The Short Run Effects of Age Based Youth Minimum Wages in Australia: A Regression Discontinuity Approach

Alex Olssen

Abstract

This paper uses based youth award rates (minimum wages) in Australia to identify the short-term effect of minimum wages on both actual wages and youth employment. Youth minimum wages are typically based on adult minimum wages, starting at 50 per cent for 16 year olds and increasing by 10 per cent per year until the full adult rate at 21 years of age. I find some evidence to suggest a 10 per cent increase in youth minimum wages results in an increase in actual wages of around 6 per cent. I find no evidence that increasing youth minimum wages affects youth employment hours.

Disclaimer

This paper uses unit record data from the Household, Income and Labour Dynamics in Australia (HILDA) Survey. The HILDA Project was initiated and is funded by the Australian Government Department of Families, Housing, Community Services and Indigenous Affairs (FaHCSIA) and is managed by the Melbourne Institute of Applied Economic and Social Research (Melbourne Institute). The findings and views reported in this paper, however, are those of the author and should not be attributed to either FaHCSIA or the Melbourne Institute.

1. Introduction

The employment effects of minimum wages are important for policy. Policy makers face a trade-off between increasing the pay of low wage workers and potentially reducing their employment opportunities. Despite this more work is still needed to evaluate direction and the magnitude of minimum wage effects on employment. Of the 102 studies covered in a comprehensive review by Neumark and Wascher (2006) around a thrid fail to find negative employment effects of minimum wages – however the authors do state that they find a higher proportion of the most credible studies report negative employment effects. Furthermore there are relatively few studies estimating the employment effects of minimum wages in Australia (most recently see Leigh (2003, 2004) for estimates of Western Australia). This research aims to estimate the causal impact of minimum wages on employment for youths in Australia. In doing so it contributes more evidence to the effect of minimum wages on youth labour market outcomes. It is also the first attempt, internationally, that I am aware of, to use the regression discontinuity design to estimate the short run causal effects of minimum wages – on both hours of employment and actual received wages.[1]

Australia has a unique policy position on minimum wages – they vary by state, occupation, and importantly for my research, by age. Most people have their wages set by an award scheme. These award schemes provide the minimum pay that people receive depending on their occupation and the state that they work in. Furthermore for youths – workers under the age of 21 – there is a sliding scale that increases with age. Award minimum wages typically begin at about 50 per cent of the adult wage for people younger than 15 or 16 and then increase by about 10 per cent per year until the age of 21 at which point the adult award minimum wage applies. For example, the 2009 pay-scale for retail staff in South Australia sets the minimum wage for a normal 17 year old worker at $8.27, which is 50 per cent of the minimum rate for an adult worker in the same job. For normal 20 year old retail staff the minimum wage is $14.88, which is 90 per cent of the adult rate. For ticket machine operators the rates were $8.50 to $14.48 for the same age groups. Both groups have minimum wage increases of 10 per cent of the adult rate per year.

Consider two groups of people, those just older than 19 and those just younger than 19. For a small difference in age we don't expect there to be much difference in the averages of those underlying characteristics that are important for employment. Stated differently, absent differences in minimum wages we would expect these two groups to have similar levels of employment. Policy means that award minimum wages increase discontinuously at birthdays by about 10 per cent. I aim to use this discontinuity in minimum wages to estimate

[1] Lee and Lemieux (2009) is a comprehensive review of regression discontinuity. They have a table of many of the applications of regression discontinuity of which they are aware.

the average causal effects of increasing minimum wages for Australian youths. I use data from the longitudinal household survey HILDA, which includes around 14,000 individuals from 2000 to 2009.

I use a regression discontinuity approach. Regression discontinuity has had more attention from the economic research community recently and there have been some notable theoretical contributions. Hanh, Todd, and van der Klaauw (2001) prove that very weak requirements are needed for identification of a causal average treatment affect.

The rest of the paper is structured as follows. Section 2 looks at the theory of regression discontinuity and how it can be applied to this research. Section 3 provides further background on Australian institutions and looks at what regression discontinuity can hope to tell us in this case. Section 4 introduces the data, provides some descriptive statistics and presents my research design. In section 5 I present graphical analyses on the effect of minimum wages on actual wages and employment. Econometric technique and results are reported in section 6. Section 7 presents some graphs to evaluate the credibility of my regression discontinuity design. Section 8 concludes.

2. Regression discontinuity Literature

This section reviews the theoretical regression discontinuity literature. It looks to develop a basic understanding of the regression discontinuity design so I can implement it in my research. It does this by considering the major theoretical contributions to econometrics of Regression Discontinuity design and discussing how the findings of these papers relate to my research. It breaks the discussion into three parts. Firstly I discuss identification and interpretation. Secondly, I discuss estimation, and finally I discuss validity testing. Before I begin I would like to point out that Lee and Lemieux (2009) provide a review of regression discontinuity. Furthermore Imbens and Lemieux (2008) is a very useful practical guide for regression discontinuity.

Regression discontinuity has been gaining attention in economic research recently. The 142nd volume of the Journal of Econometrics was devoted entirely to regression discontinuity research. Furthermore Imbens and Wooldridge devoted a lecture in their "What's New in Econometrics" lecture series for the NBER to regression discontinuity. Two major factors explain the attractiveness of the Regression Discontinuity design for economists. Firstly, under certain conditions, standard Regression Discontinuity designs have been shown to be "as good as random". Secondly, many policies treat people differently based on discrete assignment rules that make use of observed underlying characteristics – for example means based income benefits. This is exactly the situation that regression discontinuity designs can deal with.

2.1. Identification in Regression Discontinuity

It will be useful to set up a framework to discuss Regression Discontinuity. Suppose individuals are assigned to treatment or not based upon whether an observed covariate passes a known threshold – in this research assignment is based on whether a youth has recently had a birthday. Let Z_i represent the assignment variable and let z_0 represent the cut-off so that person i is treated if $Z_i \geq z_0$. Let X_i be an indicator variable that is 1 for people who are treated and 0 for people who are not. We observe Y_{i1} for treated individuals and Y_{i0} for control group individuals. We would like to know $Y_{i1} - Y_{i0}$, however, we only observe either Y_{i1} or Y_{i0}, not both. Thus we focus on average treatment effects. In this research HILDA participants are interviewed yearly so that we observe their wage and employment hours either before or after they have had their birthday. Note that we can write the observed outcome Y_i as $Y_i = \alpha + X_i\beta_i$, where $\alpha = Y_{i0}$ and $\beta_i = Y_{i1} - Y_{i0}$.

Regression discontinuity designs come in two flavours – sharp and fuzzy. In a sharp design everyone whose assignment variable passes some cut-off is treated while everyone whose assignment variable fails to pass the cut-off is not treated. Thus the probability that an individual is treated is 0 for individuals who do not pass the threshold and 1 for individuals who do pass it. A fuzzy design is similar, however, it only requires a discontinuous increase in the probability of treatment at the threshold – i.e., it allows for smaller jumps in the probability of treatment.

Hahn et al. (2001) establish the mild conditions necessary to successfully identify an Average Treatment Effect in a regression discontinuity framework. They focus on the fuzzy design and demonstrate the sharp design is a special case. I reproduce their proof which establishes the necessary conditions for identification with heterogeneous treatment effects. I fill in some steps which were helpful in my own understanding of the theorem. This case does not allow individuals to self-select into the program on the basis of prospective gains from treatment. This is not a problem for my research as I do not anticipate people can self-select into treatment as this would require them to change their age.

Firstly I state the necessary assumptions for identification in a regression discontinuity setting and explain them.

Assumption (RD): (i) The limits $x^+ = \lim_{z \to z_0^+} E[X_i \mid Z_i = z]$ and $x^- = \lim_{z \to z_0^-} E[X_i \mid Z_i = z]$. (ii) $x^+ \neq x^-$.

This is just a mathematical formalisation of the basic requirement I know I need for regression discontinuity to work, namely, that there is some discontinuity in treatment!

Assumption (A1): $E[\alpha_i \mid Z_i = z]$ is continuous at $Z = z_0$.

Assumption (A2): $E[\beta_i \mid Z_i = z]$ as a function of Z is continuous at z_0.

These two assumptions are necessary conditions for identification. They are motivated as follows. Suppose that in the absence of treatment individuals near the cut-off are similar. Then the expected outcome either side of the threshhold should be similar if either nobody is treated or if everybody is treated. Importantly the expected outcome of the control group represents the valid counterfactual for the treatment group. It will be important later to recall the definition of continuity – a function $f(x)$ is continuous at x_0 if $\lim_{x \to x_0^-} f(x) = \lim_{x \to x_0^+} f(x) = f(x_0)$ where equality implies existence.

We are now ready to reproduce the result in Hahn et al. (2001). The following theorem shows that if we have a discontinuity in treatment (Assumption (RD)) and if the expected outcome of the control group is the valid counterfactual for the treatment group (Assumption (A1) and Assumption (A2)) then we can identify the treatment effect.

Theorem: Suppose that X_i is independent of β_i conditional on Z_i near z_0. Further suppose Assumptions (RD), (A1), and (A2) hold. We have

$$E[\beta_i \mid Z_i = z_0] = \frac{y^+ - y^-}{x^+ - x^-}$$

Where $y^+ = \lim_{z \to z_0^+} E[Y_i \mid Z_i = z]$ and $y^- = \lim_{z \to z_0^-} E[Y_i \mid Z_i = z]$

Proof: Let ε be some arbitrarily small positive real number. The mean difference in outcome for people above and below the discontinuity point is

$$E[Y_i \mid Z_i = z_0 + \varepsilon] - E[Y_i \mid Z_i = z_0 - \varepsilon]$$

which, using the definition of Y_i can be rewritten as

$$E[\alpha_i + X_i\beta_i \mid Z_i = z_0 + \varepsilon] - E[\alpha_i + X_i\beta_i \mid Z_i = z_0 - \varepsilon]$$

By the linearity of the expectation operator this is just

$$\{E[X_i\beta_i \mid Z_i = z_0 + \varepsilon] - E[X_i\beta_i \mid Z_i = z_0 - \varepsilon]\}$$
$$+ E[\alpha_i \mid Z_i = z_0 + \varepsilon] - E[\alpha_i \mid Z_i = z_0 - \varepsilon]$$

And by conditional independence

$$\{E[\beta_i \mid Z_i = z_0 + \varepsilon]E[X_i \mid Z_i = z_0 + \varepsilon] - E[\beta_i \mid Z_i = z_0 - \varepsilon]E[X_i \mid Z_i = z_0 - \varepsilon]\}$$
$$+E[\alpha_i \mid Z_i = z_0 + \varepsilon] - E[\alpha_i \mid Z_i = z_0 - \varepsilon]$$

Then using Assumptions (A1) and (A2) and taking limits as $\varepsilon \to 0^+$ we get

$$\lim_{z \to z_0^+} E[Y_i \mid Z_i = z] - \lim_{z \to z_0^-} E[Y_i \mid Z_i = z]$$
$$= E[\beta_i \mid Z_i = z_0]\{\lim_{z \to z_0^+} E[X_i \mid Z_i = z] - \lim_{z \to z_0^-} E[X_i \mid Z_i = z]\}$$

Assumption (RD) ensures the term in the curly-brackets on the RHS is not zero and division by this term completes the proof. The proof also makes it clear that the sharp design estimates a lower bound for the fuzzy design. In the sharp design $x^{+} - x^{-} = 1$ while in the fuzzy design $x^{+} - x^{-} \in (0, 1)$. Thus looking at the theorem we see that the estimate in the sharp case bounds the fuzzy estimate below.

Thus Hahn et al. (2001) demonstrates the weak assumption required to identify an average treatment effect in regression discontinuity designs. How should we interpret this treatment effect? A relatively common idea is that the treatment effect is only valid for the subgroup whose assignment variable is at the threshold. This is similar to how instrumental variables estimates the effect for the subgroup captured by the instrument. However for Regression Discontinuity this interpretation of the average treatment affect is incorrect.

Lee (2008) shows that in fact regression discontinuity estimates a weighted average treatment effect across all individuals who face the treatment or the control, where the weights are heavier for individuals who are ex-ante more likely to be close to the cut-off. Lee and Lemieux (2009) explain this more intuitively. However because I only observe Z_i once for each individual it is normally impossible to know anything about person i's ex-ante distribution of Z_i. When the distributions are similar for everyone then the treatment effect estimated by regression discontinuity is close to the average treatment effect for the whole population. When the distributions are very different the estimated treatment effect for regression discontinuity is likely to be different to the average treatment effect for the population. They give an example where scholarships are awarded to people who get more than 50 per cent on an exam. If people who score 90 per cent had no ex-ante chance of getting a score near the threshold then the regression discontinuity gap is dominated by the effect of people near 50 per cent. If scores are less reliably obtained then the gap will apply to a broader sub-population.

In this data the ex-ante distribution of age at interview depends on two things, date of birth and survey date. I take date of birth to be essentially random. Survey date is not random, being heavily concentrated around September and October. However combined with date of birth the age of person i at interview date is essentially uniformly random. Thus the treatment effects I estimate arguably applies to a relatively broad proportion of youths on the award minimum rates for people in the year groups either side of the birthday of interest.

2.2. Estimation

Much of the theoretical literature recommends estimating the treatment effect in regression discontinuity by local (linear) regression (Hahn et al. (2001), Porter (2003), Imbens and Lemieux (2008)). I am interested estimating the regression function at a single boundary point and this method has been shown to work well. Imbens and Lemieux (2008) specifically suggest estimating

$$\min_{\alpha,\beta,\tau,\gamma} \sum_{i=1}^{N} 1\{z_0 - h \leq Z_i \leq z_0 + h\}(Y_i - \alpha - \beta(Z_i - z_0) - \tau X_i - \gamma(Z_i - z_0)X_i)^2$$

Where Z_i is the assignment variable, X_i an indicator which is 1 if treated, z_0 is the cut-off, and h is the band-width which determines the distance from the cut-off from which observations will be admissible. $\alpha, \beta, \tau,$ and γ are parameters which are estimated. τ is the estimated discontinuity. The $(Z_i - z_0)$ ensures I estimate the effect at the cut-off. The $\gamma(Z_i - z_0)X_i$ term allows for different linear slopes for the treated and control groups. This is important because imposing the restriction that the slope is the same for both the treatment and control group allows observations in one group to affect the estimated outcome for the other group as it approaches the cut-off. The method proposed by Imbens and Lemiuex is attractive because of its simplicity. By using the rectangular kernel I can estimate the discontinuity by appropriately using Ordinary Least Squares methods.[2] Furthermore the Robust Standard Errors are valid for inference. Imbens and Lemieux (2008) note that while more sophisticated kernels can be used that they are only likely to have much impact in situations where the results are not robust in any case. They suggest simply sticking to the rectangular kernel and reporting several band-widths. Despite potentially making little difference, I also check my results using the triangular kernel. In a regression discontinuity setting this kernel has the attractive property of weighting values near the cut-off more heavily and away from the cut-off more lightly.

A key consideration in regression discontinuity estimation is bandwidth selection. The larger the bandwidth the greater potential for my subgroups to differ. Ideally I would only use individuals on the threshold – then the treatment group and the control group should have exactly the same distributions of underlying predetermined covariates. However this is impossible because everyone who is on the threshold is treated and so there is nobody in the control group – for example if you are exactly 19 you have the 19 year old award wage minimum. Thus, as Imbens and Lemieux (2008) point out, there is an unavoidable need to

[2] I have tested careful use of OLS against -lpoly- Stata's inbuilt local linear regression estimator to ensure they give the same numerical results. A downside of -lpoly- is that it does not report standard errors

extrapolate. There is a trade-off between using more data to estimate the regression function at the cut-off more accurately and making the sub-groups less and less similar by including observations further away from the threshold. It is necessary to specify a bandwidth in order to carry out local linear regression. One method to select the bandwidth is by using a cross-validation procedure as in Ludwig and Miller (2005). A much simpler method is to estimate the local linear estimates using several bandwidths to check for sensitivity(Imbens and Lemieux (2008)). This is the method I use.

Lee and Card (2008) point out that when the assignment variable is measured discretely the standard errors can be biased. Intuitively if age is only observed in months then there will be many observations with the same month – the level of measurement of the discrete assignment variable. Thus we have no data on the difference in the age of youths within a month of their birthday. This bias in standard errors can be overcome by making appropriate parametric assumptions. The assignment variable for this research is the number of days from a birthday. This is discrete but I assume it is close enough to being continuous that it doesn't affect my results and hence I ignore this potential problem.

2.3. Validity

The mathematically necessary conditions for identification in Regression Discontinuity designs established by Hahn et al. (2001) are fundamentally untestable – I cannot test whether the expected outcome treated individuals would get if they were not treated is the same as the expected outcome of non-treated individuals. As such some researchers have looked to develop more practical means for checking the validity of a regression discontinuity design. Lee (2008) , using a very general treatment assignment mechanism, develops direct tests which should always pass when the necessary continuity assumptions hold. McCrary (2008) develops a test to investigate the possibility that people manipulate their observed assignment variable measurements – such manipulation would imply the continuity assumptions do not hold and the control group would not represent a valid counter-factual. Lee (2008) shows that if the assignment variable is continuously distributed and has some random chance element then the causal inferences from regression discontinuity designs can be almost as credible as those from randomised experiments – this is true so long as people cannot perfectly control their observed score, partial control is no problem. He proves that this treatment assignment mechanism fulfills the necessary identification condition established in Hahn et al. (2001). Furthermore he shows that this treatment assignment mechanism implies that predetermined baseline covariates should have the same distribution for those people who just cross the threshold and get treatment compared to those people who just fail to cross the cut-off. This means that one test of the validity of regression discontinuity is to check that predetermined variables do not change discontinuously at the cut-off.

In this research the assignment variable is age in years at the date of the HILDA interview – in fact the age reported to the employer, and not the age reported to HILDA, will determine the award minimum wage but it seems very difficult for an employee to lie about their age to obtain a higher award minimum wage, especially considering pay is likely to be linked to the tax office. For any individual age as of their HILDA interview date is essentially random they have no control over when the interview is conducted. Interviews take place around September every year but the exact date differs for each person from year to year. Thus I can test if the randomisation implied by regression discontinuity has worked by checking for similarity in the distributions of characteristics such as year 12 attainment which ought to be predetermined for 19, 20, and 21 year olds. Crudely I can check this by looking at the year 12 attainment rate for my control groups and treatment groups. A more sophisticated method is to look for a discontinuity in year 12 attainment for people either side of a birthday. It seems unlikely that such a discontinuity would exist suggesting that at least we are comparing youths of similar characteristics – however I will run these tests specifically in any case.

There is one last important consideration to make when thinking about my regression discontinuity design. Everyone turns 19 eventually. This means that assignment to treatment is essentially inevitable. Lee and Lemieux (2009) discuss several important implications of this. Firstly the parallel to random experiments in most Regression Discontinuities does not apply. Random experiments require some uncertainty about whether a person will be treated or not. In my case everyone will be treated, eventually. Secondly, validity tests based on looking for discontinuities in pre-determined covariates can be uninformative. If I follow a single cohort then they necessarily share the same distribution of underlying covariates. However in my research I observe people born across different years. They are not just the same group and the validity tests can hope to tell us something. Furthermore regression discontinuity designs using age as the assignment variable can only hope to identify short-run effects. If I wanted to look at the effect of an increase in award minimum wages over a 5 year period then both the under 19 initially and the over 19 initially will have had approximately 5 years of treatment. The final important consideration is that when treatment is anticipated and inevitable people may change their behaviour in order to take advantage of it. Thus the treatment effect can be muted or accentuated. To date there are no firm guidelines as to how to deal with these issues. In my research I hope manipulation is not widespread. Firstly, even if award minimum wage increases will result in firms having to pay their workers more once they reach a birthday I have no reason to expect this will make them change their pay prior to a birthday – particular on any large scale. However, it is possible that firms facing the requirement to raise a worker's wages on their birthday could use early wage increases as a reward for good employees. Importantly, it is possible that as youths approach a birthday firms give them warning that their casual work will end. This could enable firms to avoid paying any wages at the higher award minimum rates. This could reduce employment slightly for youths just before a birthday and hence cause my estimate of the treatment effect on employment to be biased downwards.

3. Background

Australia historically had a complex set of rules regarding minimum wages. A modernisation program is currently being implemented with a large set of new rules taking effect as of 1st of January 2010. However given my data is all collected by early 2009 I will not be affected by these new rules. Thus this section will describe the background of minimum wages – focusing on junior award minimum wages – before modernisation. It will introduce the institutions which have controlled minimum wages in Australia. It will describe Australia's system of minimum wages. It will look at junior award minimum wages and discuss the difficulty of identifying which youths are receiving the award minimum wage in data and how this affects my ability to know what proportion of youths get the award minimum wage.[3]

Australia has had several different organisation to regulate minimum wages. The Australian Industrial Relations Commission (AIRC) operated from 1988 until the end of 2009. The Australian Industrial Relations Commission had several roles. However between 2006 and 2009 the Fair Pay Commission administered wage setting decisions. Since the 1st of August 2009 Fair Work Australia's Minimum Wage Panel has been responsible for altering minimum wages.

These institutions have used a range of tools for setting minimum levels of pay in Australia. Prior to the recent award-modernisation, minimum pay in Australia was determined by awards and a Federal Minimum Wage (FMW). Awards specify minimum conditions of employment beyond legislated minimum terms – including minimum pay. Enterprise Awards are a special type of award specifically made to apply to a single business,

[3] Most of the information for this section comes from http://www.fwa.gov.au/ or http://www.fairwork.gov.au/ the Ibsites for Fair Work Australia and the Fair Work Ombudsman respectively. A good starting point with many information on Australia's minimum wage framework is http://www.fairwork.gov.au/resMyces/pages/Glossary.aspx the Fair Work Ombudsman glossary.

activity, project or undertaking. Australia had more than 1500 awards – a big part of award-modernisation involved the review of these awards and the subsequent creation of 122 modern awards. In addition to Enterprise awards there were many state based awards which covered workers in specific occupations. The Federal Minimum Wage – now replaced by the National Minimum Wage – provided a safety net for workers not covered by an award scheme.

Award minimum wages vary from job to job. For example, the 2009 pay-scale for retail staff in South Australia indicates a minimum wage for standard staff working – at a workplace that opens on Saturdays – of $8.27, 50 per cent of the adult rate for workers under 17, to $14.88, 90 per cent of the adult rate for 20 year olds. For ticket machine operators the rates Ire $8.50 to $14.48 for the same age groups. Both groups have minimum wage increases of 10 per cent of the adult rate per year.

In addition to these minimum wage tools Australia has a long history of special minimum wage rates for youths. For the purposes of Australian labour laws people under 21 years of age are considered youths. Youths are not covered by the National Minimum Wage –nor were the covered by the Federal Minimum Wage. Furthermore many industry awards provide special rates of pay for youths. In 1999 the Department of Employment, Workplace Relations and Small Business surveyed a census of federal awards.[4] They found that 45 per cent of awards – 755 of 1690 awards – had junior rates. Furthermore many awards lacked junior rates because they were unnecessary. This could be because requisite qualifications essentially precluded juniors, because juniors could do the job as Ill as adults, or because of safety reasons. However most awards relevant to youth workers have special junior rates. Furthermore where junior rates were included they often started at 50 per cent of the adult rate for 16 year olds and increased by 10 per cent so that 21 year olds earned the full adult rate.[5] Despite the age of this report these aspects of awards are typical throughout the period for which I have data. Because of the highly specific occupational nature of awards it is very difficult to identify which individuals are paid award minimum wages in data without specifically asking people if they are paid the award minimum rate. Typically measures of occupation are far too broad to be of use. Thus it is difficult to determine what proportion of individuals are paid award minimum wages at any one time. The Australian Bureau of Statistics publish the Employee Earnings and Hours report. It has a large section on pay

[4] Australian Government (2006), Australian Government Submission to the Fair Pay Com-mission. p 157.
[5] Australian Government (2006), Australian Government Submission to the Fair Pay Com-mission. p 157.

determination methods. However despite having information on youths it does not have a section reporting the proportions of youths on award minimum wages.[6]

Flatau et al. (2008) produced an Australian Fair Pay Commission report which estimated the number of youths who receive wages which are consistent with junior award minimum wages. They used HILDA wave 5. Their method involved collecting award minimum wages for several different awards per occupational category to get a minimum and maximum junior award rate for broad occupation categories. Then they can compare youths calculated wages to see which youths fall inside the award minimum rate ranges. In retail trade they estimate 57 per cent of 15-20 year olds earn award minimum rate consistent wages, while a further 14 per cent earn wages below the minimum award rate for the broad occupation code. For accommodation, cafe, and restaurant work they estimate 68 per cent of youths earn award minimum wage consistent rates with a further 22 per cent earning below award minimum wage consistent rates. However the ranges of wages that are award consistent can often be considerably large – for the food service industry they find a minimum hourly youth award rate of $4.54 and a maximum rate of $13.29. Flatau et al. (2008), due to measurement limitations for individual's occupation, call any youth “award wage consistent” if they work in the food industry and earn betIen $4.54 and $13.29 per hour. This is likely to overestimate the number of youths on the award wage minimum. For example some of those youths could earn $7.00 in a particular food industry job which has an award minimum wage of $4.54. HILDA wave 8 asks people for information on the method used to determine their pay. These direct observations should enable much more accurate estimates in the future. However focusing solely on wave 8 to make use of this question results in having insufficient observations for my study.

Overall Australia's system for determining minimum pay levels results in many youths being paid junior award minimum wages which are age based. Awards typically increase by 10 per cent per year of age of an individual. I hope to make use of this variation in minimum wages to identify the causal impact of minimum wages in my research.

4. Data and Summary Statistics

This section describes my data, provides summary statistics of my samples and presents my research design.

4.1. Data

[6] EEH reports can be found at http://www.abs.gov.au/

In order to identify the causal effects of changes in youth award minimum wages I use data from the longitudinal survey HILDA 8 in-confidence. The HILDA survey began in 2001 and collects information on people living in selected households each year.
The three most important variables for which I need information are age, current work hours, and current wage. These variables will let us undertake a regression discontinuity approach by contrasting how current work hours or current wages differ for youths slightly before or after a birthday. HILDA in-confidence reports data on birth dates and interview dates. Combining these allows us to calculate the age in years of any individual in the data set as (birthdate - interviewdate)=365.25 – because this does not perfectly implement the leap year rule I manually recode the age of people who are interviewed on their birthday as this helps Stata plot the discontinuity more clearly. HILDA also reports observations for current Iekly income from wage and salaried work and current weekly hours worked. Taking the ratio gives us the current wage. Individuals report whether their work is for an employer for wages or a salary, for a family business, as a volunteer, or self-employed. Because this report is only made once per individual I use main job weekly wages and hours to avoid problems from extra jobs not being with an employer for wages and salaries – I.e., I do not expect award minimum wages to have any causal impact on people who are volunteers or self-employed.

I focus on a separate subsample for each of the two causal effects I look to identify. For identification of the effect on actual wages of an increase in award minimum wages of 10 per cent I focus on only those youths who are employed for wages or a salary. However, several of the wages calculated using current weekly wages and current hours are extreme – values for youths who work for an employer for wages or a salary as high as a couple of hundred dollars per hour and as low as zero dollars per hour. Therefore in order to prevent my analysis being influenced heavily by these extreme and most likely incorrect wages[7] I focus on youths in the 5th to 95th percentile of the wage distribution. In order to identify the effect of increasing award minimum wages on employment I look at all individuals who are unemployed or are employed for a wage or salary. Furthermore with both sub-samples I exclude people who have badly inconsistent data – such as reporting their hours worked in their main job as being greater than the hours they work in all of their jobs combined. I also

[7] The HILDA user's manual suggest such extreme values may be the result of asking for current weekly wages and current weekly hours in separate parts of the questionnaire and also notes that people may not be good at thinking in terms of hours per week. N. Watson (ed) (2010), 'HILDA User Manual Release 8', Melbourne Institute of Applied Economic and Social Research University, of Melbourne, pp. 49 and pp. 94.

exclude some people who fail to answer certain questions about personal characteristics however the number excluded in this way is small. HILDA also provides information on a large number of personal characteristics for people in the dataset. These covariates are useful for two reasons. Firstly, a regression discontinuity analysis will be invalidated if the outcome variables for the non-treated group do not represent the valid counter-factual for the treated group (HTV, 2001). One way this would occur is if the two groups vary in some systematic way which affects the outcomes – the whole point of regression discontinuity designs is to compare two groups which are similar in most respects other than treatment. Having observations on personal characteristics allows us to check whether underlying differences betIen the groups is related to differences in outcomes – I.e., I could check if schooling was similar for treated and control groups. In fact Lee (2008) proves that a testable corollary of the critical assumption needed for identification of a causal effect in regression discontinuity designs that the control group represent the true counterfactual is that predetermined baseline characteristics be continuous across the assignment threshold. A second benefit to having data on personal characteristics is that if they vary smoothly over the assignment threshold then including them in regression analysis should not alter the value of my estimates but should improve their precision.

4.2. Summary Statistics

I now present summary statistics on the samples used for my analysis. Panel A of Table 1 reports selected statistics for my employment subsample, while panel B reports selected statistics for the wage subsample.

Looking firstly at panel A there are around 800 youths for each treatment group and another 800 for each control group. Each pair of columns represents the pre and post treatment groups for a certain birthday at which award minimum wages will increase. For example column (1) presents descriptive statistics on youths are just younger than 19 while column (2) represents the same statistics for the group just older than 19 who have the higher award minimum wages. Sex and ancestry are clearly predetermined. Looking at these summary statistics we see the means for these variables are similar regardless of treatment.[8] In fact comparing treatment and control groups year by year we see that the difference in means for these predetermined variables is never more than 0.03. Combined with large

[8] Lee (2008) tells us that predetermined variables should be distributed similarly in both the control group and the treatment group if regression discontinuity has worked. The similarity will be greater the smaller the bandwidths, and ultimately I can test My design validity by confirming there are no discontinuities in predetermined variables.

standard deviations it seems unlikely I will find any discontinuities in these variables, making us more confident of the validity of the research design. Interestingly for the youths close to their 21st birthday there appears to be a rather big difference in the proportion of youths with year 12 completion as their highest achievement – the percentage of youths with year 12 as their highest achievement is 56.7 for the control group in column (5) but only 49.5 for the treatment group in column (6). However, highest education is not predetermined. In fact many youths around age 21 will be completing higher education. When I test the validity of the regression discontinuity design I check that year 12 completion – which I estimate by the number of youths with year 12 education or higher – does not have any discontinuity at the threshold. Other covariates are also similar for all treatment and control groups, typically varying by around 2 per cent with the largest differences being around 5 per cent. Furthermore the composition of the baseline covariates seems highly stable across different treatment-control groups.

Panel B reports selected summary statistics for the group of youths used for the wages analysis. This time I have about 650 people in each treatment group and another 650 in each control group. Once again comparisons of each treatment group with its corresponding control group generally show little difference. The issues with highest education are the same as before. The summary statistics for weekend work are of interest. Many award wages specify loading penalties which require employers to pay more for weekend work. In all birthday pairs, the proportions of youths doing weekend work is about 5 per cent higher for the young group who are soon to have a birthday than for the old group who have just had a birthday. This may represent firms offering less Iekend work to employees who have just had their birthday, as the dollar cost of the penalties will be higher for higher wage workers. On the other the summary statistics could equally simply result from more youths moving into full-time work as they get older.

4.3. Research Design

The summary statistics in table 1 make it clear that when I split youths into groups according to those who are up to half a year above, and those half a year below a birthday the groups are highly comparable along a number of covariates. The identification strategy takes advantage of the fact that award minimum wages are typically age based for youths – increasing by around 10 per cent per birthday. I look to identify two effects. The effect of award minimum wage increases on actual wages, and their effect on employment hours. I use a regression discontinuity approach. This approach compares outcomes of interest either side of the treatment cut-off to look for a discontinuity. In theory I get the limit of a regression function approaching the cut-off at either side and compare them.

I implement this strategy graphically and using a very simple version of a local linear estimator – using the rectangle kernel. The actual wage affect of award minimum wages is something that has not been estimated for Australia before. If award wages increase but are binding for nobody then I will see no effect. If wages are set at the award level for all youths then I will expect at least a 10 per cent discontinuity in actual wages over the birthday.

One complication with this strategy is that in fact award minimum wages do not affect the entire youth labour market – for example some award minimum wages pay adult rates for anyone 18 years or older, thus people in these jobs will not face an increase in their award minimum wage when they have a birthday. As such perhaps a strict discontinuity is not best. A fuzzy discontinuity would take into account the fact that the probability of any youth facing an award minimum wage increase is the same as the proportion of youths in occupations covered by awards – whether or not those awards bind. Unfortunately this data is unavailable to us. HILDA only has information about occupation type at the industry level and this is not detailed enough to determine which youths are covered by awards. HILDA 8 includes wage determination and identifies youths who are paid exactly the award minimum wage however this would leave us with too little data for our research. Furthermore the discussion in HTV (2001) makes it clear that the estimate using a strict design is at worst a loIr bound for the effect if a fuzzy design is truly appropriate – all estimates will be divided by the jump in probability of treatment which must lie between 0 and 1.

It is important to recognise that the impact I look to identify on employment in particular is somewhat different to the impact most minimum wage research deals with. Research based on a once off change in minimum wages can look to identify the overall causal effect of minimum wages on employment. In doing so one factor which is typically thought to be important is the length of time in which the researcher allows the impact of the minimum wage to be felt. In contrast Regression Discontinuity based on the age of youths by its very nature cannot include lags. One aspect of Australian institutions that may make lags less important for youth employment effects is the fact that junior award minimum rates have been around for a long time and the relative saving from hiring a younger worker has been steady – a as junior rates have been proportions of adult rates. Hence firms have had plenty of time to optimise themselves to best take advantage of cheap labour –one possible explanation for the need for long lags to capture employment effects is an argument about slow ability to change capital and the interaction of capital and labour on productivity rates.

A further difference is in the kind of labour-labour substitution which is possible under this policy compared to a large range of other minimum wage research. If the minimum wage is raised for all youths then when looking at the effect on employment researchers must be aware of the potential for employers to substitute away from relatively more expensive youths to relatively less expensive adults who may have other advantages such as more experience – the difference in minimum wages is important. In this situation employers face different choices, including the choice to substitute down to a younger and cheaper employee. This may change the dynamics. However an advantage of the regression discontinuity approach is its quasi-experimental nature. By comparing youths close to but on differing sides of a birthday I can hopefully avoid misspecification and problems due to unobserved heterogeneity. When regression discontinuity works it can produce results similar to those found from a random experiment and hence the Average Treatment Effect identified represents a causal impact.

5. Graphical Analysis

In this section I present graphical evidence for discontinuities in actual wages and hMys of employment at youth birthdays. I present graphical analyses for youths either side of their 19th, 20th, and 21st birthdays – younger ages are problematic for this analysis as they can coincide with school leaving ages and other legislative differences in various states. All graphs show average outcome by age. When I analyse the 19th birthday I include individuals who either will turn 19 within the next 6 months or have turned 19 within the last 6 months – for the 20th and 21st birthdays I use the equivalent groups of youths. A linear model has been super-imposed which corresponds to the line produced by a local linear regression estimated at the birthday using the rectangular kernel and a half-year bandwidth. It allows for a different intercept and slopes on either side of the birthday.

5.1. Effect of Minimum Wage Increase on Actual Wages

I now provide graphical evidence on the effect of award minimum wages on actual wages. This relationship is important to my research for two reasons. Firstly, as far as I am aware, this relationship has not been estimated for Australia. HoIver, if policy makers want to increase youth award minimum wages in order to increase youth remuneration then it would be useful to know how minimum wages translate into actual wages. This relationship need not be dollar for dollar. Suppose all youths were paid well about award minimum rates. Then small increases in the award minimum would have no effect on actual wages. On the other hand if all individuals were bound by a youth award minimum wage then increasing award minimum wages by a dollar would increase actual wages by a dollar as long as firms complied.

Figure 1 provides evidence on the impact of increasing award minimum wages on the realised wages of 19 year olds. The control group constrains 641 youths who are younger than 19 and will face 18 year old minimum wages. The treatment group contains 632 youths just older than 19 who must be paid at least the 19 year old minimum award wage, which for any state/occupation award scheme is typically 10 per cent higher than the 18 year old counter-part. From the graph we see that the average wage per half-month grouping is well below $12.50 for the younger group and above $12.50 for the old group. The discontinuity at the threshold is estimated at $0.68 and with a standard error of 0.39 this is signigicant at the 10 per cent level. Given the mean actual wage for youths under 19 in My sample is $11.66 this corresponds to a 5.8 per cent increase in actual wages in response to a 10 per cent increase in the minimum wage.

Figure 2 provides evidence on the same outcome for 20 year olds. This time the control group consists of 658 youths just under 20 years of age while the treatment group consists of 662 youths just over 20. The graph shows a discontinuity of $0.66 with a standard error of 0.4001. This discontinuity is significant at the 10 per cent level. Given the mean wage for youths betIen 19.5 and 20 years of age is $13.12 in this sample the discontinuity represents a 5.0 per cent increase in actual wages – I.e., a 10 per cent increase in award wages flows through to a 5 per cent increase in actually wages for the average 20 year old youth. Importantly the graphical analysis suggests that perhaps wages rise steeply just before the 20th birthday. I will look at reasons for why this might be when I examine the econometric results more carefully later on.

Figure 3 finally presents the actual wage effect evidence for the group who are near their 21st birthday. This time the control group consists of 680 youths under 21 while the treatment group consists of 642 youths just over 21. The average wage over half-month sub-groups is typically less than $15.00 for youths with the loIr award wages but greater than $15.00 for youths with the higher award minimum wages. The estimated discontinuity is $1.06 with a standard error of 0.4139. This is significant at the 1 per cent level. This represents a 7.3 per cent increase in actual wages for youths who face the 10 per cent higher award minimum wages. In this age group we do not see the same steep increase in wages just before the birthday.

Overall this provides relatively strong evidence that higher award minimum wages have a positive effect on actual wages. The discontinuities for both all age groups are statistically significant at the 10 per cent level, while the results for youths near their 21st birthday are significant at the 1 per cent level. They suggest a 10 per cent increase in award minimum wages flows through to an average increase in actual wages of betIen 5 per cent and 7 per cent. This effect is smaller than we would see if all youths were paid the exact award minimum wage – in that case I would expect an effect closer to 10 per cent as increases in the award minimum wage would flow through to actual wages dollar for dollar except for non-compliance and spill-overs.

5.2. Effect of Award minimum Wage Increase on Hours Employed

In Figure 4 I present evidence for the impact of award minimum wages on hours of employment for youths near 19 years of age. The control group has 836 youths younger than 19 while the treatment group has 826 youths older than 19. From the graph there seems to be a lot of variation in the sample but no real difference in the typical number of hours worked betIen the groups. This is borne out by the fact that the standard error is larger than the discontinuity. I estimate a positive discontinuity showing that youths in the older group with higher award minimum wages typically work 1.59 hours more than their younger counter-parts. However with a standard error of 1.7176 this effect is not statistically different from 0.

Figure 5 presents the same evidence for 20 year olds. This time there are 834 youths in the control group and 816 youths in the treatment group. There seems to be slightly more variation in employment hours for the younger group. Once again there is no noticeable difference between the typical level of average hours between the groups. This time the estimated discontinuity is negative so that higher award minimum wages are correlated with lower hours – the estimated discontinuity is -1.27 hours. However once again the discontinuity lacks any statistical significance as the standard error is large at 1.8169.

Figure 6 reports the final piece of graphical analysis for the effects of award minimum wages. This graph presents evidence of the effect of an increase in award minimum wages by 10 per cent for youths near their 21st birthday. The control group contains 780 youths younger than 21 and the treatment group contains 782 youths older than 21. The estimated discontinuity is -1.43 hours, once again negative. However with a standard error of 1.8626 it also fails to have any statistical significance. Overall there doesn't seem to be much of an effect of higher award minimum wages on hours of employment. None of the estimated effects were statistically significant and they didn't even all have the same direction.

6. Econometric Results

This section reports my econometrics results. I use the semi-parametric local- linear estimator that is widely used in the literature. I report estimates for several bandwidths to indicate the sensitivity of my results. There is a tension in Regression Discontinuity estimation – using smaller windows of data is likely to keep treatment and control groups more similar in terms of underlying characteristics, while using larger windows improves the precision of the estimate at the discontinuity. Imbens and Lemieux (2008) suggest using the rectangular kernel and checking several bandwidths. They also note that the inclusion of observed characteristics makes estimates more efficient – in line with Lee (2008). I use several personal characteristics for controls. Specifically I estimate,

$$\min_{\alpha,\beta,\tau,\gamma,\delta} \sum_{i=1}^{N} (Y_i - \alpha - \beta(Z_i - z_0) - \tau X_i - \gamma(Z_i - z_0)X_i - \delta W_i)^2 \, K\left(\frac{Z_i - z_0}{h}\right)$$

Where Y_i is the outcome of interest, Z_i is observed age – the assignment variable. X_i is an indicator that is 1 for treated youths. A youth is treated if they are just over their birthday. This means that minimum wage awards that apply to them are on average 10 per cent higher than the awards that apply to youths who are just about to have a birthday. W_i is a set of covariates measuring various personal characteristics. $\alpha, \beta, \tau, \gamma,$ and δ are parameters which are estimated. τ represents the estimate of the discontinuity. α represents the constant for the control group. β is the estimated slope for the control group while γ is the estimated slope for the treatment group. The term $\gamma(Z_i - z_0)X_i$ allows different slops on either side of the cut-off and this prevents observations on one side of the threshold affecting the estimate of the regression function on the other side of the cut-off. W_i represents the vector of control variables which is set to 0 when observed characteristics are not included. This model can be estimated using various kernels. For the rectangular kernel we get

$$K\left(\frac{Z_i - z_0}{h}\right) = \frac{1}{2} 1\left\{\left|\frac{Z_i - z_0}{h}\right| < 1\right\} = \frac{1}{2} 1\{z_0 - h < Z_i < z_0 + h\}$$

and I can estimate by OLS on the appropriate sub-sample. The rectangular kernel estimates the local linear regression function at the cut-off using all observations whose assignment score is within h of the bandwidth and Iights them evenly. Thus as recommended in Imbens and Lemieux (2009) I implement this strategy by careful use of OLS and use robust standard errors for inference. Furthermore, I also check to see if results differ when the triangle kernel is used instead. The triangular kernel has the property of Iighting observations closer to the discontinuity more heavily than those far away. This is intuitively appealing in a regression discontinuity setting. Furthermore Imbens and Lemieux (2009) note that it has been shown that the triangular kernel is boundary optimal.

To implement the triangular kernel I use

$$K\left(\frac{Z_i - z_0}{h}\right) = \left(1 - \left|\frac{Z_i - z_0}{h}\right|\right) 1\{z_0 - h < Z_i < z_0 + h\}$$

Local linear regression can be done in Stata using the -lpoly- command. However, this command has two shortcomings. Firstly its syntax only allows two variables meaning I do not use any controls with the triangular kernel. Secondly, and more importantly, it does not produce standard errors. There are two ways around this. Imbens and Lemieux (2008) or Porter (2003) provide plug-in estimators of the variance for the fuzzy regression discontinuity design. Alternatively I can bootstrap. I implement the latter alternative. I carefully make sure I only use observations from within My bandwidth when estimating the bootstrapped standard errors.

Table 2, page presents regression discontinuity estimates for the impact of an increase in award minimum wage rates of around 10 per cent following a birthday on the actual wages received for those people who are working. It is divided into three panels. Panels A, B, and C report estimates for the discontinuity at the 19th, 20th, and 21st birthdays respectively. Column (1) represents a simple comparison of means. This is a crude first estimate of an effect of treatment. Importantly, if the outcome of interest, actual wages in this case, is positively correlated with age then I could find a difference in means even if treatment has no effect. Column (2) represents local-linear estimates using the rectangular kernel and no controls. Column (3) adds controls and column (4) reports estimates from the triangular kernel. Panel A reports estimates for the discontinuity at the 19th birthday based on 1,089 youths. I see that there is a highly significant difference in means. Youths aged 19 to 19 and a half earn on average $1.08 more than youths in the half year age bracket below 19. Do increasing award minimum wages have an effect on actual wages? To answer this I need more information. Column (2) shows estimates a positive discontinuity of $0.68 and this is significant at the 10 per cent level. Adding observed characteristics the estimate is $0.66 in column (3) and is still significant at the 10 per cent level. Adding the observed characteristics has little effect on the estimate or its precision, however the observed characteristics haven't added much in absolute terms to R-squared either. Finally using the triangular kernel the estimate is $0.67 which is very close to the estimates when using the rectangular kernel. This estimate is no longer statistically significant at even the 10 per cent level. These results suggest that a comparison of means overestimates any effect of increasing award minimum wages. All the local-linear estimates are around $0.67. Given an average wage for the younger group of $11.66 this represents a 5.7 per cent response of actual wages to an approximate increase in award minimum wages of 10 per cent. How sensitive is this result to changes in bandwidth? Firstly it is important to note that large variation in reported wages of youths required us to extend the bandwidth quite far away from the cut-off in order to get statistically significant results – if I extended the bandwidth further then I would have observations closer to another birthday than the one of interest, however given they would still have the same required award minimum wage rates maybe this would not be too problematic. In any case with a bandwidth of 4 months the estimates are between $0.64 and $0.77 with the triangular kernel providing the smallest estimate. With this amount of data none of the estimates are statistically significant. When the bandwidth is 2 months the estimates from the rectangular kernel are $0.52 and $0.49 when controls are omitted or included. Importantly the triangular kernel produces a much smaller estimate of $0.05. What is happening here? Another look at Figure 1 shows wages increasing just prior to a birthday. When the triangular kernel only includes 2 months worth of observations and heavily weights these observations near the birthday its point estimate at the birthday of the regression function for the younger group is very high as a steep straight line minimises its residuals. This effect has no statistical significance and disappears as I add more data to more precisely estimate the regression function. However it should make us more concerned about the validity of the results.

Panel B presents the same estimates as panel A. The results are similar in magnitude. The difference between means in column (1) is now $1.24 and is highly significant. Using the rectangular kernel and no controls I get an estimate of $0.66 which is significant at the 10 per cent level. Adding controls the estimate falls slightly to $0.62, the standard errors get slightly tighter too but overall there is no significance. The triangular kernel estimates also lack significance with an estimate of $0.55 which is far from being significant as the corresponding standard error is 0.43. These results seem to suggest that increasing award wages are not impacting on actual wages by much. Why could this be? Looking at Figure 2 we see a very big jump in actual wages just before the birthday. If this is due to anticipated birthdays then it would affect the estimates. Firstly it would bias the estimates downwards. Secondly, it would make them more sensitive to bandwidth selection near the time of manipulation. Why could a person's wage increase before their birthday? One possible argument runs along efficiency wage lines. If an employer will have to pay any of the award minimum wage bound workers more once they reach their birthday then they could increase wages slightly early in order to buy loyalty. Indeed with very small bandwidths I get non-significant negative estimates of the treatment effect. However for bandwidths of 2 months or more the estimates are all around $0.60.

Panel C reports the last estimates for 1,153 youths near 21 years of age. This age group shows the strongest discontinuity. The difference in means is estimated to be $1.06 and is highly significant. Using a rectangular kernel and no controls produces an estimate of $1.08, slightly bigger than just comparing means, and this estimate is significant at the 1 per cent level. Adding controls reduces the estimate to $0.92 and significance falls to 5 per cent. Observed characteristics which vary smoothly over the cut-off should not affect the estimates of the discontinuity a lot and so the small difference that arises when I add controls is a little uncomfortable. Finally estimation using the triangular kernel produces a similar estimate of $0.96 and this is also significant at the 1 per cent level. Using the estimate based on the triangular kernel and a mean wage for the control group of $14.51 I find that actual wages respond by 6.6 per cent when award minimum wages rise by about 10 per cent.

In Table 3, I report estimates for the effect on employment hours. I use the same models. This analysis use the employment subsample. Again there are three panels as before. A quick scan of the table shows no statistically significant differences in hours employed between the younger group with lower required award minimum wages and the older group with higher minimum award wages.

Panel A reports results for 19 year olds. For 19 year olds the older group with the higher award minimum wages seems to work more hours than the younger group. The average working hours is 0.65 more for the older group but this is not statistically significant. Furthermore there is no evidence for any treatment effect as most the standard errors for the local-linear estimates are around as big as the estimated effects themselves. There is a substantial difference between the estimates using the rectangular kernel with and without controls. However these estimates are very imprecise. Panel B reports similar results for 20 year olds. The mean difference is very small at 0.26 hours and completely insignificant. The estimated treatment effects are all insignificant and vary from positive to negative effects. Finally panel C presents estimates for 21 year olds. Once again there is no significant evidence of any difference in employment hours between the groups.

When looking at the impact of higher award minimum wages on actual wages there is some evidence of a positive effect. Although there was not very strong statistical significance for the 20 year olds a 10 per cent increase in award minimum wages seemed to cause an increase in actual wages of around 6 per cent. In comparison there seems no evidence at all that award minimum wages have any short term causal impact on employment.

7. Regression Discontinuity Validity

This section examines the validity of the regression discontinuity approach used in this paper. It does so in two ways. Firstly it looks for discontinuities in wages for other age groups who do not have different underlying required award minimum wages on either side of a birthday – I focus on wages because only wages showed significant treatment effects. Secondly it looks for discontinuities in predetermined covariates across the threshold for those youths for whom I look for a causal effect. Lee (2008) proved that with a very general assignment mechanism – which can account for non-perfect self-selection – that whenever regression discontinuity is valid it implies that all predetermined baseline covariates should vary smoothly over the cut-off. Thus he says that an applied researcher has essentially unlimited over-identifying restrictions to test. I look at year 12 achievement as this is arguably important for finding a job and for wage determination.

I look for discontinuities in wages at birthdays where award minimum wages do not change. This is an attempt to show that the effect identified by the discontinuity is indeed the result of the change in award minimum wages by showing there is no discontinuity when there is no change in award minimum wages. Perhaps the best people to look at are those near their 22nd birthday. These people are the most similar in age to youths who receive junior award wages. Thus hopefully the way their wages are affected by age and other characteristics (observable or not) is similar. Figure 7 presents a graph of average wages over the age intervals for people near their 22nd birthday. I also add in linear fits on either side of the birthday which are the same as the line produced by the local-linear estimator with the rectangular kernel estimating at the birthday. The discontinuity is estimated at 0.487 which is about two-thirds as big as the smaller discontinuities for youths when award minimum wages differed either side of the birthday. Furthermore the estimate is not even close to being statistically significant. Adding other characteristics as controls also fail to get a statistically significant discontinuity. So while there was a very big and highly significant discontinuity in wages at the 21st birthday there is no significant discontinuity at the 22nd birthday in this data. I have also produced graphs and estimates for people near their 25th, 35th, 45th, and 55th birthdays. There are no significant discontinuities at any of these birthdays under any of the local-linear estimation techniques used above.

Secondly I look for discontinuities in year 12 achievement. Year 12, the final year of high-school is typically completed by youths aged 18. Hence year 12 achievement should be predetermined for youths aged 19, 20, or 21. If regression discontinuity has worked then there should be near randomisation close to the border (Lee (2008)). Year 12 achievement should be distributed similarly in treatment and control groups and there should be no discontinuity in year 12 achievement rates at the age threshold. I constructed a variable for year 12 achievement by giving individuals a 1 for year 12 achievement if their highest education is year 12 or higher and a 0 if it is year 11 or lower. This is necessary because if individuals gain higher education then using year 12 as highest education is not predetermined, it will change as people upgrade their highest education. This constructed variable will still identify a person as having achieved year 12 even if their highest education changes from year 12 to Cert or to some sort of university qualification. However it is not an exact measure of year 12 achievement as it is possible for people to have higher education without ever completing year 12 –although this is unlikely for people aged 19, 20, or 21. Because I am primarily interested in checking the validity of the wage analysis I use the same methods for constructing my sub-sample as I used for that analysis although the qualitative result is the same even if I do not do this. Figure 8 presents a graph showing average year 12 achievement for youths near their 19th birthday. The fitted lines are the same as those from a local-linear estimate at the birthday using the rectangular kernel. From the graph there is no evidence of a discontinuity. The estimated discontinuity is 0.0261 with a standard error twice as large at 0.0504. Figure 9 shows the relevant graph for 20 year olds. Once again using the rectangular kernel there is no evidence of a discontinuity. The estimated treatment effect is 0.0773 and the standard error is 0.0485. The effect is not significant at even the 10 per cent level. Finally Figure 10 depicts the same graph for 21 year olds. There is no discontinuity in year 12 achievement. The estimated treatment effect is -0.0161 and the relevant standard error is almost three time larger at 0.0459. There are potentially unlimited over-identifying restrictions which can be tested in this way. I chose year 12 achievement because I imagine it could affect employment and wages. I find no evidence for a discontinuity in this predetermined covariate.

Thus I have provided two pieces of evidence to support the interpretation of the estimates of the treatment effect of increasing award minimum wages as causal. Firstly I have shown that there is no discontinuity in wages for people either side of a birthday for which minimum wages do not change. In particular I showed the graph for 22 year olds. Secondly I have shown that the predetermined covariate for year 12 achievement is smooth over the threshold.

8. Conclusion

Australia has a complex set of rules which govern the minimum pay a person can receive. Part of these rules involve special rates for youths which are typically fixed proportions of the corresponding adult award minimum wage. These youth rates increase by about 10 per cent per birthday and cover a large proportion of working youths in Australia. Recent estimates suggest up to 71 per cent of youths in retail work and 90 per cent of youths in the food industry earn wages that are consistent with or below relevant award wage minimums. As such this paper aims to use the variation in youth award minimum rates to identify the causal affect of a change in youth minimum wages on actual wages received and hours of employment worked.

Using a regression discontinuity approach I find evidence that increasing youth award minimum wages significant increases actual youth wages in the short run. I find no evidence that changes in minimum youth awards have any effect on short run youth employment.

This paper makes use of the fact that there is a well known assignment mechanism by which award minimum wages are increased – award minimum rates are higher for every year of age. Thus I implement a regression discontinuity design. One important consideration is that the effect I identify is different to much of the minimum wage literature. The effect is a strictly local and short-term effect. Furthermore employers have the option to substitute to younger, cheaper labour in the situation I have studied, which is different from much of the literature which uses an across the board increase in a minimum wage – potentially a youth minimum wage – to identify the effect.

I have presented some evidence to evaluate the validity of the regression discontinuity design used. I find no evidence of discontinuities in wages for age groups where award minimum wages do not differ. I also find that year 12 achievement – a predetermined covariate – varies smoothly over the cut-off. Checking predetermined covariates for smoothness over the cut-off is a form of over-identifying restriction, and for year 12 achievement the data passes. While there are potentially unlimited such over-identifying restrictions I could run I used year 12 achievement as it is clearly predetermined and also could plausibly affect both wages and employment.

My results are mixed. An important first proviso is that due to the inevitability of treatment agents optimising behaviour can affect the estimates of this research.

I find significant effects of award minimum wages on actual wages for all groups in the basic local-linear regression model using the rectangular kernel. In this specification actual wages respond by around 6 per cent when award minimum wages increase by around 10 per cent. However, these estimates are not as robust as they could be. All estimates show some sensitivity to bandwidth – particularly at very low bandwidths. Local-linear estimates using the triangular kernel are less sensitive than estimates using the rectangular kernel. However, even these estimates can be sensitive if the bandwidth is very small. Furthermore, the estimates have less significance with the triangular kernel. Importantly, this sensitivity is statistically insignificant and disappears very quickly. With a bandwidth from 2 to 6 months estimates are more robust.

I find no evidence of an short-term employment effect. Even mean hours worked is not statistically significantly different for individuals in the higher award minimum wage group relative to those in the lower award minimum wage group. Adding controls and changing the kernel all fail to find any evidence of an employment effect.

References

Australian Government, 2006: Australian government submission to the fair pay commission.

Flatau, P., A. M. Dockery, and T. Stromback, 2008: The economic and social circumstances of Australian young people aged 15-20 years. Australian Fair Pay Commission Melbourne.

Hahn, J., P. Todd, and W. Van der Klaauw, 2001: Identification and estimation of treatment effects with a regression-discontinuity design. Econometrica, 69 (1), 201-09.

Imbens, G. W. and T. Lemieux, 2008: Regression discontinuity designs: A guide to practice. Journal of Econometrics, 142 (2), 615-635.

Lee, D. S., 2008: Randomized experiments from non-random selection in U.S. house elections. Journal of Econometrics, 142 (2), 675-697.

Lee, D. S. and D. Card, 2008: Regression discontinuity inference with specification error. Jouynal of Econometrics, 142 (2), 655-674.

Lee, D. S. and T. Lemieux, 2009: Regression discontinuity designs in economics. NBER Working Papers 14723, National Bureau of Economic Research, Inc.

Leigh, A., 2003: Employment effects of minimum wages: Evidence from a quasi-experiment. Australian Economic Review, 36 (4), 361-373.

Ludwig, J. and D. L. Miller, 2005: Does head start improve children's life chances? Evidence from a regression discontinuity design. NBER Working Papers 11702, National Bureau of Economic Research, Inc.

McCrary, J., 2008: Manipulation of the running variable in the regression discontinuity design: A density test. Journal of Econometrics, 142 (2), 698-714.

Neumark, D. and W. Wascher, 2006: Minimum wages and employment: A review of evidence from the new minimum wage research. NBER Working Papers 12663, National Bureau of Economic Research, Inc.

Porter, J., 2003: Estimation in the regression discontinuity model. Mimeo. Department of Economics, University of Wisconsin.

Table 1: Selected descriptive statistics (means with standard deviations in parentheses)

Age group	(1) 18.5–19 years	(2) 19–19.5 years	(3) 19.5–20 years	(4) 20–20.5 years	(5) 20.5–21 years	(6) 21–21.5 years
A. Employment subsample						
Female	0.505 (0.500)	0.525 (0.500)	0.512 (0.500)	0.526 (0.500)	0.504 (0.500)	0.535 (0.499)
English speaking ancestry	0.018 (0.133)	0.027 (0.161)	0.022 (0.145)	0.029 (0.169)	0.031 (0.173)	0.043 (0.204)
Other ancestry	0.063 (0.244)	0.059 (0.236)	0.073 (0.261)	0.069 (0.253)	0.073 (0.260)	0.092 (0.289)
Year 12	0.559 (0.497)	0.586 (0.493)	0.578 (0.494)	0.547 (0.498)	0.567 (0.496)	0.495 (0.500)
Cert	0.090 (0.286)	0.123 (0.329)	0.127 (0.333)	0.185 (0.389)	0.162 (0.368)	0.196 (0.397)
University	0.010 (0.097)	0.015 (0.120)	0.028 (0.164)	0.042 (0.200)	0.059 (0.236)	0.084 (0.278)
Observations	836	826	834	816	780	782
B. Wage subsample						
Female	0.507 (0.500)	0.491 (0.500)	0.491 (0.500)	0.497 (0.500)	0.475 (0.500)	0.531 (0.499)
English speaking ancestry	0.019 (0.136)	0.024 (0.152)	0.021 (0.144)	0.030 (0.171)	0.035 (0.185)	0.039 (0.193)
Other ancestry	0.033 (0.178)	0.040 (0.195)	0.052 (0.222)	0.045 (0.208)	0.060 (0.238)	0.073 (0.260)
Year 12	0.596 (0.491)	0.641 (0.480)	0.638 (0.481)	0.582 (0.494)	0.601 (0.490)	0.523 (0.500)
Cert	0.092 (0.289)	0.134 (0.341)	0.119 (0.323)	0.187 (0.390)	0.169 (0.375)	0.202 (0.402)
University	0.009 (0.096)	0.013 (0.112)	0.024 (0.154)	0.053 (0.224)	0.060 (0.238)	0.104 (0.306)
Weekend work days	0.401 (0.490)	0.358 (0.480)	0.407 (0.492)	0.355 (0.479)	0.363 (0.481)	0.301 (0.459)
Other work days	0.256 (0.437)	0.275 (0.447)	0.233 (0.423)	0.276 (0.448)	0.259 (0.438)	0.290 (0.454)
Observations	641	632	658	662	680	644

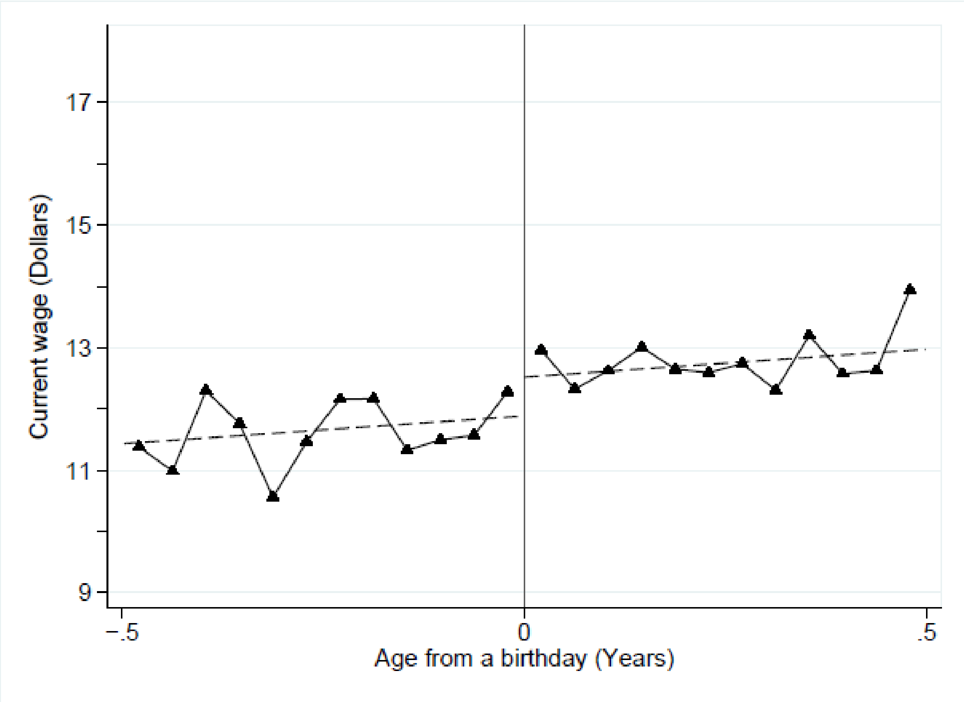

Discontinuity at threshold = 0.6815; with std. err. = 0.3945

Figure 1: The effect of an increase of award minimum wages on actual wages for 19 year olds. Sample is restricted to those youths who earn a regular wage or salary from an employer.

Discontinuity at threshold = 0.6634; with std. err. = 0.4001

Figure 2: The effect of an increase of award minimum wages on actual wages for 20 year olds. Sample is restricted to those youths who earn a regular wage or salary from an employer.

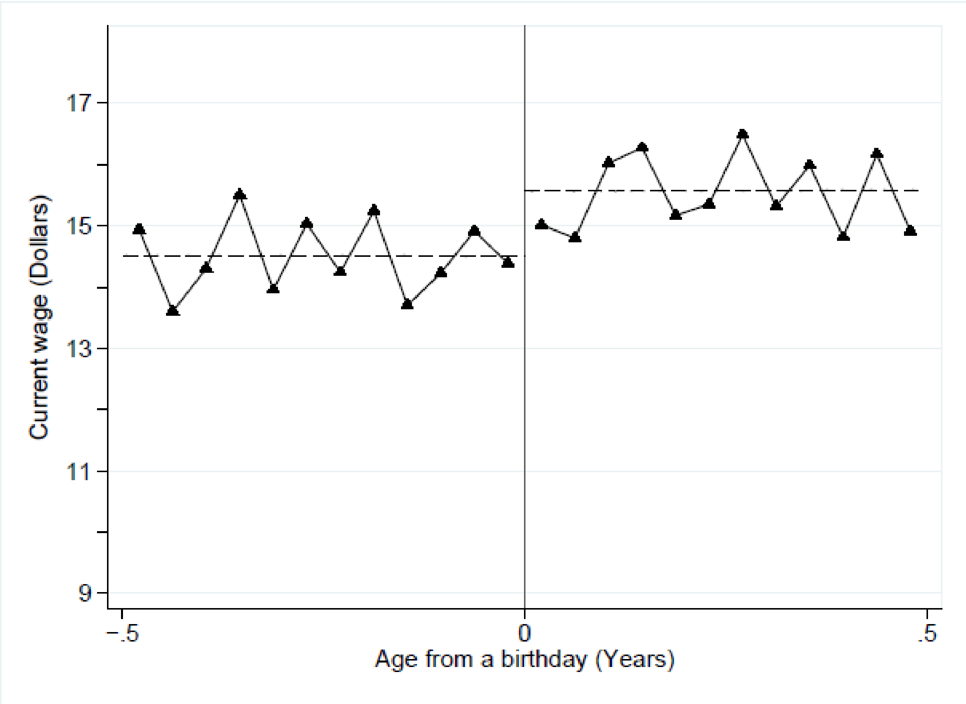

Discontinuity at threshold = 1.0595; with std. err. = 0.4139

Figure 3: The effect of an increase of award minimum wages on actual wages for 21 year olds. Sample is restricted to those youths who earn a regular wage or salary from an employer.

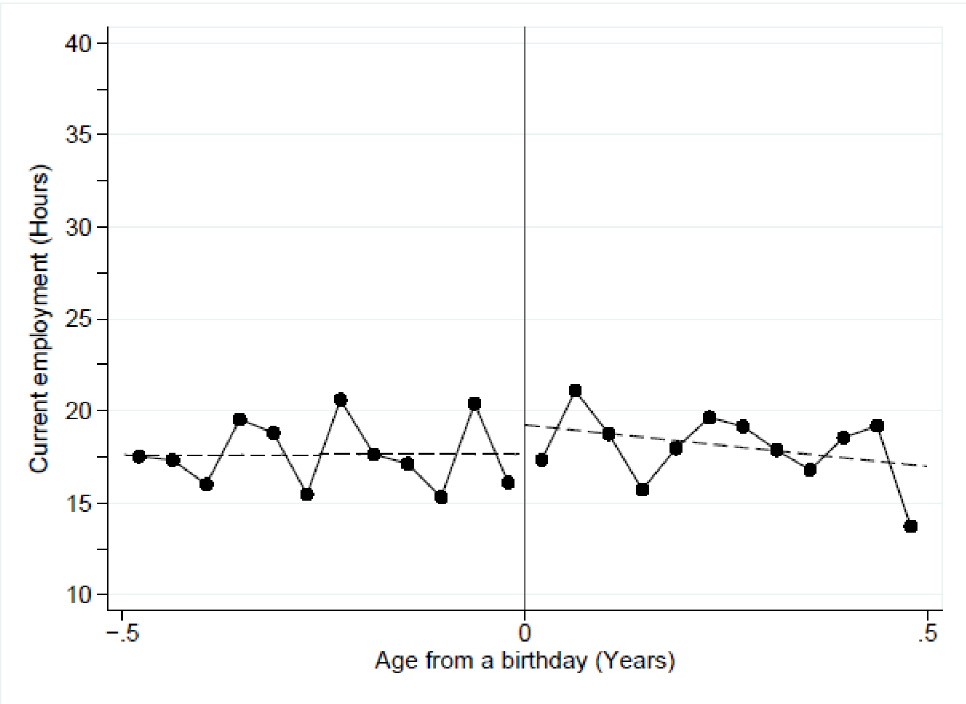

Discontinuity at threshold = 1.5911; with std. err. = 1.7176

Figure 4: The effect of an increase of award minimum wages on hours of employment for 19 year olds. Sample only excludes those youths with extreme wages.

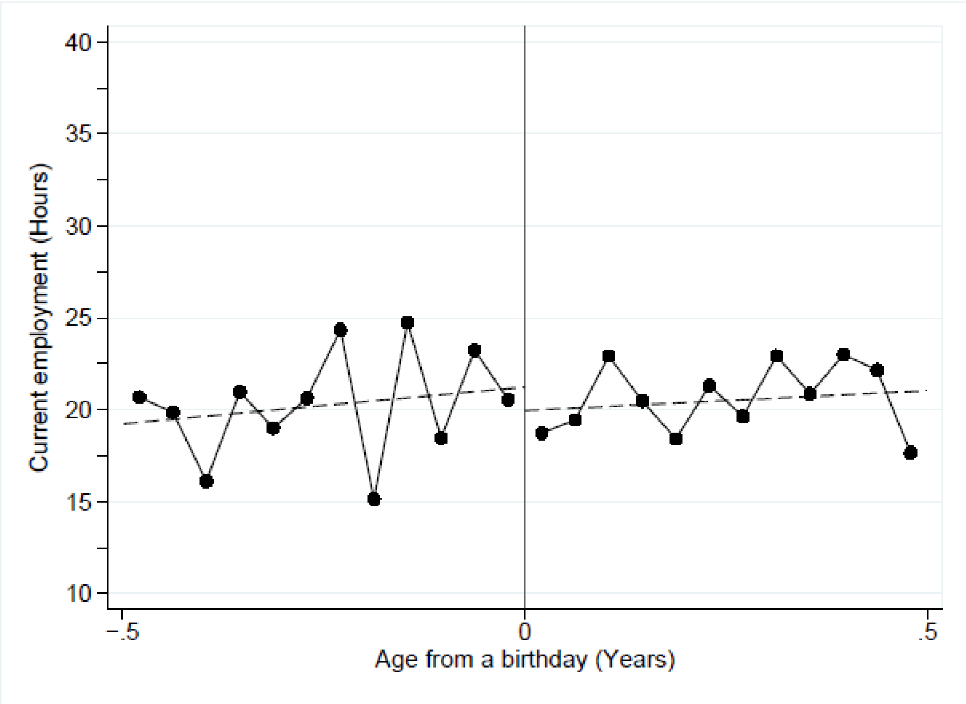

Discontinuity at threshold = -1.2728; with std. err. = 1.8169

Figure 5: The effect of an increase of award minimum wages on hours of employment for 20 year olds. Sample only excludes those youths with extreme wages.

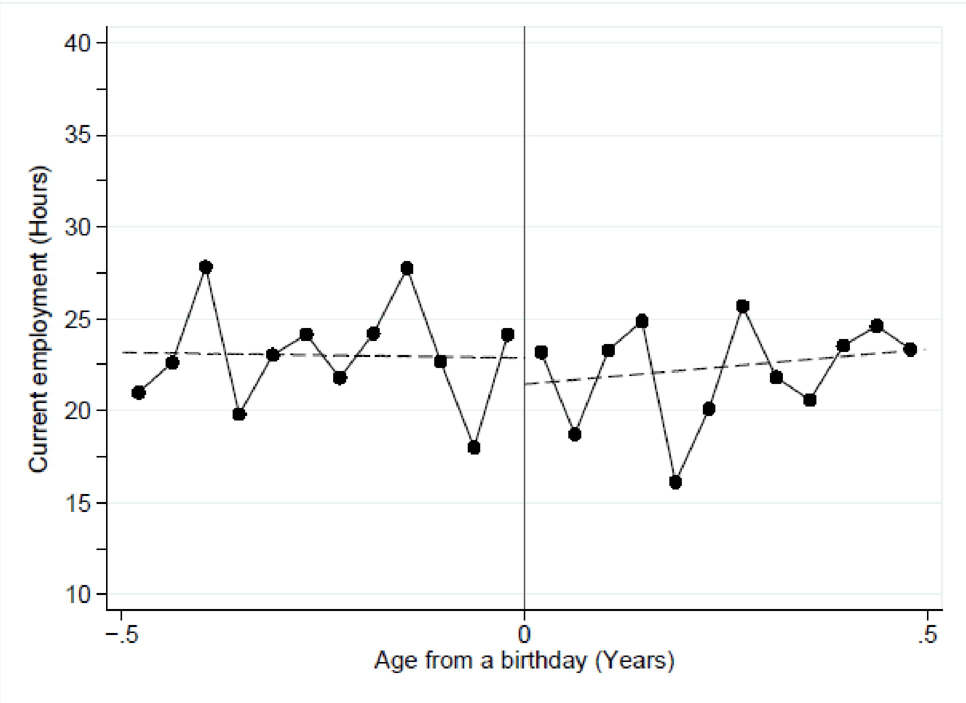

Discontinuity at threshold = -1.4317; with std. err. = 1.8626

Figure 6: The effect of an increase of award minimum wages on hours of employment for 21 year olds. Sample only excludes those youths with extreme wages.

Table 2: Estimates of the impact of award minimum wage increases on actual wages

	(1)	(2)	(3)	(4)
A. 19 year olds				
Discontinuity	1.0767	0.6815	0.6625	0.6654
	(0.19504)***	(0.3945)*	(0.3943)*	(0.4460)
Polynomial Order	0	1	1	1
Controls	No	No	Yes	No
Observations	1,089	1,089	1,089	1,089
R-Squared	0.0273	0.0287	0.1167	
B. 20 year olds				
Discontinuity	1.2479	0.6634	0.6205	0.5515
	(0.1994)***	(0.4001)*	(0.3885)	(0.4349)
Polynomial Order	0	1	1	1
Controls	No	No	Yes	No
Observations	1,140	1,140	1,140	1,140
R-Squared	0.0326	0.0354	0.1504	
C. 21 year olds				
Discontinuity	1.0633	1.0836	0.9266	0.9607
	(0.2043)***	(0.4144)***	(0.3953)**	(0.4639)***
Polynomial Order	0	1	1	1
Controls	No	No	Yes	No
Observations	1,153	1,153	1,153	1,153
R-Squared	0.0223	0.0223	0.1450	

Columns (1) to (3) report robust standard errors in parentheses (clustered by xwaveid).
Columns (4) reports bootstrapped standard errors (clustered by xwaveid, reps=1000).
* Significant at 10 per cent.
** Significant at 5 per cent.
*** Significant at 1 per cent.

Table 3: Estimates of the impact of award minimum wage increases on employment hours

	(1)	(2)	(3)	(4)
A. 19 year olds				
Discontinuity	0.6517	1.8218	0.5960	1.5773
	(0.8421)	(1.6855)	(1.1385)	(1.9790)
Polynomial Order	0	1	1	1
Controls	No	No	Yes	No
Observations	1,667	1,667	1,667	1,667
R-Squared	0.0004	0.0011	0.5686	
B. 20 year olds				
Discontinuity	0.2612	-1.3040	0.0752	-2.0963
	(0.8786)	(1.7833)	(1.1801)	(1.9717)
Polynomial Order	0	1	1	1
Controls	No	No	Yes	No
Observations	1,654	1,654	1,654	1,654
R-Squared	0.0001	0.0007	0.5332	
C. 21 year olds				
Discontinuity	-0.6105	-1.4288	0.6868	-.6156
	(0.8870)	(1.8289)	(1.2800)	(2.0316)
Polynomial Order	0	1	1	1
Controls	No	No	Yes	No
Observations	1,565	1,565	1,565	1,565
R-Squared	0.0003	0.0008	0.5087	

Columns (1) to (3) report robust standard errors in parentheses (clustered by xwaveid).
Columns (4) reports bootstrapped standard errors (clustered by xwaveid, reps=1000).
* Significant at 10 per cent.
** Significant at 5 per cent.
*** Significant at 1 per cent.

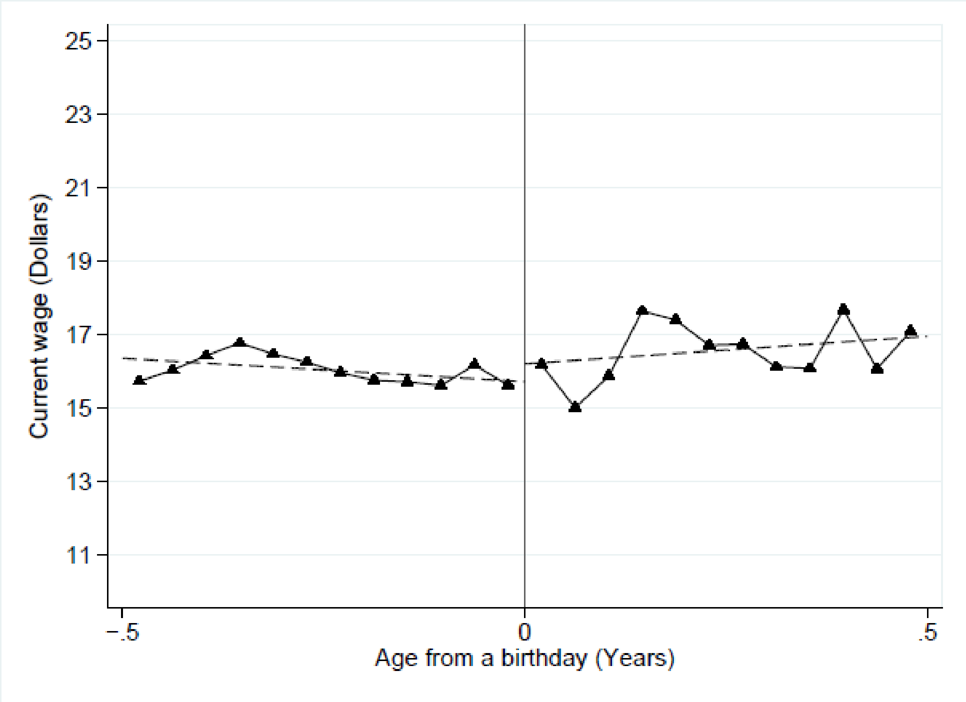

Discontinuity at threshold = 0.4874; with std. err. = 0.4380

Figure 7: The effect of an increase of award wages on actual wages for 22 year olds. Sample is restricted to those youths who earn a regular wage or salary from an employer.

Discontinuity at threshold = 0.0261; with std. err. = 0.0504

Figure 8: The effect of an increase in award minimum wages on year 12 achievement for 19 year olds. Sample is restricted to those youths who earn a regular wage or salary from an employer.

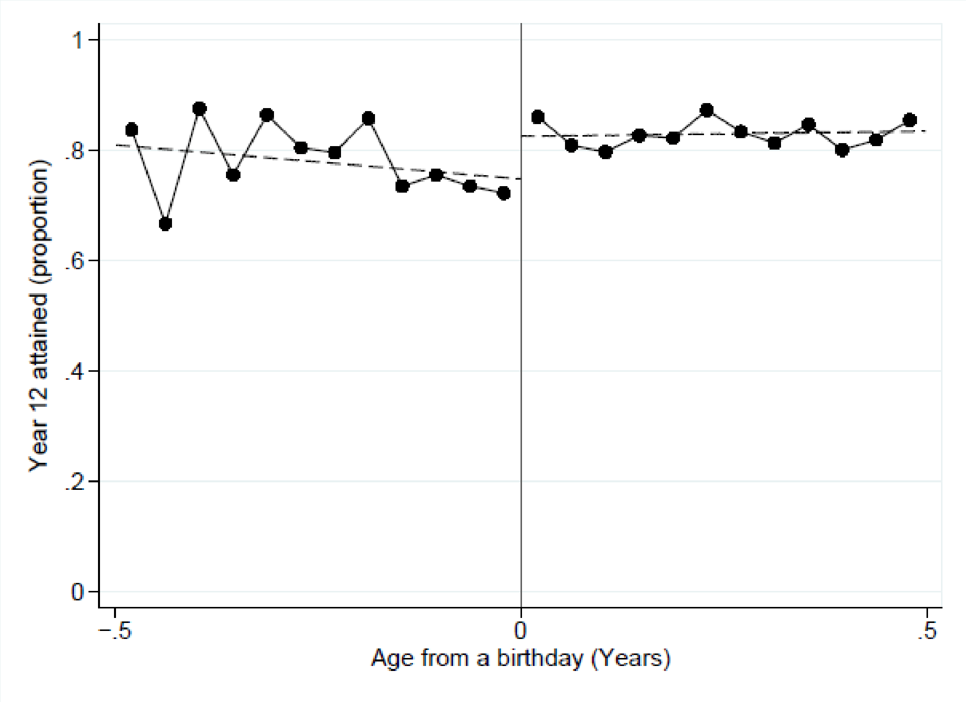

Discontinuity at threshold = 0.0773; with std. err. = 0.0485

Figure 9: The effect of an increase in award minimum wages on year 12 achievement for 20 year olds. Sample is restricted to those youths who earn a regular wage or salary from an employer.

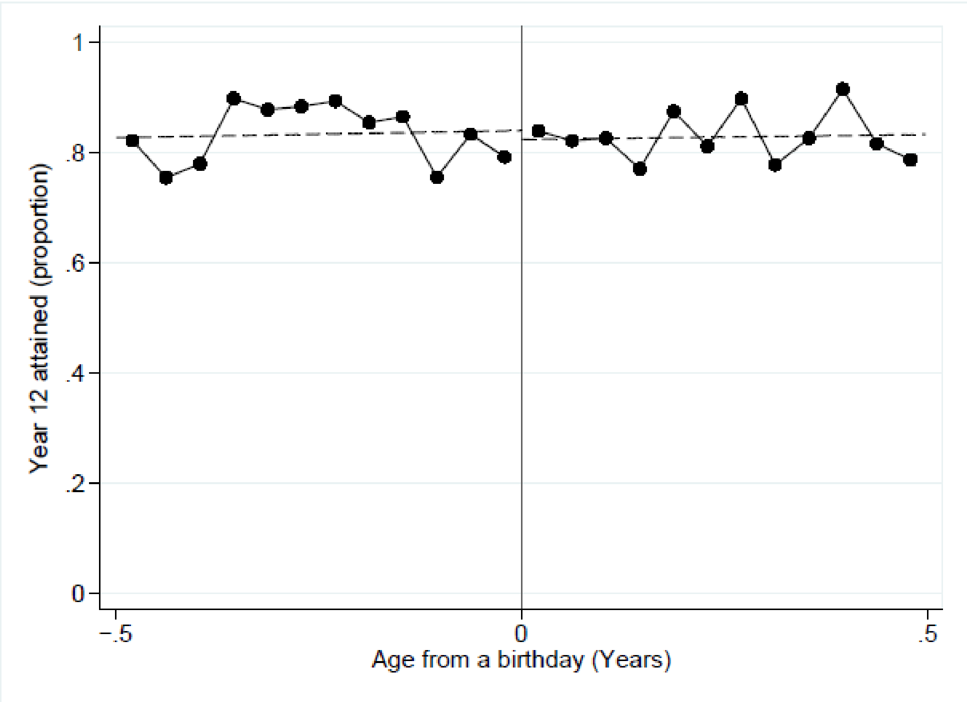

Discontinuity at threshold = -0.0161; with std. err. = 0.0459

Figure 10: The effect of an increase in award minimum wages on year 12 achievement for 21 year olds. Sample is restricted to those youths who earn a regular wage or salary from an employer.