which to choose, and perhaps some internal wage rigidities.

These results are corroborated by prior Markov analysis of the transitions between unemployment, being out of the labour force, part-time work, and full-time work (PC 2020a). We do not see any strong increase in the transitions from part-time work to full-time work after 2008. Again, these results suggest that unlucky young workers who fail to secure full-time work in a weak labour market do not find it easier to move out of part-time work.

5. Robustness

We undertake robustness checks to ensure that unobserved heterogeneity is not biasing the results. If there are differences in the young people who participate in labour markets over the sample period, the results will be biased. This is a strong possibility given that many young Australians live with their parents, and thus many of them have the option to not participate in labour markets if conditions are poor.

As described above, our Mincer regressions control for observed characteristics, and in the Markov analysis, matching techniques are used to compare cohorts with similar observed characteristics. Therefore our results will not be driven by differences in observed characteristics; but there remains the risk that unobserved differences will drive the results.

In order to address these concerns, we use a Rosenbaum sensitivity test to assess whether the data suggests the presence of unobserved heterogeneity between earlier and later cohorts.

Although the unconfoundedness assumption is untestable directly, Rosenbaum bounds are useful to indirectly assess whether matched samples are free from hidden bias that affects both assignment and the dependent variable (Rosenbaum 1987, 2002). They rely on the observation that the probability of treatment can only differ between two observations that have the same covariates if there is unobserved heterogeneity. Rosenbaum bounds (or procedures based on them) have been used in a variety of social research that implement matching techniques (Jordan 2012; Lamo and Messina 2010; Loughran et. al. 2015; Mavromaras et. al. 2009; McGuinness and Sloane 2011; Peel and Makepeace 2012).

The 2004-2006 graduate cohort is matched to the 2001–2003 cohort, using the matching technique described above. A dataset is created of the 2004-2006 cohort and the matched people from 2001-2003. Occupational scores after one year are regressed on the twenty-six demographic variables and on a dummy for belonging to the 2004-2006 cohort, which is the ‘treatment’. The average treatment effect on the treated is calculated, i.e. the decline in occupational score for the 2004-2006 cohort. The same exercise is undertaken for each cohort, and for occupational scores after four years, to assess whether cohorts caught up after four years.

A Rosenbaum sensitivity test calculates a sensitivity parameter, Γ , that measures the degree of departure from a randomised assignment to a treatment group. The value of gamma is interpreted as the tipping point beyond which we cannot reject that the results are driven by unobserved heterogeneity. For instance, if gamma is 1.1 then to attribute a higher occupational score to an unobserved covariate rather than to the treatment, that unobserved covariate would need to increase the odds ratio of treatment by a factor of 1.1. It tests how sensitive a test that rejects the null is to unobserved heterogeneity.²¹

The sensitivity tests (table B.5) show that the average treatment effect on the treated for the 2010–2012 graduate cohort was statistically significant and negative for the first year (at the 99 per cent level) and fourth year (at the 95 per cent level) after graduation. These support the finding that these cohorts either started in lower-scored

²¹ Rosenbaum bounds are a ‘worst-case’ scenario because they assume the unobserved variable affects both assignment and the outcome equally. In cases where hidden bias has a strong effect on assignment but a weak effect on the outcome, the confidence interval of the average treatment effect would still not include zero even in cases where Rosenbaum bounds suggest the results are sensitive (DiPrete and Gangl 2004).

occupations or were less likely to climb the occupational ladder. The average treatment effect on the treated for the 2013–2014 graduate cohort was statistically significant and negative for the first year (at the 95 per cent level) after graduation and negative but not statistically significant for the fourth year.

The Rosenbaum tests show that the statistically significant results have gamma values between 1.2 and 1.4, which means the odds ratio would need to be increased by 20–40 per cent before the null hypothesis (that there is no average treatment effect) cannot be rejected. The most robust result is the 2010–2012 cohort one year after graduation, which has a gamma value of 1.2 (for the p-values) and 1.4 (for the point estimate).

The full set of matching characteristics along with balance tests and Rosenbaum bounds are presented in appendix B. The Rosenbaum scores suggest that unobserved heterogeneity is unlikely to play a large role in our results, but they do not rule out hidden bias altogether.

Conclusion

The weak labour market from 2008 to 2018 was reflected primarily in young workers finding lower-scored occupations and earning lower wage rates than earlier generations (PC 2020a). Movement down the jobs ladder has negative implications for the lifetime earnings of those young workers. It also has implications for young people with low educational attainment, as they face competition from young people with more education. Some young people were pushed ‘off the ladder’, as indicated by the rise in long-run unemployment (PC 2020a).

Our results show that there is some potential for long-term effects from the weak labour market of 2008 to 2018. We show that young workers looking for work after the GFC experienced a dip in the occupational score of the jobs they secured. With Markov analysis, we are able to show that there was no large increase in upward movements; there is no evidence that young people who took lower-scored jobs found better jobs in the recovery. The absence of those upward transitions suggests that the generation who sought work during the 2008–2018 weak labour market faced worse long-term career prospects than their predecessors. Taken together, the bimodal distribution of occupational scores, the stability of the Markov transition matrices and the worsening of initial occupation outcomes suggest that poor initial outcomes can have long-term effects on young workers. In effect, young workers can be trapped in lower-score occupations.

What policies could address this challenge is a vexed question. While hiring credits are an effective policy to respond to scarring from unemployment, they may be less effective in addressing occupational scarring. Low-score occupations are likely to be more responsive to hiring credits than high-score occupations, unfortunately. To address occupational scarring, governments may need to understand the reason for an imbalance between demand and supply for high-score occupations in each country. The experience of Australia from 2008 to 2018 is likely to be a short-term imbalance due to the rapid expansion of the university sector (PC 2020a), and labour demand should adjust over

time. But in all contexts, the matching between demand and supply can be improved by providing better information to students about their job prospects according to the course and institution they select. The Productivity Commission has uncovered evidence that student course choice may reflect misinformation (PC 2019) and has recommended that more information be made available to assist students choosing a vocational education course (PC 2020b).

References

- ABS (Australian Bureau of Statistics) 2020, Labour Force, Australia, Nov 2019, Cat. no. 6202.0, Canberra.
- Altonji, J., Kahn, L. and Speer, J. 2016, 'Cashier or Consultant? Entry Labor Market Conditions, Field of Study, and Career Success', *Journal of Labor Economics*, vol. 34, no. S1, pp. S361–S401.
- Arulampalam, W. 2001, 'Is Unemployment Really Scarring? Effects of Unemployment Experiences on Wages', *The Economic Journal*, vol. 111, no. 475, pp. 585–606.
- , Gregg, P. and Gregory, M. 2001, 'Unemployment Scarring', *The Economic Journal*, vol. 111, no. 475, pp. 577–584.
- van den Berge, W. 2018, 'Bad Start, Bad Match? The Early Career Effects of Graduating in a Recession for Vocational and Academic Graduates.', *Labour Economics*, vol. 53, pp. 75–96.
- BLS (U.S. Bureau of Labour Statistics) 2020, Seasonally Adjusted Unemployment Rate (LNS14000000), Labor Force Statistics from the Current Population Survey, <https://data.bls.gov/timeseries/LNS14000000> (accessed 12 March 2020).
- Borland, J. 2020, Labour Market Snapshot #53: Is It Time to Say Goodbye to the Rate of Unemployment?, <https://sites.google.com/site/borlandjum/labour-market-snapshots> (accessed 29 January 2020).
- Buddelmeyer, H., Lee, W.-S. and Wooden, M. 2010, 'Low-Paid Employment and Unemployment Dynamics in Australia', *Economic Record*, vol. 86, no. 272, pp. 28–48.
- Chalmers, J. and Hill, T. 2007, 'Marginalising Women in the Labour Market: "Wage Scarring" Effects of Part-time Work', *Australian Bulletin of Labour*, vol. 33, no. 2, pp. 180–201.
- Coelli, M. and Wilkins, R. 2009, 'Credential Changes and Education Earnings Premia in Australia', *The Economic Record*, vol. 85, no. 270, pp. 239–259.
- Craig, B.A. and Sendi, P.P. 2002, 'Estimation of the Transition Matrix of a Discrete-time Markov Chain', *Health Economics*, vol. 11, no. 1, pp. 33–42.

- Deutscher, N. 2019, Job-to-Job Transitions and the Wages of Australian Workers, Treasury Working Paper, 2019–07, The Treasury, Canberra.
- Diamond, A. and Sekhon, J. 2013, ‘Genetic Matching for Estimating Causal Effects: A General Multivariate Matching Method for Achieving Balance in Observational Studies’, *The Review of Economics and Statistics*, vol. 95, no. 3, pp. 932–945.
- DiPrete, T.I. and Gangl, M. 2004, ‘Assessing Bias in the Estimation of Causal Effects: Rosenbaum Bounds on Matching Estimators and Instrumental Variables Estimation with Imperfect Instruments’, *Sociological Methodology*, vol. 34, pp. 271–310.
- Downes, P., Hanslow, K. and Tulip, P. 2014, The Effect of the Mining Boom on the Australian Economy, Research Discussion Paper, 2014–08, Reserve Bank of Australia, Canberra.
- Fok, Y.K., Scutella, R. and Wilkins, R. 2015, ‘The Low-Pay No-Pay Cycle: Are There Systematic Differences Across Demographic Groups?’, *Oxford Bulletin of Economics and Statistics*, vol. 77, no. 6, pp. 872–896.
- Forsythe, E. 2019, ‘Careers Within Firms: Occupational Mobility Over the Lifecycle’, *Labour*, vol. 33, no. 3, pp. 241–277.
- Fouarge, D. and Muffels, R. 2009, ‘Working Part-Time in the British, German and Dutch Labour Market: Scarring for the Wage Career?’, *Schmollers Jahrbuch: Journal of Applied Social Science Studies*, vol. 129, no. 2, pp. 217–226.
- Ganzeboom, H.B.G., P.M. de Graaf and D.J.Treiman (1992). ‘A Standard International Socio-economic Index of Occupational Status’, *Social Science Research* 21, pp. 1-56.
- Guvenen, F., Kuruscu, B., Tanaka, S. and Wiczer, D. 2020, ‘Multidimensional Skill Mismatch’, *American Economic Journal: Macroeconomics*, vol. 12, no. 1, pp. 210–244.
- Haldane, A. 2019, ‘Climbing the Jobs Ladder’, Speech, Bank of England, Glanford Park Stadium, Scunthorpe.
- Heckman, J. 1979, ‘Sample Selection Bias as a Specification Error’, *Econometrica*, vol. 47, no. 1, pp. 153–161.
- Ho, D.E., Imai, K., King, G. and Stuart, E.A. 2011, ‘MatchIt: Nonparametric Preprocessing for Parametric Causal Inference’, *Journal of Statistical Software*, vol. 42, no. 8, pp. 1–28.
- Jordan, K 2012, Juvenile Transfer and Recidivism: a Propensity Score Matching Approach, *Journal of Crime and Justice*, vol. 35, no. 1, pp. 53–67.
- Kahn, L. 2010, ‘The Long-Term Labor Market Consequences of Graduating from College in a Bad Economy’, *Labor Economics*, vol. 17, no. 2, pp. 303–316.

- Kalb, G. and Meekes, J. 2019, 'Wage Growth Distribution and Decline Among Individuals: 2001–2017', presented at Reserve Bank of Australia's Conference on Low Wage Growth, Sydney, 4 April.
- Keele, L.J. 2014, rbounds: Perform Rosenbaum Bounds Sensitivity Tests for Matched and Unmatched Data. R Package Version 2.1., <https://CRAN.R-project.org/package=rbounds> (accessed 7 January 2020).
- Knights, S., Harris, M.N. and Loundes, J. 2002, 'Dynamic Relationships in the Australian Labour Market: Heterogeneity and State Dependence', *Economic Record*, vol. 78, no. 242, pp. 284–298.
- Lamo, A. and Messina, J. 2010, 'Formal Education, Mismatch and Wages after Transition: Assessing the Impact of Unobserved Heterogeneity using Matching Estimators', *Economics of Education Review*, vol. 29, no. 6, pp. 1086–1099.
- Loughran, T. Wilson, T., Nagin, D. and Piquero, A. 2015, 'Evolutionary regression? Assessing the problem of hidden biases in criminal justice applications using propensity scores', *Journal of Experimental Criminology*, vol. 11, pp. 631–652.
- Mavromaras, K., McGuinness, S. and Fok, Y. 2009, 'Assessing the Incidence and Wage Effects of Overskilling in the Australian Labour Market', *Economic Record*, vol. 47, no. 23, pp. 2413–2429.
- Mavromaras, K., Sloane, P. and Wei, Z. 2015, 'The Scarring Effects of Unemployment, Low Pay and Skills Under-Utilisation in Australia Compared', *Applied Economics*, vol. 85, no. 268, pp. 60–72.
- McMillan, J., Beavis, A. and Jones, F.L. 2009, 'The Ausei06: A New Socioeconomic Index for Australia', *Journal of Sociology*, vol. 45, no. 2, pp. 123–149.
- McGuinness, S., Sloane, P. 2011, 'Labour Market Mismatch Among UK Graduates; An Analysis Using REFLEX Data', *Economics of Education Review*, vol. 30, no. 1, pp. 130–145.
- Mosthaf, A. 2014, 'Do Scarring Effects of Low-Wage Employment and Non-Employment Differ Between Levels of Qualification?', *Scottish Journal of Political Economy*, vol. 61, no. 2, pp. 154–177.
- Naidoo, D., Packard, T. and Auwalin, I. 2015, Mobility, Scarring and Job Quality in Indonesia's Labor Market, Policy Research Working Paper, No. 7484, World Bank, Washington.
- Oreopoulos, P., von Wachter, T. and Heisz, A. 2012, 'The Short- and Long-Term Career Effects of Graduating in a Recession', *American Economic Journal: Applied Economics*, vol. 4, no. 1, pp. 1–29.
- PC (Productivity Commission) 2019, The Demand Driven University System: A Mixed Report Card, Commission Research Paper, Canberra.

- Author Manuscript
- (Productivity Commission) 2020a, Why did young people's incomes decline?, Commission Research Paper, Canberra.
- (Productivity Commission) 2020b, National Agreement for Skills and Workforce Development Review. Commission Study Report, Canberra.
- Peel, M. and Makepeace, G 2012, 'Differential Audit Quality, Propensity Score Matching and Rosenbaum Bounds for Confounding Variables', *Journal of Business Finance and Accounting*, vol. 39, no. 5, pp. 606–648.
- Pennington, A. and Stanford, J. 2019, *The Future of Work for Australian Graduates: The Changing Landscape of University Employment Transitions in Australia*, The Australia Institute.
- Ralston, K., Feng, Z., Everington, D. and Dibben, C. 2016, 'Do Young People Not in Education, Employment or Training Experience Long-Term Occupational Scarring? A Longitudinal Analysis Over 20 Years of Follow-Up', *Journal of the Academy of Social Sciences*, vol. 11, no. 2–3, pp. 203–221.
- Rosenbaum, P.R. 1987, 'Sensitivity Analysis for Certain Permutation Inferences in MAtched Observational Studies', *Biometrika*, vol. 74, pp. 13–26.
- 2002, *Observational Studies*, Springer-Verlag.
- Sekhon, J. 2011, 'Multivariate and Propensity Score Matching Software with Automated Balance Optimization: The Matching Package for R', *Journal of Statistical Software*, vol. 42, no. 7, pp. 1–52.
- Sicherman, N. 1991, "'Overeducation" in the Labor Market', *Journal of Labor Economics*, vol. 9, no. 2, pp. 101–122.
- Stewart, A., Stanford, J. and Hardy, T. 2018, *The Wages Crisis in Australia. What It Is and What to Do About It*, The University of Adelaide Press.
- Storesletten, K., Telmer, C.I. and Yaron, A. 2004, 'Cyclical Dynamics in Idiosyncratic Labor Market Risk', *Journal of Political Economy*, vol. 112, no. 3, pp. 695–717.
- World Bank 2020, World Bank National Accounts Data, and OECD National Accounts Data Files, GDP growth (annual %) - Australia, <https://data.worldbank.org/indicator/NY.GDP.MKTP.KD.ZG?locations=AU> (accessed 12 March 2020).

Table 1 Summary statistics of pre-matched sample

<i>Variable</i>	<i>Grad cohort: 2001 - 2003</i>	<i>Grad cohort: 2004 - 2006</i>	<i>Grad cohort: 2007 - 2009</i>	<i>Grad cohort: 2010 - 2012</i>	<i>Grad cohort: 2013 - 2014</i>
	Mean (St. dev)	Mean (St. dev)	Mean (St. dev)	Mean (St. dev)	Mean (St. dev)
Age	34.66 (12.26)	34.46 (12.97)	35.04 (13.1)	33.82 (12.81)	32.36 (12.15)
Experience	13.96 (11.87)	13.64 (11.75)	14.11 (11.91)	12.97 (11.8)	11.83 (11.35)
Indigenous	0.03 (0.16)	0.02 (0.13)	0.02 (0.15)	0.03 (0.18)	0.04 (0.19)
Female	0.55 (0.5)	0.54 (0.5)	0.56 (0.5)	0.56 (0.5)	0.53 (0.5)
Married	0.54 (0.5)	0.50 (0.5)	0.56 (0.5)	0.52 (0.5)	0.55 (0.5)
Multiple jobs	0.11 (0.31)	0.08 (0.27)	0.09 (0.29)	0.09 (0.29)	0.09 (0.29)
Born in English speaking country	0.85 (0.36)	0.89 (0.31)	0.90 (0.3)	0.88 (0.32)	0.89 (0.32)
Regional unemployment	5.79 (0.93)	4.68 (1)	4.51 (1.01)	5.14 (0.67)	5.98 (0.83)
Proportion of life unemployed	0.01 (0.04)	0.02 (0.04)	0.02 (0.05)	0.02 (0.04)	0.02 (0.05)
Major city	0.66 (0.47)	0.65 (0.48)	0.65 (0.48)	0.61 (0.49)	0.68 (0.47)
Regional	0.32 (0.47)	0.34 (0.47)	0.33 (0.47)	0.38 (0.49)	0.31 (0.46)
Remote	0.02 (0.13)	0.01 (0.12)	0.02 (0.15)	0.01 (0.12)	0.01 (0.11)
Does not speak English well or at all	0.00 (0.05)	0.01 (0.09)	0.00 (0.05)	0.01 (0.08)	0.01 (0.08)
Speaks English well or very well	0.97 (0.16)	0.99 (0.09)	1.00 (0.05)	0.99 (0.09)	0.99 (0.08)
Diploma or Cert III or IV	0.39 (0.49)	0.44 (0.5)	0.45 (0.5)	0.52 (0.5)	0.53 (0.5)
University	0.32 (0.47)	0.27 (0.44)	0.28 (0.45)	0.26 (0.44)	0.27 (0.45)
Other qualification	0.28 (0.45)	0.29 (0.46)	0.27 (0.44)	0.21 (0.41)	0.20 (0.4)
Employed full-time	0.53 (0.5)	0.49 (0.5)	0.55 (0.5)	0.48 (0.5)	0.45 (0.5)
Employed part-time	0.25 (0.44)	0.29 (0.46)	0.25 (0.43)	0.29 (0.46)	0.30 (0.46)
N	695	941	888	1292	1238

Means of important control variables for each graduation cohort. Standard deviations are reported below each mean in brackets.


Table 2 Selected ANZSCO occupations and their quartile

<i>ANZSCO occupation sub-group description</i>	<i>AUSEI06 score</i>	<i>20-24 quartile</i>	<i>25-34 quartile</i>
Farm, Forestry and Garden Workers	11.0	Q1	Q1
Factory Process Workers	12.1	Q1	Q1
Mobile Plant Operators	14.7	Q1	Q1
Cleaners and Laundry Workers	20.4	Q1	Q1
Store persons	20.8	Q1	Q1
Food Trades Workers	21.2	Q1	Q1
Food Preparation Assistants	22.0	Q1	Q1
Construction and Mining Labourers	23.3	Q1	Q1
Farmers and Farm Managers	34.0	Q2	Q1
Hospitality Workers	34.7	Q2	Q1
Carers and Aides	34.9	Q2	Q1
Construction Trades Workers	36.4	Q2	Q1
Inquiry Clerks and Receptionists	37.3	Q2	Q2
Clerical and Office Support Workers	37.7	Q2	Q2
General Clerical Workers	41.9	Q2	Q2
Personal Assistants and Secretaries	44.8	Q2	Q2
Hospitality, Retail and Service Managers	45.4	Q2	Q2
Other Clerical and Administrative Workers	47.1	Q2	Q2
Protective Service Workers	47.8	Q2	Q2
Numerical Clerks	48.8	Q2	Q2
Sports and Personal Service Workers	49.8	Q2	Q2
Sales Representatives and Agents	50.7	Q2	Q2
Office Managers and Program Administrators	57.4	Q3	Q2
Engineering, ICT and Science Technicians	57.7	Q3	Q2
Health and Welfare Support Workers	60.0	Q3	Q2
Arts and Media Professionals	68.1	Q3	Q3
Specialist Managers	71.4	Q3	Q3
Business, Human Resource and Marketing Professionals	77.3	Q3	Q3
Chief Executives, General Managers and Legislators	78.2	Q3	Q3
Design, Engineering, Science and Transport Professionals	81.2	Q3	Q4
ICT Professionals	81.3	Q3	Q4
Legal, Social and Welfare Professionals	84.4	Q4	Q4
Education Professionals	84.9	Q4	Q4
Health Professionals	85.1	Q4	Q4

ANSIC occupational categories with their associated AUSEI06 score along with their estimated quartile for each age group. For example farmers and farm managers are in Q2 for 20-24 year olds but Q1 for 25-34 year olds.

Figure 1 **Workers aged 25-54 are in higher-scored occupations in 2018**


Empirical probability density functions (PDFs) of occupational scores



Age-based estimated distributions of the AUSIE06 score in 2001 and in 2018. The distribution is weighted using responding person weights in the HILDA dataset.


Figure 2 **Graduates under 35 are more likely to be in lower-scored occupations in 2018**

Empirical probability density functions (PDF) of occupational scores



Age-based estimated distributions of the AUSIE06 score in 2001 and in 2018. The distribution is weighted using responding person weights in the HILDA dataset. Each age group has been separated into groups based on educational attainment. The 'bachelor and above' group includes graduates with a bachelor and post-graduate qualifications. The other group includes all other post-school qualifications, including VET certificates.

Figure 3 Occupational scores for graduates have not improved
Average occupational score, first five years after graduation, 2001–2018



The average occupational score of graduates for each graduation cohort over time. The occupation of the graduate in their graduation year (year 0) is included as the first data point. Each observation is weighted by their responding person weight from HILDA.

Table 3 Mincer regressions: Equations (1) and (2)

Coefficient	Dependent variable: occupational score	
	(1)	(2)
	β (s.e)	β (s.e)
Years since graduation	3.350 *** (1.258)	

Years since graduation Squared	-0.656 **	
	3.350 ***	
Graduation cohort: 2004 - 2006		0.054
		(2.194)
Graduation cohort: 2007 - 2009		-0.743
		(3.718)
Graduation cohort: 2010 - 2012		-1.430
		(2.975)
Graduation cohort: 2013 - 2015		-5.881 **
		(2.683)
Years since graduation x Graduation cohort:		
2001 - 2003		2.651 *
		(1.610)
2004 - 2006		2.710
		(2.149)
2007 - 2009		3.152 ***
		(1.017)

2010 - 2012	1.485 *
	(0.883)
2013 - 2015	4.675 ***
	(1.642)
Years since graduation Sq. x Graduation cohort:	
2001 - 2003	-0.490 *
	(0.272)
2004 - 2006	-0.769
	(0.513)
2007 - 2009	-0.743 **
	(0.370)
2010 – 2012	-0.276
	(0.225)
2013 – 2015	-0.825 **
	(0.343)
# observations	14311
	14311


R squared	0.323	0.324
F statistic	161.947	126.633

OLS regression results for regressions (1) and (2). Regression (1) includes mincer-style specification where experience is modelled via years since graduation and years since graduation squared variables. Regression (2) extends this regression by including cohort effects and cohort interaction effects. Standard errors are reported below coefficient estimates and are clustered at the 1-digit ANZSCO occupation level.
 * p = 0.1, ** p = 0.05, *** p = 0.01.


Figure 4 Longer-term cohort effects do not show improvement

Plot of the quadratic function in experience^a from regression (2)

(a) Shorter-term cohort effects (0 to 4 years after graduation)



(b) Longer-term cohort effects (0 to 4 years after graduation)



^a The marginal effect of years since graduation is plotted for each cohort for their first 5 years of graduation. That is, $c_g \hat{\alpha}_2 + e_i c_g \hat{\alpha}_3 + e_i^2 c_g \hat{\alpha}_4$ is calculated, where years since graduation is set to 0, 1, 2, 3, 4.

Table 4 IV regression main results: equations (3) and sensitivity tests


Coefficient	Dependent variable: occupational score				
	OLS	OLS	IV (u.e. 21)	IV (u.e. 22)	IV (u.e. 23)
	(S1)	(S2)	(S3)	(3)	(S4)
	β (s.e)	β (s.e)	β (s.e)	β (s.e)	β (s.e)
Years since graduation	3.93*** (1.42)		-15.84 (13.63)	-14.52 (12.8)	-21.18* (11.83)
Years since graduation sq.	-0.61*** (0.2)		11.31** (5.28)	10.18** (5)	9.97** (4.45)

Regional unemployment in graduation year	-1.20*		-7.46	-9.31*	-9.57*
	(0.62)		(5.3)	(5.36)	(5.11)
Years since grad * Regional ue in grad year	0.01		-2.51	-1.99	-0.52
	(0.24)		(3.66)	(3.41)	(3.03)
Years since 22		2.53***			
		(0.76)			
Years since 22 sq.		-0.17**			
		(0.07)			
Regional unemployment when 22		-3.13			
		(2.24)			
Age * Regional ue when 22		0.11			
		(0.08)			
N	7836	7831	7807	7831	7836
F-statistic	52.11	69.34	26.4	28.41	29.53
Degrees of freedom	50	45	42	42	42

Instrumental variable regression results for regressions (3) and some sensitivity tests (S1 to S4). Regression (3) includes mincer-style specification where experience is modelled via years since graduation and years since graduation squared variables and extends regression (2) (table 3). We instrument our experience variable — years since graduation — with years since reaching the age of 22

(with sensitivity tests for ages 21 and 23). We also instrument the regional unemployment rate at year of graduation with the regional unemployment rate at the age of 22 of the region the graduate lived in when they completed secondary school. Standard errors are reported below coefficient estimates. are clustered at the 1-digit ANZSCO occupation level. * $p = 0.1$, ** $p = 0.05$, *** $p = 0.01$.

Figure 5 Young people are likely to stay in the same occupation quartile
Probability of transition between occupation quartiles each year^a, unmatched sample




Baseline Markov probability of transition between occupation quartiles each year for the unmatched samples. Occupational transitions are observed for the four years after graduation — that is, someone who graduated in 2007 is observed between 2008 and 2011. These transitions are used to calculate a probability transition matrix, $\hat{M}$, where the probability of transitioning from state i to state j is the number of transitions from state i to state j divided by the number of people originally in state i (Craig and Senti

2002). **a** The occupation quartile in each year of the four years after graduation was recorded and then year to year transitions were pooled across the four-year period. **b** Transitions weighted by HILDA's representative person weights.

Figure 6 **Young people are more likely to stay in occupations in the two bottom quartiles of occupational score**

Changes in the probability of transition between occupation quartiles each year^a, comparisons between treatment and matched control groups



Changes in the probability of transition between occupation quartiles each year, comparisons between treatment and matched control groups. Once each individual i in the treatment cohort has been matched to an individual k in the 2001-2003 cohort, we have the matching weights. For example, if person k was matched ω_k times, then every transition person k makes counts as ω_k transitions, not one transition. We calculate a weighted transition matrix for each treatment group as well as for each control group, which we then use to calculate a matrix of transition probabilities, $\bar{M}$, for each treatment and control group. The weighted observed transitions will carry across so that the transitions probabilities are also weighted by

the matching weights. The final step is to take the difference of the probability matrices of each treatment group and their respective control group: $\Delta \hat{M} = \hat{M}_T - \hat{M}_C$.

Appendix A

Table A.1 Markov transition matrices: workers aged 20–24

Row: quartile from	Q1	Q2	Q3	Q4
Column: quartile to				
	<i>p</i> (<i>s.e</i>)	<i>p</i> (<i>s.e</i>)	<i>p</i> (<i>s.e</i>)	<i>p</i> (<i>s.e</i>)
	Graduated from 2004 to 2006			
Q1	0.625*** (0.044)	0.158*** (0.032)	0.144*** (0.032)	0.073*** (0.023)
Q2	0.087*** (0.027)	0.746*** (0.042)	0.089*** (0.027)	0.078*** (0.025)
Q3	0.002 (0.006)	0.157*** (0.041)	0.795*** (0.046)	0.046** (0.023)
Q4	0.000	0.039	0	0.961***

	(0.000)	(0.029)	(0)	(0.029)
	Graduated from 2007 to 2009			
Q1	0.641***	0.204***	0.112***	0.043**
	(0.043)	(0.037)	(0.028)	(0.019)
Q2	0.127***	0.692***	0.104***	0.077***
	(0.028)	(0.038)	(0.024)	(0.023)
Q3	0.011	0.169***	0.75***	0.07***
	(0.011)	(0.038)	(0.044)	(0.026)
Q4	0	0.04	0.099***	0.861***
	(0)	(0.025)	(0.037)	(0.044)
	Graduated from 2010 to 2012			
Q1	0.7***	0.176***	0.067***	0.056***
	(0.032)	(0.026)	(0.017)	(0.015)
Q2	0.149***	0.663***	0.142***	0.045***
	(0.026)	(0.034)	(0.025)	(0.015)
Q3	0.092***	0.164***	0.635***	0.109***
	(0.024)	(0.031)	(0.038)	(0.026)

Q4	0.005 (0.008)	0.035* (0.019)	0.082*** (0.029)	0.878*** (0.034)
Graduated from 2013 to 2014				
Q1	0.718*** (0.029)	0.135*** (0.022)	0.075*** (0.017)	0.072*** (0.017)
Q2	0.22*** (0.029)	0.563*** (0.036)	0.174*** (0.027)	0.043*** (0.014)
Q3	0.064*** (0.021)	0.154*** (0.03)	0.61*** (0.042)	0.172*** (0.033)
Q4	0.009 (0.011)	0.049* (0.025)	0.115*** (0.036)	0.828*** (0.043)

Baseline Markov probability of transition between occupation quartiles each year for the unmatched samples. Occupational transitions are observed for the four years after graduation — that is, someone who graduated in 2007 is observed between 2008 and 2011. These transitions are used to calculate a probability transition matrix, $\hat{M}$, where the probability of transitioning from state i to state j is the number of transitions from state i to state j divided by the number of people originally in state i (Craig and Sendi 2002). **a** *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. **b** s.e calculated using bootstrap methodology.

Table A.2 Markov transition matrices: workers aged 25–34

Row: quartile from	Q1	Q2	Q3	Q4
Column: quartile to				
	<i>p</i> (s.e)	<i>p</i> (s.e)	<i>p</i> (s.e)	<i>p</i> (s.e)
Graduated from 2004 to 2006				
Q1	0.751*** (0.033)	0.156*** (0.029)	0.047*** (0.018)	0.047*** (0.017)
Q2	0.109*** (0.027)	0.684*** (0.04)	0.151*** (0.03)	0.055*** (0.019)
Q3	0.108*** (0.031)	0.085*** (0.028)	0.714*** (0.045)	0.093*** (0.03)
Q4	0.027** (0.013)	0.029** (0.014)	0.084*** (0.023)	0.86*** (0.028)
Graduated from 2007 to 2009				
Q1	0.779*** (0.038)	0.137*** (0.031)	0.013 (0.026)	0.07** (0.036)
Q2	0.091*** (0.024)	0.771*** (0.035)	0.09*** (0.029)	0.048* (0.026)
Q3	0.022 (0.015)	0.102*** (0.03)	0.771*** (0.279)	0.105 (0.278)
Q4	0 (0)	0.074*** (0.021)	0.103 (0.301)	0.824*** (0.301)
Graduated from 2010 to 2012				
Q1	0.833*** (0.026)	0.109*** (0.022)	0.051*** (0.015)	0.008 (0.006)
Q2	0.059*** (0.015)	0.796*** (0.026)	0.071*** (0.016)	0.074*** (0.017)
Q3	0.034** (0.015)	0.114*** (0.026)	0.718*** (0.038)	0.135*** (0.028)
Q4	0.002 (0.003)	0.029*** (0.011)	0.068*** (0.016)	0.901*** (0.02)
Graduated from 2013 to 2014				
Q1	0.674*** (0.029)	0.232*** (0.026)	0.061*** (0.014)	0.033*** (0.012)
Q2	0.063*** (0.014)	0.728*** (0.024)	0.129*** (0.017)	0.08*** (0.014)
Q3	0.165*** (0.031)	0.161*** (0.031)	0.623*** (0.041)	0.051*** (0.018)
Q4	0.009 (0.006)	0.059*** (0.014)	0.032*** (0.011)	0.9*** (0.018)

Baseline Markov probability of transition between occupation quartiles each year for the unmatched samples. Occupational transitions are observed for the four years after graduation — that is, someone who graduated in 2007 is observed between 2008 and 2011. These transitions are used to calculate a probability transition matrix, $\hat{M}$, where the probability of transitioning from state i to state j is the number of transitions from state i to state j divided by the number of people originally in state i (Craig and Sendi 2002). **a** *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. **b** s.e calculated using bootstrap methodology.

Table A.3 **Markov transition matrices, change from 2001–2003: workers aged 20–24**

Row: quartile from Column: quartile to	Q1	Q2	Q3	Q4
	p (s.e)	p (s.e)	p (s.e)	p (s.e)
Graduated from 2004 to 2006				
Q1	0.08*** (0.009)	-0.159*** (0.015)	0.027 (0.021)	0.052*** (0.017)
Q2	-0.102*** (0.017)	0.121*** (0.011)	-0.085*** (0.021)	0.066** (0.026)
Q3	0.002 (NaN)	-0.102*** (0.03)	0.107*** (0.015)	-0.007 (0.017)
Q4	0 (NaN)	-0.185*** (0.033)	-0.078 (0.071)	0.262*** (0.016)
Graduated from 2007 to 2009				
Q1	0.009 (0.009)	-0.03* (0.016)	0.045** (0.019)	-0.024 (0.017)
Q2	-0.034** (0.015)	-0.011 (0.009)	0.004 (0.014)	0.041** (0.017)
Q3	-0.013 (NaN)	0.02 (0.023)	0.024 (0.016)	-0.031 (0.031)
Q4	0 (NaN)	0.04 (0.031)	0.013 (0.052)	-0.053*** (0.019)
Graduated from 2010 to 2012				
Q1	0.091*** (0.007)	-0.076*** (0.012)	-0.025** (0.012)	0.01 (0.013)
Q2	0.054*** (0.011)	-0.073*** (0.008)	0.036** (0.016)	-0.017 (0.013)
Q3	0.028 (0.02)	-0.12*** (0.021)	0.038*** (0.013)	0.053*** (0.019)
Q4	0.005 (NaN)	-0.182*** (0.022)	0.04 (0.031)	0.138*** (0.013)
Graduated from 2013 to 2014				
Q1	0.138*** (0.006)	-0.094*** (0.01)	-0.104*** (0.014)	0.06*** (0.012)
Q2	0.062*** (0.012)	0.036*** (0.009)	-0.108*** (0.015)	0.01 (0.012)
Q3	0.03 (0.022)	-0.08*** (0.017)	-0.033*** (0.012)	0.083*** (0.016)
Q4	0.009 (NaN)	-0.115*** (0.031)	-0.037 (0.032)	0.144*** (0.013)

Transitions calculated using methodology in Craig and Sendi (2002). We calculate a weighted transition matrix for each treatment group as well as for each control group,

which we then use to calculate a matrix of transition probabilities, $\hat{M}$, for each treatment and control group. The weighted observed transitions will carry across so that the transitions probabilities are also weighted by the matching weights. The difference of the probability matrices of each treatment group and their respective control group is: $\Delta\hat{M} = \hat{M}_T - \hat{M}_C$. ^a *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. ^b s.e. calculated based on pooled variance of treatment and control transition matrices. ^c Significance based on t-test of difference between treated and control transition matrices. ^d NaN in s.e. arise when there are no observed transitions in either the treatment or control.

Table A.4 Markov transition matrices, change from 2001–2003: workers aged 25–34

Row: quartile from	Q1	Q2	Q3	Q4
Column: quartile to				
	<i>p</i> (<i>s.e</i>)	<i>p</i> (<i>s.e</i>)	<i>p</i> (<i>s.e</i>)	<i>p</i> (<i>s.e</i>)
Graduated from 2004 to 2006				
Q1	-0.13*** (0.007)	0.056*** (0.015)	0.039** (0.02)	0.035* (0.018)
Q2	0.022 (0.015)	0.129*** (0.01)	-0.151*** (0.031)	0 (0.039)
Q3	0.066*** (0.025)	0.003 (0.022)	-0.067* (0.038)	-0.003 (0.094)
Q4	0.027* (0.015)	-0.051*** (0.022)	-0.002 (0.038)	0.026 (0.094)

	(0.016)	(0.015)	(0.11)	(0.029)
	Graduated from 2007 to 2009			
Q1	-0.076***	0.035*	0.009	0.032
	(0.008)	(0.019)	(0.03)	(0.038)
Q2	-0.013	0.19***	-0.152***	-0.026
	(0.013)	(0.009)	(0.029)	(0.032)
Q3	0.005	0.009	-0.038	0.024
	(0.017)	(0.018)	(0.06)	(0.151)
Q4	0	-0.009	-0.025	0.035
	(NaN)	(0.021)	(0.143)	(0.048)
	Graduated from 2010 to 2012			
Q1	0.08***	-0.062***	0.048***	-0.066***
	(0.008)	(0.014)	(0.016)	(0.007)
Q2	-0.154***	0.239***	-0.062***	-0.023*
	(0.013)	(0.006)	(0.011)	(0.012)
Q3	0.022	0.063***	-0.067***	-0.018
	(0.016)	(0.015)	(0.024)	(0.053)


Q4	0.002	-0.08***	-0.051	0.13***
	(NaN)	(0.017)	(0.057)	(0.017)
Graduated from 2013 to 2014				
Q1	-0.176***	0.109***	0.052***	0.015
	(0.006)	(0.012)	(0.015)	(0.012)
Q2	-0.127***	0.131***	-0.005	0.001
	(0.01)	(0.005)	(0.01)	(0.012)
Q3	0.14***	0.071***	-0.022	-0.189***
	(0.023)	(0.015)	(0.021)	(0.066)
Q4	0.009	-0.018*	-0.155**	0.164***
	(0.007)	(0.01)	(0.075)	(0.018)

Transitions calculated using methodology in Craig and Sendi (2002). We calculate a weighted transition matrix for each treatment group as well as for each control group, which we then use to calculate a matrix of transition probabilities, $\hat{M}$, for each group. The weighted observed transitions will carry across so that the transitions probabilities are also weighted by the matching weights. The difference of the probability matrices of each treatment group and their respective control group is: $\Delta\hat{M} = \hat{M}_T - \hat{M}_C$. ^a *** p < 0.01; ** p < 0.05; * p < 0.1. ^b s.e calculated based on pooled variance of treatment and control transition matrices. ^c Significance based on t-test of difference between treated and control transition matrices. ^d NaN in s.e. arise when there are no observed transitions in either the treatment or control.

Appendix B

Figure B.1 2004–2006 graduates


Absolute standardised means before and after genetic matching procedure.



a Cut-off value = 0.1 (dotted line). The closer to zero, the better the balance of the sample is on these observables. Distance measures the overall balance of the sample.

Figure B.2 **2007–2009 graduates**


Absolute standardised means before and after genetic matching procedure.



a Cut-off value = 0.1 (dotted line). The closer to zero, the better the balance of the sample is on these observables. Distance measures the overall balance of the sample.

Figure B.3 **2010–2012 graduates**


Absolute standardised means before and after genetic matching procedure.



a Cut-off value = 0.1 (dotted line). The closer to zero, the better the balance of the sample is on these observables. Distance measures the overall balance of the sample.

Figure B.4 **2013–2015 graduates**

Absolute standardised means before and after genetic matching procedure.



a Cut-off value = 0.1 (dotted line). The closer to zero, the better the balance of the sample is on these observables. Distance measures the overall balance of the sample.

Table B.5 Sensitivity analysis of matching results: Rosenbaum bounds

Average treatment effect on the treated and Rosenbaum bounds for the matched sample of each graduate cohort.

<i>Cohort and years after graduation</i>	<i>Regression results (average treatment effect on the treated)</i>	<i>Rosenbaum sensitivity test for Wilcoxon Signed Rank P-Value</i>	<i>Rosenbaum Sensitivity Test for Hodges-Lehmann Point Estimate</i>
	β (s.e)	γ [c.i interval]	γ [c.i interval]
2004 – 2006			
1 year after	-0.96 (1.22)	1.0 [0.15, 0.15]	1.2 [-3, 0.9]
4 years after	-0.4 (1.57)	1.0 [0.38, 0.38]	1.1 [-2, 0.6]
2007 – 2007			
1 year after	-0.84 (1.37)	1.0 [0.16, 0.16]	1.2 [-3, 0.5]
4 years after	0.89 (1.49)	1.0 [0.34, 0.34]	1.1 [0, 1.55]
2010 – 2012			
1 year after	-3.56*** (1.11)	1.2 [0, 0.14]	1.4 [-7, 0.55]
4 years after	-2.59** (1.22)	1.2 [0, 0.29]	1.3 [-6, 0.2]
2013 – 2015			
1 year after	-2.73** (1.16)	1.2 [0, 0.33]	1.3 [-6, 0.45]
4 years after	-0.96 (1.22)	1.0 [0.15, 0.15]	1.2 [-3, 0.9]

A Rosenbaum sensitivity test calculates a sensitivity parameter, Γ , that measures the degree of departure from a randomised assignment to a treatment group. The value of gamma is interpreted as the tipping point beyond which we cannot reject that the results are driven by unobserved heterogeneity. For instance, if gamma is 1.1 then to attribute a higher occupational score to an unobserved covariate rather than to the treatment, that unobserved covariate would need to increase the odds ratio of treatment by a factor of 1.1.