

DISCUSSION PAPER SERIES

IZA DP No. 10727

**Does Reducing Unemployment Benefits During
a Recession Reduce Youth Unemployment?
Evidence from a 50% Cut in Unemployment
Assistance**

Aedín Doris
Donal O'Neill
Olive Sweetman

APRIL 2017

DISCUSSION PAPER SERIES

IZA DP No. 10727

Does Reducing Unemployment Benefits During a Recession Reduce Youth Unemployment? Evidence from a 50% Cut in Unemployment Assistance

Aedín Doris

Maynooth University

Donal O'Neill

Maynooth University and IZA

Olive Sweetman

Maynooth University

APRIL 2017

Any opinions expressed in this paper are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but IZA takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The IZA Institute of Labor Economics is an independent economic research institute that conducts research in labor economics and offers evidence-based policy advice on labor market issues. Supported by the Deutsche Post Foundation, IZA runs the world's largest network of economists, whose research aims to provide answers to the global labor market challenges of our time. Our key objective is to build bridges between academic research, policymakers and society.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

Does Reducing Unemployment Benefits During a Recession Reduce Youth Unemployment? Evidence from a 50% Cut in Unemployment Assistance¹

We use administrative data to examine the effect of a 50% benefit cut for young unemployed workers in Ireland during the Great Recession. Because the cut applied only to new benefit claims, claimants whose unemployment start dates differed by a matter of days received very different benefits; we exploit this fact in our Regression Discontinuity and Difference-in-Difference analyses. While we find no impact on unemployment duration for those aged 20–21, the benefit cut significantly reduced duration for 18 year olds, with an estimated elasticity close to one. We consider possible explanations for our findings and also examine long-run effects.

JEL Classification: J64

Keywords: unemployment assistance, labour supply, regression discontinuity

Corresponding author:

Donal O'Neill
Rhetoric House
Maynooth University
Maynooth, Co. Kildare
Ireland
E-mail: donal.oneill@nuim.ie

¹ We are grateful to Terry Corcoran (DSP) for providing us with the DSP longitudinal data used in our analysis and for many useful discussions in relation to this research. We are also grateful to Tim Callan, Paul Devereux, Paul Redmond and participants at the Labour Market Council conference on program evaluation (Dublin Castle February 2017) and seminar participants at Maynooth University for helpful comments on an earlier version of this paper.

1. Introduction

While no age group was spared the effects of the Great Recession, younger workers were hardest hit, with unemployment rates for 15-25 year olds exceeding 30% in some OECD countries (van Ours 2015). There is strong evidence that suggests that unemployment when young has particularly adverse long-run effects, especially for disadvantaged youths (Bell and Blanchflower 2011). As a result, policies aimed at tackling youth unemployment have become a key priority of policymakers in recent years and reforms of the unemployment benefit system have been prominent in these discussions (OECD 2010). Proposals include reductions in benefit generosity to improve work incentives (OECD 1994) and stronger enforcement of job search and training requirements, enforced by the threat of benefit sanctions, in return for adequate support (OECD 2013).

While there is a large literature on the labour supply effects of unemployment benefit reforms, much of this work focuses on the responses of prime-age workers and was carried out before the Great Recession. Schmieder et al. (2016) call for further studies encompassing broader groups of workers. Despite the substantial interest in youth unemployment, little is known about the impact of unemployment benefit reform on young claimants during severe economic downturns. In this paper, we examine the labour market response of claimants aged 18-21 to a substantial cut in unemployment assistance payments introduced in Ireland in response to the Great Recession. To do this, we use a quasi-experimental approach that exploits the fact that only new claimants were affected by cuts that reduced the weekly benefit rate from €204.30 to €100, a reduction of over 50%. As a result, people whose unemployment start dates differed by a matter of days were subject to very different benefit rates.

To carry out the analysis, we use administrative data on welfare duration covering every new unemployment claim initiated between 2007 and 2014. These data provide the start and end dates of every unemployment spell that commenced during this eight year period. In addition, the data contain information on earnings, as well as the destination states for completed unemployment spells. To identify the causal effect of the benefit cuts, we use both Regression Discontinuity and Difference-in-Difference approaches. The ability to combine the clean quasi-experimental nature of the intervention with rich administrative data on the entire population of claimants at a time of

substantial changes to the benefit system provides a unique opportunity to identify the impact of benefits cuts on young people during the Great Recession.

Despite the substantial reductions in benefits that were implemented, we find no evidence that these cuts affected the unemployment durations of those aged 20-21. However, for the very youngest claimants, we find robust evidence that the benefit cut substantially shortened unemployment durations. For 18 year olds we find that the cut in benefits reduced unemployment durations by over a year, implying a significant duration elasticity of 1.04. The corresponding elasticity for 19 year olds is similar at 1.08; however, the treatment effect is not precisely estimated for this group.

We begin in the next section with a review of the literature analysing the impact of benefit changes on labour market outcomes. Section 3 outlines relevant features of the Irish welfare system and describes the changes made by the government to the system during the Great Recession. Section 4 discusses the econometric specification and identification assumptions used in our analysis, while Section 5 describes our data in more detail. Our main results are presented in Sections 6. In Section 7 we examine the relative importance of alternative destination states in explaining our overall result and provide a competing risk decomposition of the total effect. These results also provide the basis for a discussion of possible explanations for the differences in treatment effects across age groups. In Section 8 we examine the long-run effects of the benefit cut. Section 9 concludes our analysis.

2. Literature Review

While there has been some analysis of the effects of active labour market policies on youth unemployment in a number of countries (e.g. Fougere et al. 2000, Jensen et al. 2003, Carling and Larson 2005, Bell and Blanchflower 2011, Banerji et al. 2014), our focus is on the effectiveness of benefit cuts as a policy measure. In the standard static labour supply model, cuts in unemployment benefits shift an individual's budget constraint, resulting in an income effect that reduces an individual's reservation wage rate, thus increasing exits out of unemployment into employment. The effect of benefits on unemployment duration can also be analysed using a job search model of unemployment (Mortensen 1977). In a simple version of this model, job seekers

receive offers from a known cumulative wage distribution. The arrival rate of offers depends on a worker's productivity and the general state of the economy, as well as on how hard the job-seeker searches. Unemployed individuals receive an unemployment payment, which they lose once they start working. The optimal search strategy consists of a level of effort that depends negatively on benefits and positively on the state of the economy and a reservation wage that depends positively on both of these factors. In this model, the effect of a cut in benefits is difficult to identify if macroeconomic conditions are changing at the same time.

There is a large body of empirical work looking at the effect of unemployment payments on unemployment duration. Atkinson and Micklewright (1991) and Layard et al. (1991) provide summaries of early empirical work in this field, with Layard et al. (1991) noting that the elasticity of duration with respect to benefits typically ranged from 0.2 to 0.9. However, much of this early work relied on cross-sectional variation in benefit receipt; this approach may be biased if there are unobserved characteristics that are correlated with both benefit receipt and unemployment duration. To avoid the potential endogeneity of benefit receipt, more recent work has tended to exploit natural experiments that arise following changes to benefit rates and/or the duration of payments. Given that our analysis concerns benefit cuts, we focus on the results from studies examining changes to benefit rates. Krueger and Meyer (2002) and Tatsiramos and Van Ours (2014) provide summaries of this work. Most of the papers cited report estimates between 0.5 and 1.0 for the elasticity of duration with respect to benefits. However, some studies find estimates outside this range. For example, Hunt (1995) finds no significant effect of a benefit cut on the overall probability of exiting unemployment in Germany. On the other hand, Carling et al. (2001) report a benefit elasticity of 1.6 for Sweden. This effect is large compared to earlier findings and the overall effect is driven by particularly large effects for those aged less than 25. More recently, Card et al. (2015a) and Kyyra and Pesola (2017) use a regression kink design to estimate benefit elasticities in Austria and Finland respectively. Both studies report elasticities in the range of 1.5 to 2 but many of their reported elasticities are imprecisely estimated and not robust to changes in specification.

There is some evidence that benefit elasticities tends to vary with economic conditions. For example, Arulampalam and Stewart (1995) conduct separate analyses for cohorts entering unemployment in 1978, during a period of low unemployment, and in 1987, during a recession,

and find that benefits have a much lower effect on unemployment duration during the recession. More recently, Kroft et al. (2016) estimate a benefit elasticity of 0.99 when the state unemployment rate is below the US national average, but 0.28 when it is above the average.

To date, only a small number of papers have considered the impact of benefit cuts during the Great Recession. For example, Rebello-Sanz and Rodriguez-Planas (2016) examine the impact of a reduction in the replacement rate in Spain in 2012. They find that the reform reduces mean unemployment duration by 5.7 weeks, implying an elasticity of 0.86. Card et al. (2015b) use administrative data for the state of Missouri over the period 2003-2013. In contrast to the work of Kroft et al. (2016), they find that unemployment durations became more responsive to benefit levels during the Great Recession, with an elasticity of 0.65- 0.9 during the recession compared to about 0.35 pre-recession.

The empirical work discussed above concerns the effects of unemployment insurance (UI) payments rather than the unemployment assistance (UA) payments that are the focus of our paper. UA payments differ from UI payments in potentially important ways. Firstly, they are not time limited. In general, this would make claimants more responsive to cuts in UA, since the expected value of the benefits foregone will be higher than if it were time limited. However, in the case being examined in our paper, where payments were reduced only for young claimants, the opposite time profile is implied. For example, the large benefit reforms we consider only applied to 18-21 year olds. A 21 year old may have expected their payments to increase once he reached his 22nd birthday. This time profile will raise the reservation wage above what it otherwise would have been and so may attenuate the effect of the benefit cut.

In addition, UA payments are typically paid to claimants with less favourable characteristics: individuals who have exhausted insurance-based payments and those who have insufficient insurance contributions to qualify for them. Therefore, the individuals who qualify for assistance payments tend to be younger, lower educated and have less labour market experience. This may again lead to UA benefit effects that differ from those found for UI reforms. First, low skilled workers have little scope to add lower segments of the labour market to their job search possibilities as the spell progresses, so even if their reservation wage falls, they have less capacity to change their behaviour in a way that improves their exit probabilities. On the other hand,

younger workers tend to face a wage offer distribution with a low variance, and this can be expected to lead to larger benefit responses (Narendranathan et al., 1985); any fall in the reservation wage will tend to have a larger impact on the probability of leaving unemployment if situated in a dense part of a wage offer distribution because it brings in many more potential job offers. These differences in wage offer distributions may be more pronounced when there is an age-related sub-minimum wage, resulting in distributions that are truncated at different thresholds for workers of different ages.

There has been some analysis of the effects of social assistance (SA) on the duration of unemployment using the natural experiment approach. UA and SA are similar to the extent that they are of open-ended duration and means-tested, but unlike UA, receipt of SA tends not to be conditional on unemployment. Bargain and Doorley (2011) study the effect of the French *Revenu Minimum d'Insertion* (RMI), an income maintenance payment available to all individuals aged over 25.² Using a regression discontinuity approach, they find significant effects for single unskilled men, with employment rates falling by 7-10 percentage points at the 25 year-old threshold. They find no significant effect for higher-skilled men.

Lemieux and Milligan (2007) examine the effect of SA on labour supply in Quebec. Like Bargain and Doorley (2011), they use a regression discontinuity approach, exploiting an entitlement threshold at age 30 for identification. They find that entitlement to the benefit reduces the probability of employment by 3-5 percentage points³. Fortin et al. (2004) also estimate the effect of SA on labour supply in Quebec, this time using a 1989 reform that removed the age threshold, thus increasing the payment for those aged under 30 by 145%. They find significant results for those aged 22-29, with a duration elasticity of around 0.25. The effect for 18-21 year olds is not statistically significant. However, identification for this group is complicated by other reforms implemented at the same time.

Finally, Walsh (2015) provides an initial evaluation of the benefits cuts that we analyse. In contrast to our analysis, he has to rely on survey data with a limited panel component, making it difficult to follow claimants over a given spell. In addition, given the nature of his data, Walsh

² For a related analysis of this programme see Chemin and Wasmer (2014).

³ A related study on Denmark by Jonassen (2013) reports findings in line with Lemieux and Milligan (2007) and Bargain and Doorley (2011).

cannot identify directly those claimants who were eligible for benefits, nor can he identify which claimants were subjected to the cuts, making it difficult to identify the treatment effect of interest. He finds no evidence of a higher rate of transition from unemployment to employment for those affected by the cuts in the rate of payment of UA; however, he is careful to say that the results are suggestive rather than definitive, given the drawbacks of the data available to him.

3. The Irish Welfare System and the Great Recession

The Irish unemployment benefit system consists of two types of payments: Jobseeker's Benefit and Jobseekers Allowance. Jobseekers Benefit (JB) is a UI payment given for up to nine months to claimants who satisfy specific insurance contribution conditions. Claimants who have exhausted their entitlement to JB or who have not accumulated sufficient contributions to be eligible for JB are entitled to apply for Jobseeker's Allowance (JA), a UA payment that is means-tested and payable indefinitely provided the claimant continues to be available for work.

A relatively unusual feature of the Irish benefit system prior to 2009 was that all qualified individuals aged over 18 were entitled to the full JA payment, even if they had never worked. Scarpetta et al. (2010) report that in two thirds of OECD countries, school leavers are not eligible for unemployment payments unless they have worked a certain period of time, typically one year. Even in countries that do allow UA payments for young job seekers, most do not pay the full adult rate.

Data from the Irish Central Statistics Office (CSO) show that in January 2007, 31,000 people aged less than 25 were registered as unemployed. By October 2013, this had more than doubled, increasing to 64,700. Although JB receipt is not linked to a person's age, younger workers are less likely to have accumulated sufficient insurance contributions and so are more likely to be on JA. Throughout the period 2007 to 2014, the majority of claimants under 25 were in receipt of JA, rising from about 60 per cent in early 2009 to over 90 per cent in 2014.

Ireland was one of the countries worst affected by the Great Recession, with unemployment rising from 4.5% in 2007 to 12.2% in 2009 and peaking at 15% in 2012. The effects of the financial crisis felt elsewhere were compounded in Ireland by the bursting of a property bubble and the near-collapse of the banking system. A combination of falling tax revenue from the construction sector

and a decision to guarantee all bank liabilities resulted in the government facing severe borrowing difficulties, which lead to the introduction of draconian austerity measures⁴.

Of particular focus in this paper is a series of substantial reductions in JA paid to younger workers. The stated rationale for the cuts given by the Government was to “ensure that young people are better off in education, employment or training than claiming.”⁵ However, the necessity of cutting spending in order to reduce the government deficit also played an important role in the timing of these cuts. In 2009, total spending on JA amounted to about €2b, accounting for 3.25% of total public expenditure. It is clear that reductions in JA can generate significant savings to the exchequer even in the absence of behavioural changes.

The benefit cuts are detailed in Table 1. Prior to the cuts, all JA claimants were paid a basic rate of €204.30 a week. The first in the series of cuts was targeted at claimants aged 18 and 19, who had their weekly rate cut to €100 in May 2009. Further cuts followed in January 2010, when workers aged 20 and 21 also had their benefit cut to €100.⁶ The cuts only applied to claimants who entered after the date of the legislation, with claimants entering prior to the legislation remaining on the old rate. As a result, people whose unemployment start dates differed by a matter of days were subject to very different benefit rates. In addition to exemptions for existing claimants, new claimants were exempted from the cuts if they had a dependent child, if they had had a spell of unemployment in the previous 12 months or if they were transferring from Disability Allowance. Given the nature of these conditions, the proportion of eligible claimants exempted from the benefit cuts varied across age groups. We explicitly account for the possibility of exemptions in our econometric analysis.

The cuts in benefits outlined above are very large relative to many of those examined previously. For example, Carling et al. (2001) examine benefit cuts of the order of 6%, while Hunt (1995) considers cuts of between 3% and 7%. The cuts of 51% implemented in Ireland during the Great Recession are almost an order of magnitude bigger than these. In addition to their size, the

⁴ Ireland subsequently sought and accepted a rescue package from the Troika of the EU, ECB and IMF but the policy measures analysed in this paper predate this agreement.

⁵ <http://www.welfare.ie/en/pressoffice/Pages/pr231013.aspx>

⁶ At this time, there was also a series of smaller cuts for those aged 22-25. However, many of the eligible pool in these age groups were exempt from the benefit reduction. We do not analyse these cuts here.

restriction of the cuts to new entrants has the advantage of providing quasi-experimental variation in JA rates that we exploit in order to establish the causal effect of benefits on unemployment.

4. Econometric Specification

In this paper, we use two identification strategies to estimate the causal impact of unemployment benefit on unemployment duration. We use both a Regression Discontinuity (RD) approach and a Difference-in-Difference (DiD) estimator, both of which exploit the fact that the cuts applied only to new entrants after a well-defined date. The RD approach is used to estimate the overall effect, while the DiD approach is used in combination with a hazard model to look at the impact of the benefit cuts on the timing of exits. These approaches are discussed in more detail below.

4a. Regression Discontinuity Design

Regression Discontinuity (RD) Design is a well-established and popular approach for identifying causal effects in economics (for a review, see Imbens and Lemieux 2007). The idea behind RD is that assignment to the treatment is determined either completely (sharp RD) or partly (fuzzy RD) by the value of a predictor or running variable (S) being on either side of a fixed threshold (s_0). A key requirement of the RD approach is that the probability of receiving treatment jumps discontinuously at the cut-off, thus inducing variation in treatment status that is uncorrelated with potential confounding variables. The running variable may be associated with potential outcomes provided this relationship is smooth. Under these assumptions, any discontinuity in the estimated relationship between the running variable and the outcome at s_0 is interpreted as evidence of a causal effect of the treatment.

Formally, let $Y(1)$ and $Y(0)$ denote the potential outcomes associated with and without treatment, in our case unemployment duration. With a sharp RD design, unit i is assigned to the control group if $S_i < s_0$ and to the treatment group if $S_i \geq s_0$. We are interested in estimating the average treatment effect at the threshold:

$$\alpha = E[Y_i(1) - Y_i(0) | S_i = s_0] \quad (1)$$

Under mild continuity conditions (Hahn et al. 2001), this estimand is identified, with

$$\alpha = \alpha_+ - \alpha_-$$

where $\alpha_+ = \lim_{s \downarrow s_0} \alpha(s)$, $\alpha_- = \lim_{s \uparrow s_0} \alpha(s)$, $\alpha(s) = E[Y_i | S_i = s]$.

Following much of the literature (e.g. Gelman and Imbens 2014), we estimate α using kernel-based local linear regressions on either side of the threshold. The estimation of these local linear regressions is facilitated by two key features of our data that make it ideal for our analysis: having access to the population of claimants, resulting in a large number of observations; and having the exact start date of every claim, allowing specification of the running variable in days rather than in weeks or months. We discuss these issues in more detail in Section 5.

In choosing the bandwidth for the local linear regression, there is a trade-off between bias and efficiency. In our analysis, we follow the literature and choose a triangular kernel and the Mean Squared Error optimal bandwidth suggested by Calonico et al. (2014). We also examine the sensitivity of our results to alternative choices of the bandwidth, namely half the optimal bandwidth and twice the optimal bandwidth, as well as the non-bias-adjusted optimal bandwidths proposed by Imbens and Kalyanaraman (2011).

In the sharp RD design discussed above, receipt of treatment is a deterministic function of the running variable. The approach is easily adapted to situations where the running variable causes a discontinuity in the probability of receiving the treatment rather than a deterministic switch, resulting in a fuzzy RD design. In this case, the running variable acts as an instrumental variable for treatment status. The resulting estimator is a Wald estimator in which the estimated discontinuity in outcomes at s_0 is divided by the corresponding discontinuity in the probability of treatment. Hahn et al. (2001) discuss the identification conditions needed in a fuzzy RD design. In addition to the continuity assumptions needed for the sharp design, an additional independence assumption is needed: the running variable must only affect outcomes through its effect on treatment status.

In our analysis, we exploit the fact that in the Irish reform, people entering unemployment prior to a fixed date are exempt from the reform, while many of those entering after that date receive the lower benefits. Because of the availability of administrative data, we can record the

exact date on which the unemployment spell began. Therefore, our running variable measures the recorded time in days between when an individual enters unemployment and the date the legislation is implemented; the running variable is coded as negative for people entering before the legislation and positive for those entering after the legislation. In this way, the RD approach identifies the treatment effect using people entering within a very small window of the threshold. However, people entering at other times are used to estimate the relationship between the running variable and outcomes on either side of the threshold and in this way contribute to the final estimate. For the RD approach, the outcome variable is the duration of unemployment in weeks and we use data six months before and six months after the reform to estimate the RD parameter.

4b. Hazard Functions

When examining unemployment durations, it is quite common to conduct the analysis in terms of hazard functions, which provide information on the timing of exits out of unemployment. As a complement to the RD approach outlined in the previous section, we also use the hazard functions approach to examine the impact of the benefit cut on the timing of exits from unemployment. We follow previous work (Meyer 1990) and specify a continuous time reduced form proportional hazards model with a flexible baseline hazard:

$$h_i(t) = h_0(t)\exp[\mathbf{X}_i(t)' \boldsymbol{\beta}] \quad (2)$$

where $h_0(t)$ is the baseline hazard at time t , $\mathbf{X}_i(t)$ a vector of possibly time-varying covariates for individual i at time t and $\boldsymbol{\beta}$ a vector of unknown parameters.

For a sample of N individuals, the likelihood function can be written as:

$$L(h_0, \boldsymbol{\beta}) = \prod_{i=1}^N \{1 - \exp(-\exp(\mathbf{X}_i(t_i)' \boldsymbol{\beta}) \cdot \gamma(t_i + 1))\}^{c_i} \exp\left(-\sum_{d=1}^{t_i} [\exp(\mathbf{X}_i(d-1)' \boldsymbol{\beta}) \cdot \gamma(d)]\right)$$

where c_i is a censoring indicator with $c_i = 1$ for a completed (uncensored) spell and zero otherwise and $\gamma(d) = \int_{d-1}^d h_i(u) du$. The likelihood function is maximised with respect to $\gamma(d)$ and $\boldsymbol{\beta}$ under the restriction that the baseline hazard pieces $\gamma(d)$ are non-negative.

The key to our empirical approach is the specification of $X_i(t)' \boldsymbol{\beta}$. As with the RD approach, we compare individuals entering before and after the legislation. For example, consider the benefit cut introduced for 18 year olds on May 1, 2009. We regard 18 year olds who commenced a spell during the month of May 2009 as the treatment group and 18 year olds who commenced a spell during April 2009 as the control group. To account for any seasonal effects that may cause durations for those entering in April to differ from those entering in May, we adopt a Difference-in-Difference specification, which also includes spells from the same months in 2008. Specifically, we estimate:

$$\mathbf{X}_i(t)' \boldsymbol{\beta} = \mathbf{Z}_i(t)' \boldsymbol{\theta} + \alpha_1 \text{May}_i + \delta D_{2009,i} + \phi \text{May}_i D_{2009,i} \quad (3)$$

$\mathbf{Z}_i(t)$ is a vector of covariates including nationality, education and a constant; May_i is a dummy variable indicating entry to unemployment in the month in which the treatment was introduced; $D_{2009,i}$ is a dummy variable indicating entry into unemployment in the treatment year. The parameter of interest is ϕ , which measures the change in the hazard resulting from the cut in benefit payments to 18 year olds.

5. Data

To carry out our analysis, we use the Jobseekers Longitudinal Database (JLD) provided by the Department of Social Protection (DSP), the government department responsible for the benefit system. This is an administrative dataset that includes every claimant who received a JA or JB payment from 2004. Having access to the entire population of claimants provides sufficient observations to conduct the non-parametric analyses required for the RD approach.

The data provide records for the exact start and end date of every new claim, allowing us to calculate the duration in days for the entire population of new JA claims. The age of each individual at the start of their claim is also recorded. The availability of both age and the start date of the spell allows us to identify whether an individual was in one of the groups targeted by the benefit cuts considered in this paper. As discussed earlier some individuals, such as claimants with dependent children, were exempt from these benefit cuts. Crucially, the data provided by the DSP identifies whether or not a claimant was actually subject to a cut. Furthermore, knowing the exact start date of the unemployment claim means that the running variable in our RD analysis can be

defined in days. As noted by Lee and Card (2008), this can substantially reduce specification error compared to cases in which the running variable is only available in coarse intervals.

For claims that have ended, the data also include information on the destination state, allowing us to consider competing risk explanations of our findings. We also have information on an individual's gender and nationality. In addition, the DSP have collated the available data on education for the majority of individuals in the JLD. We control for these factors in our analysis. For this population of claimants, the data also include information on annual earnings and weeks worked for every year in which the individual worked. We use these earnings data when discussing potential explanations for our findings.

In total, between 2007 and 2014, there were 278,065 claims by individuals aged 18-21, of which 83% were for JA. For the analysis in this paper, we focus mainly on young male claimants, the group most affected by the Great Recession. However, we also report key results for women for comparison. As noted earlier, for the RD approach we use data for six months before and after the reform to estimate the RD parameter. As a result, the number of claims used in this analysis ranges from about 4,000 to over 6,000. For the DiD analysis we use two months of data for two consecutive years, resulting in the number of observations varying from about 1,000 to 1,400.

Summary statistics for the groups used in our RD analysis are provided in Table 2. Across all groups, about 95% of claimants have Irish nationality. The variable labelled "Low Education" denotes that the claimant did not finish second level education. In the Irish system, approximately 15% of recent cohorts of school leavers did not complete second level education. The figures in Table 2 show that for our sample of JA claimants, this number is much higher, indicating that, as expected, JA claimants are much less educated than their peers. The variable labelled "No Previous Employment Spell" denotes that the claimant did not have a job prior to the unemployment spell of interest. The proportion with no prior work experience clearly falls with age, going from 40-52% for 18 year olds to 12-14% for the 21 year olds. The fact those who enter unemployment later in 2009 are somewhat better educated and are less likely to have had a prior job reflects the deterioration in the labour market during that year.⁷

⁷ We consider the impact, if any, of these differences for our analysis later in the paper.

The treatment status variable indicates that the claimant was subject to the legislated benefit cut. For 18 year-old claimants, a substantial majority were subject to the cut. The proportion affected decreases with age because, as discussed earlier, more older claimants qualify for the exemptions specified in the legislation. The table also indicates that a small number of people are recorded as having their benefit cut before the legislation came into effect. This appears to be due to short delays in processing claims. Finally, the table shows the average spell duration for each group. The first point to note is how long unemployment spells were during this period. For all groups, average unemployment duration was over a year and for the youngest group it was approximately two years. This reflects the depressed nature of the Irish labour market at this time. It also noteworthy that for each age group, average durations were shorter for those entering after the legislation, suggesting a potential effect of the cuts. In the remainder of the paper, we examine whether these differences represent causal effects.

6. Results

6a. Regression Discontinuity Design

As noted earlier, we analyse four welfare changes in total: the benefit cuts for 18 and 19 year olds in 2009 and the cuts for 20 and 21 year olds in 2010. Initial results for the RD analysis are shown in Figures 1a-1d. These graphs provide an exploratory visual description of the RD design, prior to the more formal analysis outlined in Section 4a. The figures show regression discontinuity plots for each of the age groups, where the running variable is days before or after the introduction of the benefit cut. Each figure consists of two graphs. The left hand ones are regression discontinuity plots of treatment status, where the treatment variable takes the value one if a claimant was subject to the legislated cut and zero otherwise. The points represent the proportion treated within each of fifty equally spaced bins on either side of the threshold. To assist with visual interpretation, we also show estimates of global fourth order polynomials fitted to these data. These higher order polynomials are simply an exploratory visual aid; the statistical inference conducted later follows recommended procedures, estimating the discontinuity using local linear regressions. The vertical lines indicate the RD threshold given by the date of the legislation⁸.

⁸ Looking at the treatment status of claimants, it appears that the effective date in 2009 for the benefit cut for 18 and 19 year olds was April 29 rather than May 1. 95% of those who began their claim on April 29 and 87% of those who began on April 30 had their benefit cut by the legislated amount. These percentages are similar to the 86% who had their benefit cut on May 1 and compare to 14% and 18% of those who commenced their spell on April 27 and 28 respectively. The DSP have indicated that the most likely

Examination of these graphs allows us to explore the bite of the legislation for the relevant treatment, which feeds into the denominator of our fuzzy RD estimator. The right hand graphs present RD plots of unemployment duration, with each point now representing average duration within a bin. These graphs illustrate the change in unemployment duration upon introduction of the legislation and provide the basis for the estimated numerator of the fuzzy RD.

Looking first at the RD plots for treatment status, we see clear evidence of a discontinuity at the threshold in each case. However, as discussed above, the bite of the legislated cut clearly varies with age. The graphs suggest that the likelihood of treatment increased by between 60 and 70 percentage points for 18 year olds, by about 50 percentage points for 19 year olds and by 20 to 30 percentage points for 20 and 21 year olds. While the differential bite of the various treatments is taken into account in the fuzzy RD so that the estimates of all treatment effects are consistent, the smaller bite of cuts for older workers makes precise estimation of the effects for these groups more difficult.

Turning to the RD plots for unemployment duration, we see very little evidence of a causal effect of the benefit cuts on unemployment duration for 20 and 21 year olds. However, for 18 and 19 year olds, the RD plots suggest that unemployment durations fell substantially when benefits were cut, with the effect particularly pronounced for 18 year olds. For this group, there is a forty week reduction in unemployment duration at the threshold.

While the RD graphs provide an easy visual presentation of the RD design, a more formal analysis is needed to establish the statistical significance of the causal effects. The results of this analysis are given in Table 3.⁹ As noted in Section 4a, when estimating these effects we follow the recent literature and estimate local linear regressions to the left and right of the threshold and report the results for the optimal bandwidth proposed by Calonico et al. (2014).¹⁰ The first row provides the estimated effect of the legislation on the likelihood of receiving a benefit cut, while the second row provides the fuzzy RD estimate of the causal effect of the benefit cuts on

explanation was extended claim processing times due to the large inflow of young claimants during this time period. To allow for this, we choose April 29 as the threshold for our 18 and 19 year olds. For 20-21 year olds, we set January 4 2010 as the threshold. Since January 1 was a national holiday, very few claimants were processed from January 1-3. For all age groups, we have experimented with a number of other reasonable thresholds and the key results are robust to these alternatives.

⁹ Corresponding results for women are presented in Table A1 of the Appendix. The results are similar to those for men.

¹⁰ We have also estimated all our models using twice and half the optimal bandwidth, as well as the optimal nonbias-adjusted bandwidths proposed by Imbens and Kalyanaraman (2011). The results discussed below are robust to the choice of bandwidth.

unemployment duration. These findings confirm the results of the RD graphs. For those aged 20 and 21, we find no effect of benefit cuts on unemployment duration. However, there is evidence of a strong negative effect for 18 and 19 year olds. For these two groups of claimants, we find that the benefit cut resulted in unemployment durations falling by 61 weeks and 50 weeks respectively. However, only the 18 year-old effect is statistically significant.

We can use the RD results to estimate a benefit duration elasticity for each of these groups and these are reported in the last row of Table 3. For both 18 and 19 year olds, the falls in duration combined with the 50.9% reductions in benefits imply an elasticity of just over one. This estimate is consistent with the range of estimates reported in the previous literature. The elasticities for 20 and 21 year olds are much smaller and, given the RD results, statistically insignificantly different from zero. As discussed earlier, the age variation in the estimated effects may reflect the fact that the older age groups were more likely to be exempt from the benefit cut but may also reflect behavioural differences across age groups. We investigate this further in Section 7. However, we first examine the robustness of the estimated effect on 18 year olds.

6b: Robustness Checks

Two robustness checks are typically carried out to examine the validity of the RD design assumptions. The first repeats the analysis for years in which there is no treatment. If the identification strategy is valid, we should observe no effect in these years. The second focuses on the year the legislation is implemented but examines alternative thresholds that do not correspond to the legislation date. Again, in the absence of any other treatment, we should observe no effect at these alternative thresholds. In our analysis, we consider an additional third robustness check. Since JB claimants were not subject to the benefit cuts, we should see no significant effect at the threshold for this group. Here we examine the robustness of our results for 18 year olds using all three approaches.

To carry out the first check, we focus on 18 year-old claimants entering in 2008 and 2010, years in which there was no legislative change. The resulting RD plots for unemployment duration are shown in Figure 2. In contrast to the results for 2009, there is no evidence of a reduction in duration at the May threshold in either 2008 or 2010. In both years, the point estimates using the optimal bandwidth are small and statistically insignificant.

To examine the sensitivity of our results to alternative 2009 thresholds, we repeat the analysis using thresholds that are a month earlier (April 1) and a month later (June 1). The RD plots are given in Figure 3. Again, in contrast to the RD plot when the correct threshold is used, these alternative thresholds give no indication of a discontinuity. The point estimates are -0.63 and -1.55, respectively, compared to the point estimate of -61 obtained with the May threshold.

The RD plot for the population of 18 year old JB claimants in 2009 is given in Figure 4. As expected, given the time-limited nature of JB, the durations of these claims are much shorter than the durations of JA claims. However, in contrast to the findings for JA claimants, we see no evidence of a discontinuity in JB durations at the May threshold. This suggests that the benefit cut for JA claimants did not have a spill-over effect on JB claimants, who were not subject to the legislation.¹¹

These robustness checks support the identifying assumptions underlying our RD estimation.¹² Following Calonico et al. (2016), we have also considered the impact of including covariates in our RD analysis to account for any compositional changes around the threshold. Controlling for previous employment history, education and nationality of the claimants had very little effect on our results.

6c. Effects at the extensive margin

As mentioned earlier, one of the stated aims of the benefit cuts was to ensure that young people were better off in education than in unemployment. Accordingly, it is possible that the benefit cuts had effects at the extensive margin, reducing the numbers entering unemployment by encouraging young people to stay in school. To examine this, we adjust the RD design used above and check for discontinuities in the density of the running variable itself, rather than in unemployment durations.¹³ If 18 year olds remained in education longer following the reduction in benefits, we would expect to see a discontinuous fall in the density of those entering

¹¹ This is in contrast to Levine (1993) who found that the generosity of UI benefits in the U.S. appeared to decrease the unemployment duration of those who did not receive UI.

¹² We also carried out all these robustness checks for women and found results similar to men in almost every case. The exception was the validity check using claimants entering in 2008. For women, we found a significant RD effect in 2008. However, this was due to a small number of relatively high durations for those entering in the weeks before the threshold and the significance of the effect was not robust to the choice of bandwidth.

¹³ McCrary (2007) suggests carrying such a test to check the validity of the RD assumptions. However, in our set-up the test has added independent interest, capturing possible effects of the treatment at the extensive margin.

unemployment at the threshold. We provide the estimated density in Figure 5. The points represent the proportion of all claimants entering unemployment in each bin. We see no evidence that the reduced JA resulted in lower entries into unemployment. The estimated density is continuous at the threshold, with no statistically significant change following the benefit cut. This suggests that these cuts had no additional effect on unemployment over and above their effect on the duration of spells reported earlier.

6d. Unemployment Hazard analysis

To gain further insight into the results presented from the RD analysis above, we report the results from the DiD hazard function approach. We begin by presenting Kaplan Meier non-parametric hazard functions for the control and treatment groups, for both the pre-intervention and the intervention years. As noted earlier, for this analysis the treatment groups consist of those commencing a spell in May for 18 and 19 year olds and in January for 20 and 21 year olds. The control groups are those entering one month earlier. The intervention year is 2009 for 18 and 19 year olds and 2010 for 20 and 21 year olds. For ease of presentation we only report the graphs for two of our four treatments – 18 year olds in 2009 (Figure 6) and 20 year olds in 2010 (Figure 7). Looking first at Figure 6, we see very little difference in the hazard functions for 18 year olds entering in 2008, when there was no treatment. However, this changes dramatically in 2009 when the hazard for those entering in May, after the legislated benefit cut, is uniformly higher. 18 year olds subject to the benefit cut were more likely to leave unemployment in every week following the commencement of their spell. In contrast, looking at Figure 7, there appears to be little difference in the behaviour of the 20 year old control and treatment groups in either 2009 or 2010, suggesting once again that the treatment had very little effect on this age group.

To examine these changes more formally, we estimate the hazard DiD specification model given by equations (2) and (3). The results for all treatments are presented in Table 4. The results shown are for the proportional hazard model, specifying a quadratic in duration to capture a nonlinear baseline hazard. Looking at the control variables, it appears that nationality had little impact on the likelihood of exit; however, not surprisingly, lower educated workers were less likely to exit across all age groups. The key parameter is the coefficient on the interaction term between year and month of entry. Once again, we see a significant effect of the legislation for 18

year olds but not for any other group. The results from the estimated hazard imply that 18 years olds entering after the legislation were 42%¹⁴ more likely to exit their JA spell than those in receipt of the higher benefits.

7. Competing Risk Decomposition

In Section 6, we found that while the cut to JA had a small and insignificant effect on unemployment durations for those aged 20-21, there was robust evidence of a substantial and significant effect for 18 year old claimants. To examine these findings in more detail, we extend the previous analysis of unemployment duration by considering the state to which claimants exited. Given the depressed nature of the labour market in 2009/10, it may have been easier for claimants to move from JA to education or training than to employment so the significant effect we found for 18 year olds may be driven by exits to training. In addition, as noted earlier, since there are differences in the skill levels of 18 and 21 year old claimants, differential exits to training may potentially explain the age differences in our estimated effects.

To account for exit states, we carry out a competing risk decomposition of the difference in mean unemployment duration between the treatment and control groups. The difference in average duration between the treated and the control group is given by

$$\Delta Y = \bar{Y}_t - \bar{Y}_c$$

where t indicates treatment group and c denotes control group. In the case of three exit states denoted by 1, 2 and 3, where the proportion leaving into each of the three states for group i is given by f_{i1}, f_{i2} and f_{i3} , we can write the overall difference as

$$\begin{aligned} \Delta Y &= (f_{T1}\bar{Y}_{T1} + f_{T2}\bar{Y}_{T2} + f_{T3}\bar{Y}_{T3}) - (f_{C1}\bar{Y}_{C1} + f_{C2}\bar{Y}_{C2} + f_{C3}\bar{Y}_{C3}) \\ &\equiv \left(\frac{N_{T1}}{N_T} \bar{Y}_{T1} + \frac{N_{T2}}{N_T} \bar{Y}_{T2} + \frac{N_{T3}}{N_T} \bar{Y}_{T3} \right) - \left(\frac{N_{C1}}{N_C} \bar{Y}_{C1} + \frac{N_{C2}}{N_C} \bar{Y}_{C2} + \frac{N_{C3}}{N_C} \bar{Y}_{C3} \right) \end{aligned}$$

where N_T and N_C are the total number of claimants in the treatment and control groups respectively. N_{Tk} and N_{Ck} refer to the number exiting to state k from these groups and \bar{Y}_{ik} is the average duration for those in group i who exit to state k .

¹⁴ This is calculated as $(\exp(0.35)-1)*100$.

Suppose we observe spells over a period of D weeks. Then we can write $\overline{Y_{Tk}}$ as $\sum_{d=1}^D \frac{N_{Tk}^d}{N_{Tk}} \cdot d$, where N_{Tk}^d is the number exiting to state k from the treatment group in week d . The overall difference can then be rewritten as

$$\Delta Y = \left(\left\{ \left[\sum_{d=1}^D \left\langle \frac{N_{T1}^d}{N_T} - \frac{N_{C1}^d}{N_C} \right\rangle \cdot d \right\} \right\} + \left\{ \left[\sum_{d=1}^D \left\langle \frac{N_{T2}^d}{N_T} - \frac{N_{C2}^d}{N_C} \right\rangle \cdot d \right\} \right\} + \left\{ \left[\sum_{d=1}^D \left\langle \frac{N_{T3}^d}{N_T} - \frac{N_{C3}^d}{N_C} \right\rangle \cdot d \right\} \right\} \right)$$

The terms inside the curly brackets represent the contributions of each of the exit states to the overall difference in duration.¹⁵ In this way, the decomposition allows us assess the relative importance of alternative exit states.

In our data, there are 22 recorded exit states. When carrying out the decomposition, we follow DSP guidelines and aggregate these into four categories: work, education and training (hereafter referred to as training), inactivity and other.¹⁶ The results of the decompositions for each of our four age groups are given in Table 5. An estimate of the overall treatment effect is presented in the first row. Here the overall effect is estimated as the difference between the average duration of those entering unemployment in the month before the legislation and those entering in the month after. The size of the effects are similar to those reported in Section 7 using the RD approach.¹⁷ The remaining four rows report the contributions of each of the exit states. Looking at the results for 18 year olds, we see that no one exit state dominates the overall effect. While the contribution of exits to inactivity is relatively small, the other three exit states – training, work and “other” – all contribute substantially to the overall effect. Exits to training, while important for this group, are not the dominant determinant of the overall effect. The same three states are important when considering 19 year olds, although all of the effects are smaller than the corresponding effects for 18 year olds. The results for 20 year olds show how differences across exit states may be obscured in the overall effect: 20 year olds subject to the benefit cut exit to training more quickly than those not subject to the cut, but they are slower to exit to the “other” destination. These opposing effects

¹⁵ It is worth noting that $\frac{N_{Tk}^d}{N_t}$ is the slope of the treatment group’s Cumulative Incidence Function for exit state k at duration d (see for example Coviello and Bogges (2004) and Kalbfleisch and Prentice (2002)). For further details see O’Neill (2017).

¹⁶ Many of those in the “other” category were recorded as “no reason stated”. Some of these claimants had earnings records that suggested that they had exited to work. We experimented by allocating these claimants into the work category, but this had little effect on the reported results.

¹⁷ The difference with the RD estimate in Table 3 is largely due to the fact the treatment effect reported in Table 5 takes no account of the fuzzy nature of the treatment.

combine for a very small overall effect. Each of the individual exit state effects for 21 year olds is small, with none exceeding four weeks.

As discussed earlier, the government's stated motivation for the benefit cuts was to ensure that education, employment or training were preferable to unemployment. The results presented in Table 5 show that education, training and work all contributed to the overall reduction in unemployment durations for 18 year olds, but to a lesser extent for older claimants. We now explore possible explanations for this.

The training and education courses provided to JA claimants are aimed at providing them with the skills needed in the workplace as well as skills typically received during the senior cycle of second level education. A comparison of the characteristics of 18 and 21 year old JA claimants reveals substantial differences in skills at the start of their unemployment spells. From Table 2 we see that 45% of 18 year olds had not finished second level education, and 45% did not have a prior employment spell; the corresponding figures for the 21 year olds are 29% and 13%.¹⁸ It is clear, therefore, that the 21 year-old claimants are more job ready than their 18 year-old counterparts and so may have fewer effective training options open to them. This may explain why the older claimants did not respond to the benefit cut by taking up training places to the same extent as 18 year olds.

In the remainder of this section, we suggest a possible explanation for the stronger effect of the benefit cuts on exits to work for 18 year olds, which focuses on the interaction between the benefit system and the minimum wage. In 2009/10 the Irish national minimum wage was €8.65 per hour.¹⁹ However, the minimum wage system allows for a number of sub-minima, some of which are age related. Workers in their first year of employment after turning 18 are entitled to 80% of the full rate, while those in the second year of employment are entitled to 90%. As noted above, a substantial minority of 18 year-old claimants had no previous employment spell and so would be entitled only to the lower 80% rate. Of those who did have an employment spell, the median duration was just over a year, so this group would be entitled to the 90% rate. In contrast, almost 90% of 21 year olds had a previous employment spell and, conditional on having had a

¹⁸ The percentages given here are weighted averages of the relevant pre- and post-legislation percentages in Table 2.

¹⁹ More details about the Irish national minimum wage can be found at <http://www.lowpaycommission.ie/>

spell, the median duration of this spell was over 2 years. Consequently, many of these older claimants were entitled to the full adult minimum wage. Differences in the effective minimum wage rates facing these claimants therefore leads to truncation of the wage offer distribution at different points. As a result, 18 year-old claimants are much more likely to receive wage offers from the lower part of the wage distribution. A benefit cut that reduces the reservation wage will lead to more of these low wage offers becoming acceptable.²⁰

To get an indication of the potential effect of the minimum wage, we use the available data on annual earnings and weeks worked to construct a weekly wage rate for each claimant in their exit year.²¹ This can be compared to the corresponding weekly minimum wage at the time of the benefit cut. The 2009 adult hourly rate of €8.65 translates into a weekly rate of €337 for a 39 hour work week. This implies a reduced rate of €270 for those in their first year of employment and €303 for those in their second year. The average of the constructed weekly wage for 18 year olds exiting to work is €296, close to the 80% rate of €270.²² The corresponding rate for 21 year olds is €336, substantially higher than the 18 year-old rate and close to the adult weekly minimum wage rate of €337.

Figure 9 plots the density of accepted wages for those entering unemployment during the month before (control) and the month after (treatment) the benefit cut. The results for 18 year olds are provided in the left panel and those for 21 year olds in the right panel. It is clear that, both before and after the benefit cut, the accepted wage distributions for 18 year olds have substantially more mass below the full-time adult minimum wage rate of €337 than the distributions for 21 year olds. Furthermore, there is a notable shift to the left in the mode of the 18 year olds' density after the reduction in benefits. The new mode for the 18 year-old treatment group is €230, substantially below the old mode of €310; when the benefit rate is cut by €100, the mode in the accepted wage distribution falls by €80. No such pattern is evident in the densities for 21 year olds.

²⁰ On the other hand, if the benefit cut reduces the duration of unemployment and so reduces human capital depreciation, this may lead to higher wage offers and therefore higher accepted wages despite the fall in the reservation wage (see Schmieder et al. 2016).

²¹ Our data do not contain information on hours worked and therefore we cannot construct an hourly wage rate. It should also be noted that the weekly rate is calculated using earnings over the entire exit year. We cannot identify the earnings actually received on exiting unemployment.

²² The €270 rate is based on 39 hour week, which likely explains why some observed wages are lower than this.

Since we do not have information on hourly wages and cannot precisely identify earnings on exit from unemployment, this analysis is only suggestive. However, it is consistent with a role for the minimum wage in explaining why 18 year olds subject to the benefit cut moved into jobs more quickly than 21 year olds. Although the benefit cuts may reduce all claimants' reservation wages to the same extent, this reduction may move more 18 year olds into low paying jobs, jobs that their 21 year-old counterparts are not offered due to minimum wage legislation.

8. Long-Run Effects

In this section of the paper, we consider the potential long-run effects of the benefit cuts. Since our earlier analysis reveals a significant initial effect only for 18 year olds, we focus on this group in this final section. Recent evidence in the programme evaluation literature indicates that the effects of active labour market programmes are strongest in the long-run (Card et al. 2015c). The same could be true for the effects of benefit cuts; if the shorter initial durations prevent human capital depreciation, then claimants not only find jobs more quickly but are also more likely to remain employed. On the other hand, the long-run effects of benefit cuts could be weaker than the short-run effects if the cuts force people to end their job search prematurely and move into low paying, low quality, transitory jobs.²³ This could lead to substantial churning between states, which would weaken the long-run impact of the cuts. For those entering training schemes following benefit cuts, the long-run effects will also depend on the effectiveness of these schemes.

To assess the long-run impact of the cuts, we repeat the RD approach used in Section 6, but rather than using the duration of the initial spell as the dependent variable, we use the total duration of recorded unemployment from the start of the initial spell through December 31, 2014. Thus the total duration combines the duration of the initial spell with that of any subsequent spell of unemployment occurring during the five year period after the cut. The result, shown in Figure 10, suggests that those subject to the cut had shorter total durations. The point estimate obtained from the robust RD approach for this long-run effect is approximately -50 weeks and is statistically significant. This is slightly smaller than the short-run effect of -60 weeks reported in Table 3. The similarity of these estimates suggests that 18 year olds subject to the benefit cut did not receive a

²³ Lalive (2007), Card et al. (2007), Van Ours and Vodopivec (2008) and Schmieder et al (2016) all examine extensions to the duration of UI benefits directly and find little evidence that extended UI duration improve subsequent job match quality.

significant additional boost to employment prospects as a result of having left unemployment early. On the other hand, the short-run effect was maintained to a large extent, so that churning does not seem to have been an issue. However, this estimated effect is a combination of both the initial effect and any potential long run effect. In the remainder of this section, we focus on long-run effects.

We repeat our earlier RD analysis for a number of long-run outcomes. We first examine the duration of employment that follows the initial unemployment spell for those exiting to work. A similar measure has been used by Card et al. (2007), Van Ours and Vodopivec (2008) and Schmieder et al. (2016) to examine job-match quality. The results are given in the first column of Table 6 and indicate a positive but statistically insignificant effect of the benefit cut on subsequent employment duration. The remaining analyses concentrate on 2014, the last year for which new unemployment spells are recorded in our data. Focusing on this year allows us to examine the impact of the benefit cut on claimants five years later. We focus on three measures: whether the claimant was unemployed in October 2014; the total time spent unemployed in 2014; and the weekly wages reported in 2014 for those with positive earnings. Again, the results provide weak evidence of a long-run effect. The RD results suggest that the benefit cut in 2009 reduced the likelihood of being unemployed five years later by six percentage points; reduced the total time spent unemployed in 2014 by just over five weeks; and increased average weekly wages in 2014 by €4. However, none of these long-run effects was statistically significant.

9. Conclusion

This paper adds to the existing literature on the labour supply effects of changes in unemployment benefits by evaluating the impact of an unusually large cut in benefits on unemployment duration during the Great Recession. While most existing studies focus on middle-aged workers, our study provides additional information on benefit responsiveness for very young labour market participants, a group that are of particular policy interest.

Despite the very large cuts implemented during this period, we were unable to detect a significant labour supply response for those aged 20-21. In contrast, we find that the benefit cut substantially reduced unemployment duration for 18 year olds. For this age group, who are predominantly low educated and have little previous employment experience, we estimate a

significant duration elasticity of 1.04. This implies a reduction in unemployment durations of over a year. We find a similar, though less precisely estimated, elasticity for those aged 19.

Our results provide clear evidence of a labour supply response to lower unemployment benefits for the youngest claimants, even during a recession. The responses we uncover will clearly have generated immediate savings for the government. However, it is important to recognise that our analysis provides only weak evidence that the cuts had beneficial long-run labour market effects. The potential poverty and health risks associated with the cuts should also be considered when discussing these benefit reform.

To examine the effects of the benefit cuts in more detail, we decompose the overall effect into the components due to different exit states. This analysis shows that no one exit state dominated the large treatment effect estimated for 18 year olds. Furthermore, the smaller effect for older age groups was evident across all exit states. We consider two possible explanations for the differences in behaviour across age groups. To explain differences in exits to training schemes, we focus on differences in the job-readiness of the claimants. To explain differences in exits to work, on the other hand, we consider potential interactions between unemployment benefits, wage offers and minimum wage legislation. Our findings in this regard provide a useful reminder that the impact of changes to unemployment benefits may depend on other labour market policies and that such interactions between policies should be taken into account when considering benefit reforms.

Table 1: Weekly Rate (€) of Jobseeker’s Allowance for New Claimants, 2009-2010

Age	2009		2010
	January	May	January
18	204.30	100.00	100.00
19	204.30	100.00	100.00
20	204.30	204.30	100.00
21	204.30	204.30	100.00
22	204.30	204.30	150.00
23	204.30	204.30	150.00
24	204.30	204.30	150.00
25+	204.30	204.30	196.00

Table 2: Variable Means for New Claimants by Age and Date of Entry to Unemployment

	Age 18 May 1 2009		Age 19 May 1 2009		Age 20 January 1 2010		Age 21 January 1 2010	
	Entry in Six Months <i>Before</i>	Entry in Six Months <i>After</i>	Entry in Six Months <i>Before</i>	Entry in Six Months <i>After</i>	Entry in Six Months <i>Before</i>	Entry in Six Months <i>After</i>	Entry in Six Months <i>Before</i>	Entry in Six Months <i>After</i>
Nationality Irish	0.95	0.94	0.94	0.94	0.93	0.94	0.94	0.94
Low Education [†]	0.49	0.40	0.35	0.30	0.31	0.28	0.30	0.28
No Previous Employment Spell	0.40	0.52	0.19	0.29	0.17	0.21	0.12	0.14
Affected by Benefit Cut	0.02	0.87	0.01	0.56	0.00	0.42	0.00	0.33
Unemployment Spell Duration (Weeks)	115	80	90	74	78	62	75	58
<i>N</i>	3430	2893	2418	2180	1954	1963	1986	2267

[†] Indicates not having completed secondary schooling; education data are available for approximately 90% of claimants

Table 3: Fuzzy Regression Discontinuity Results for Four Benefit Cuts
Standard Errors in Parentheses

	Age 18 2009	Age 19 2009	Age 20 2010	Age 21 2010
First Stage: Effect on Proportion Treated	0.65*** (0.05)	0.37*** (0.07)	0.25** (0.10)	0.18** (0.08)
Effect of Treatment on Unemployment Duration	-60.96** (23.11)	-49.81 (37.92)	-16.73 (82.52)	2.28 (65.66)
<i>N</i>	6323	4598	3917	4253
<i>Elasticity Calculations</i>				
Mean Unemployment Spell Duration Before Treatment (Weeks)	114.77	90.42	78.03	75.05
Estimated Duration Change (%)	-53.11	-55.09	-21.44	+3.03
Benefit Change (%)	-50.9	-50.9	-50.9	-50.9
Estimated Elasticity	1.04	1.08	0.42	-0.06

Notes: *** Denotes significant at the 1% level. ** Denotes significant at the 5% level.

Table 4: Difference-in-Difference Hazard Function Results for Four Benefit Cuts
Standard Errors in Parentheses

	Age 18 2009	Age 19 2009	Age 20 2010	Age 21 2010
Treatment Month	-0.049 (0.081)	-0.057 (0.095)	-0.033 (0.079)	-0.030 (0.088)
Treatment Year	-0.007 (0.080)	-0.035 (0.093)	0.179** (0.087)	0.263*** (0.095)
Treatment Month x Treatment Year	0.349*** (0.109)	0.163 (0.129)	-0.032 (0.119)	0.164 (0.126)
Nationality Irish	-0.115 (0.118)	-0.009 (0.141)	0.190 (0.116)	0.106 (0.124)
Low Education	-0.539*** (0.057)	-0.317*** (0.070)	-0.364*** (0.064)	-0.376*** (0.066)
<i>t</i>	-0.012*** (0.001)	-0.014*** (0.001)	-0.008*** (0.001)	-0.008*** (0.001)
<i>r</i> ² /100	0.004*** (0.0002)	0.004*** (0.0003)	0.002*** (0.0004)	0.002*** (0.0003)
Constant	-3.789*** (0.136)	-3.716*** (0.161)	-4.089*** (0.132)	-4.167*** (0.144)
<i>N</i>	1388	986	1214	1070

Notes: Reference year for Difference-in-Difference estimation is one year earlier in each case. *** Denotes significant at the 1% level. ** Denotes significant at the 5% level.

Table 5: Competing Risks Decompositions for Four Benefit Cuts

	Age 18 2009	Age 19 2009	Age 20 2010	Age 21 2010
Overall Treatment Effect	-38.96**	-15.58*	2.70	-8.63
<i>Decomposition</i>				
Training & Education	-11.75	-6.87	-5.47	-3.17
Work	-14.47	-4.65	0.51	-2.62
Inactivity	-2.33	-0.84	-1.20	-0.98
Other	-10.39	-3.19	8.87	-1.85

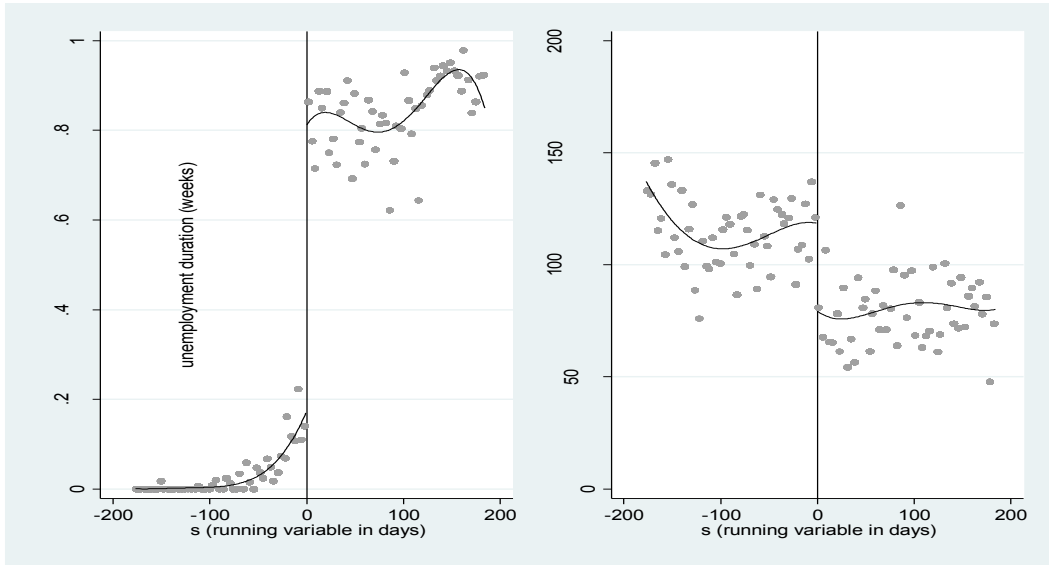
Notes: ** Denotes significant at the 5% level. * Denotes significant at the 10% level.

Table 6: Fuzzy Regression Discontinuity Results on Long-Run Outcomes for 18 Year Olds.
Standard Errors in Parentheses

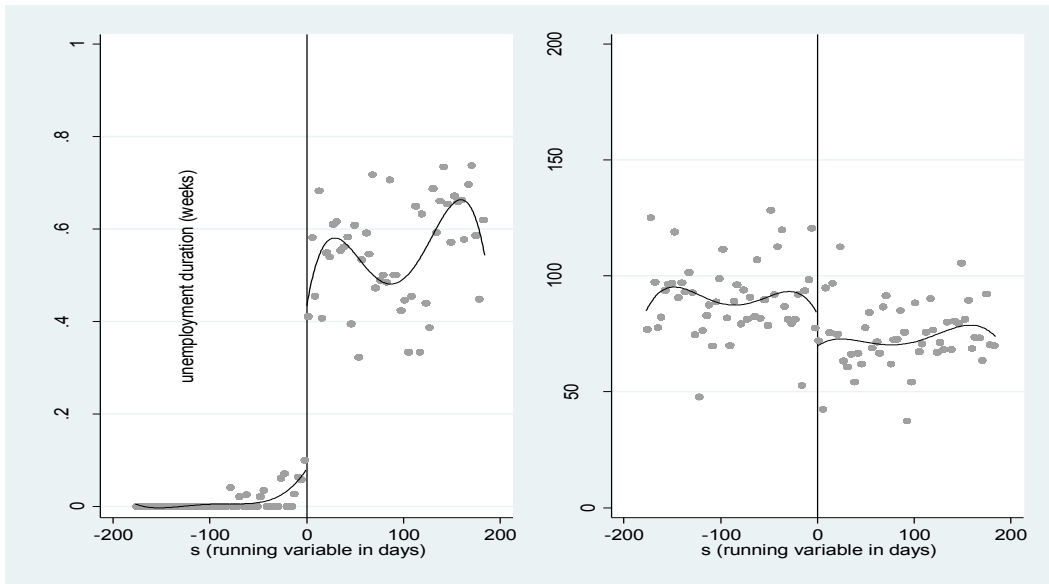
	Duration of Post- Unemployment Job	Incidence of Unemployment, October 2014	Total Unemployment, 2014 (Weeks)	Weekly Wages, 2014
Treatment Effect	8.80 (45.26)	-0.06 (0.086)	-5.13 (3.99)	4.01 (89.60)

Figure 1: Regression Discontinuity Graphs for Various Benefit Cuts
 Proportion Treated (left panel) and Average Unemployment Duration (right panel) for Entrants
 to Unemployment Six Months Before and After Threshold Date

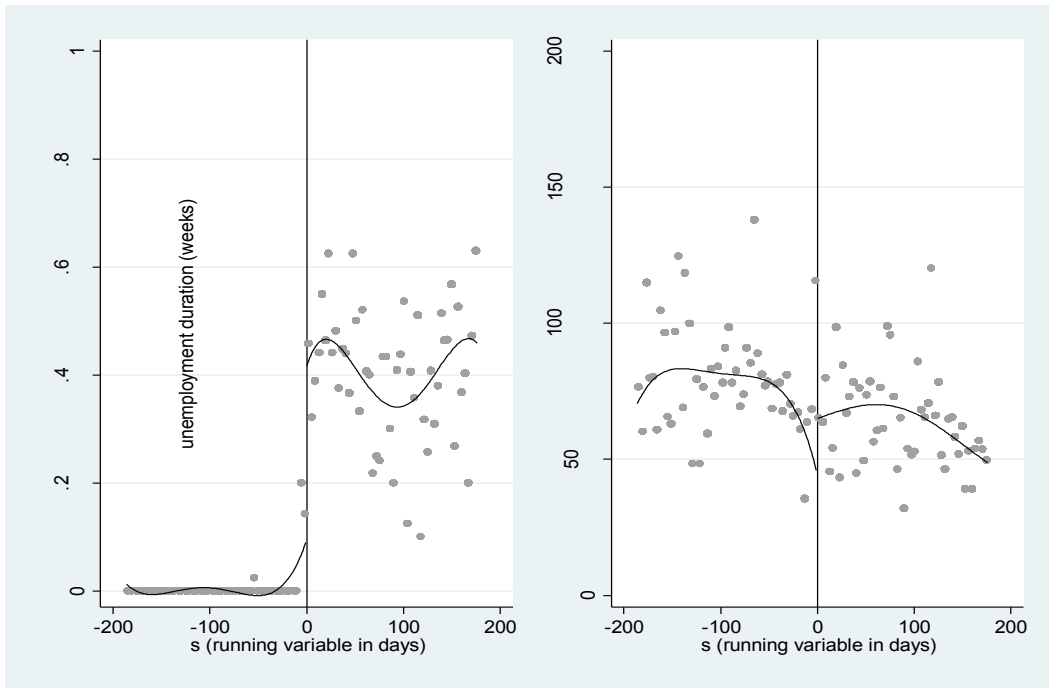
(a) Benefit Cut for 18 Year Olds, May 1 2009



(b): Benefit Cut for 19 Year Olds, May 1 2009



(c) Benefit Cut for 20 Year Olds, January 1 2010



(d) Benefit Cut for 21 Year Olds, January 1 2010

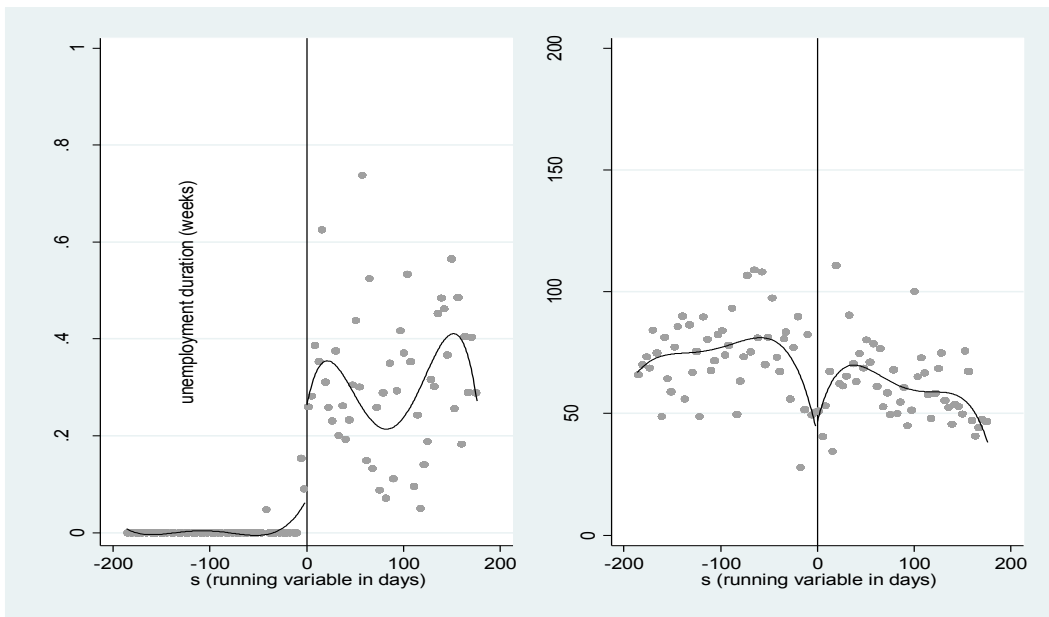


Figure 2: Regression Discontinuity Graphs of Average Unemployment Duration, 18 Year Old Entrants to Unemployment Six Months Before and After May 1 2008 (left panel) and May 1 2010 (right panel)

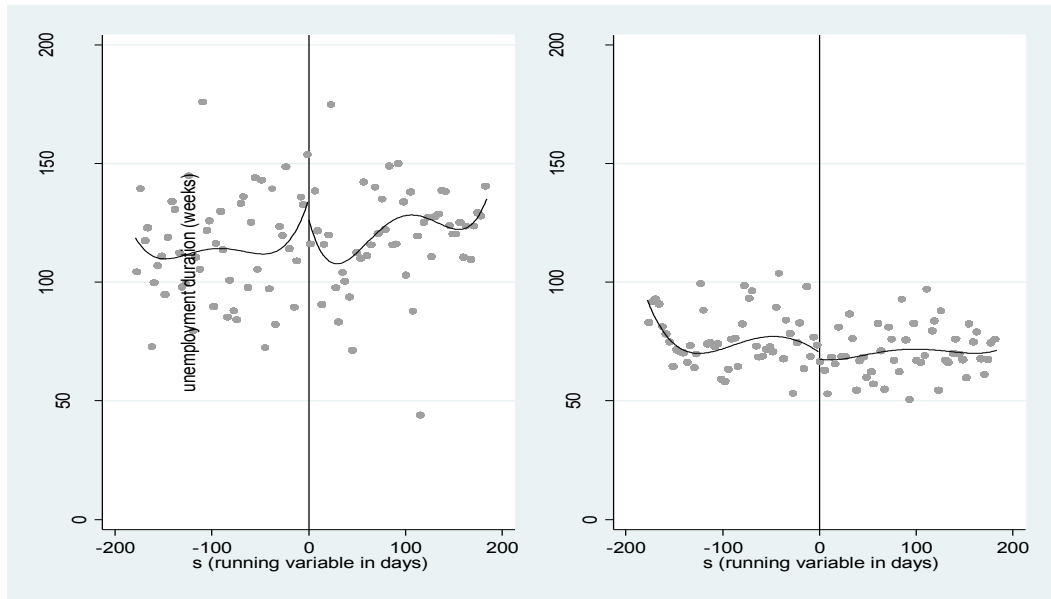


Figure 3: Regression Discontinuity Graphs of Average Unemployment Duration, 18 Year Old Entrants to Unemployment Six Months Before and After April 1 2009 (left panel) and June 1 2009 (right panel)

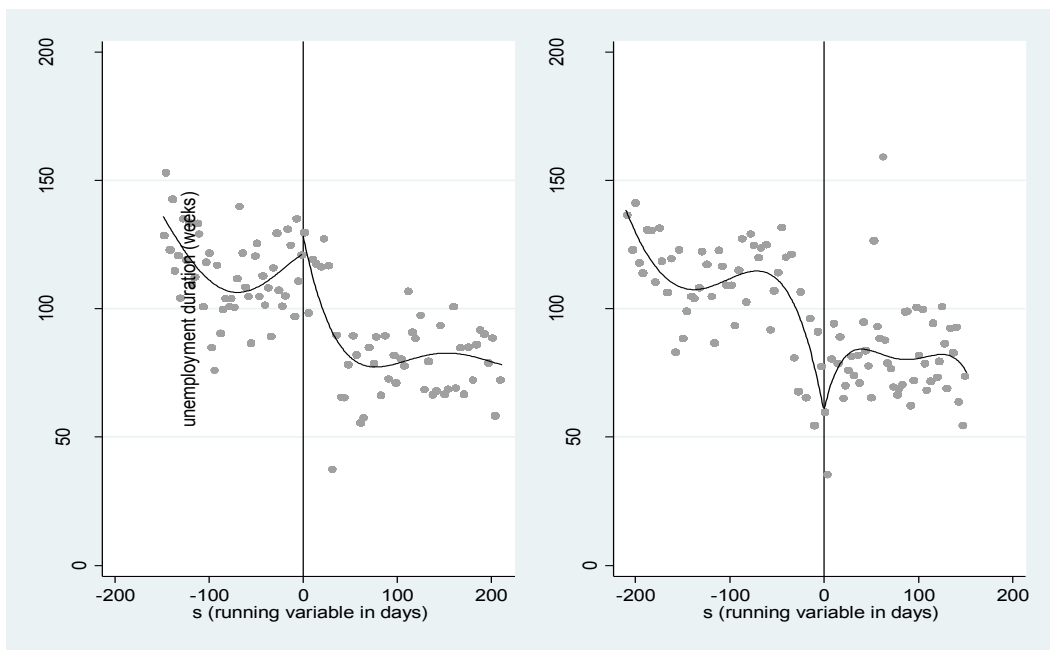


Figure 4: Regression Discontinuity Graph of Average Unemployment Duration, 18 Year Old Entrants to Unemployment Six Months Before and After May 1 2009 for UB Claimants

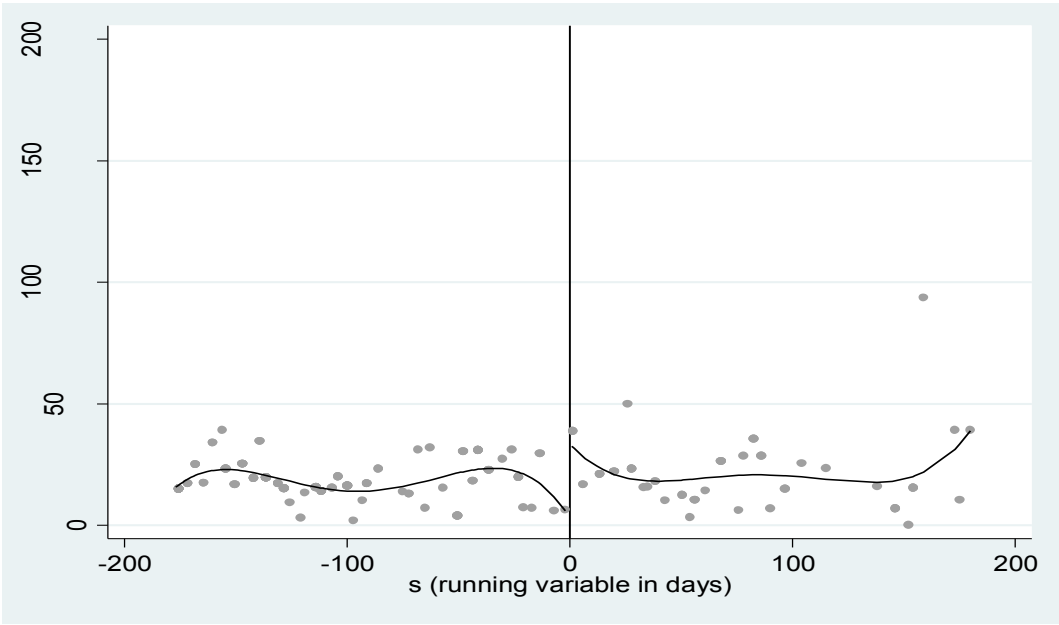


Figure 5: Regression Discontinuity Graph of Density of Entries to Unemployment, 18 Year Old Entrants to Unemployment Six Months Before and After May 1 2009

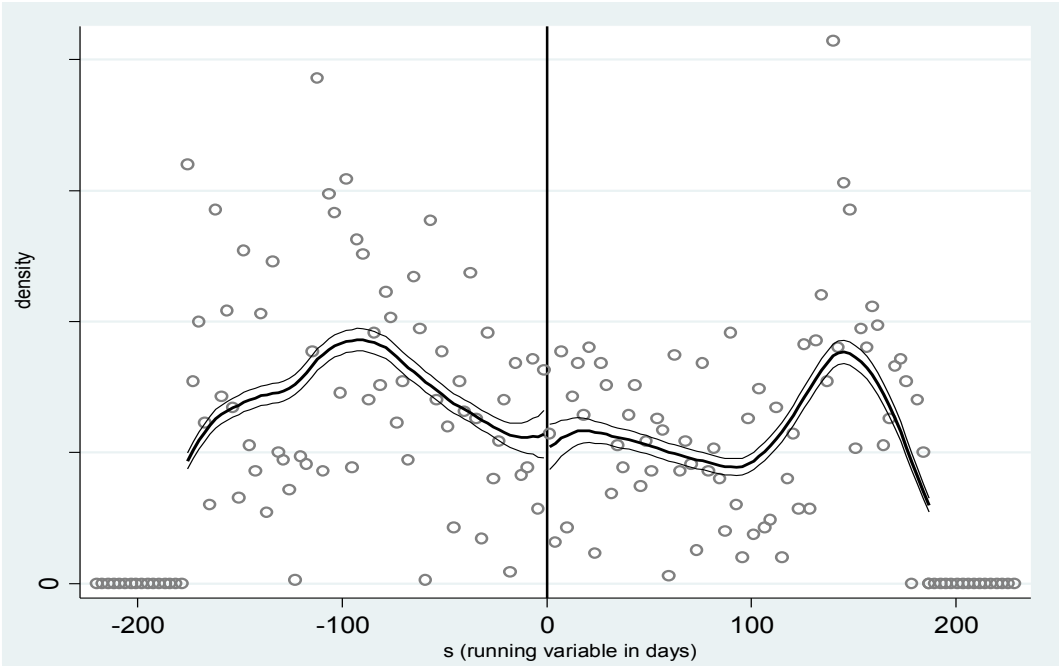


Figure 6: Kaplan-Meier Unemployment Exit Hazard Functions, 18 Year Old Entrants to Unemployment One Month Before and After May 1, 2008 (left panel) and 2009 (right panel)

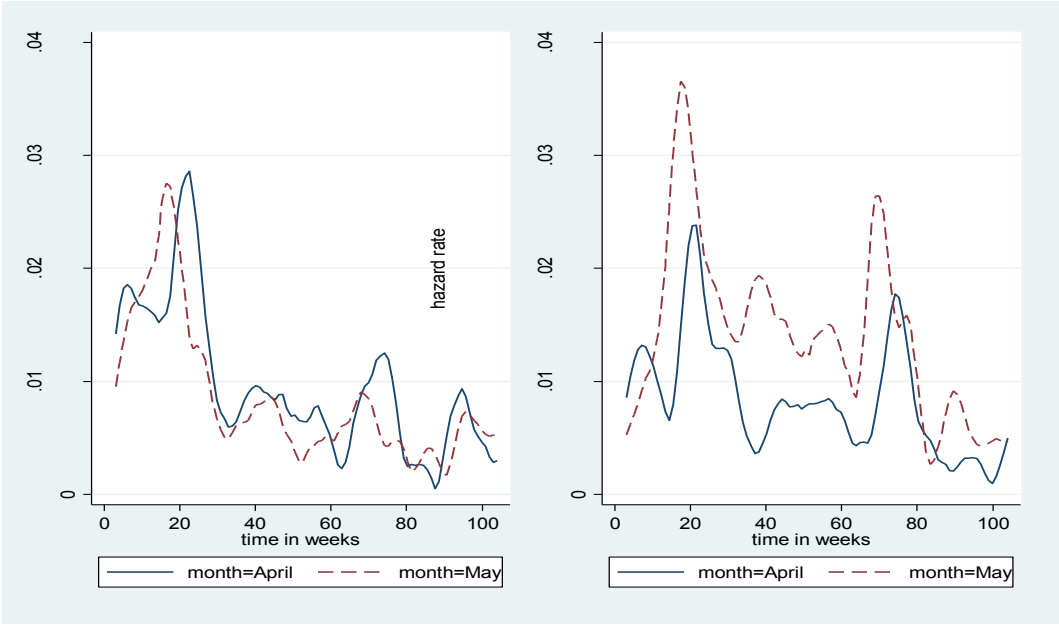


Figure 7: Kaplan-Meier Unemployment Exit Hazard Functions, 20 Year Old Entrants to Unemployment One Month Before and After January 1, 2009 (left panel) and 2010 (right panel)

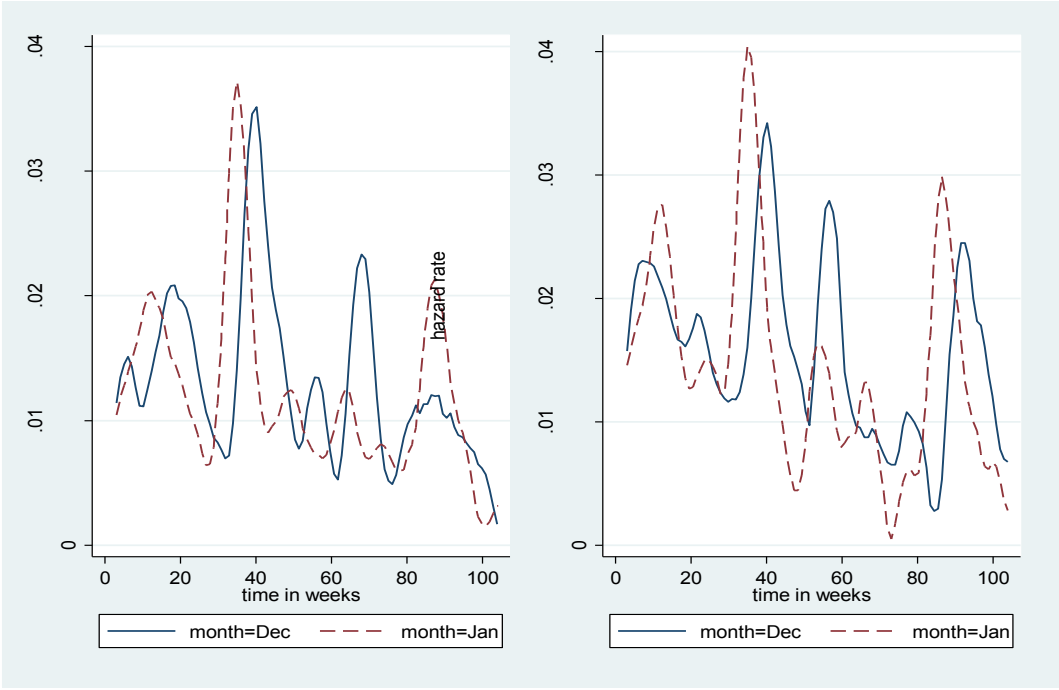


Figure 8: Kernel Densities of Weekly Wages in Year of Exit from Unemployment by Treatment and Control Groups, 18 Year Olds (left panel) and 21 Year Olds (right panel)

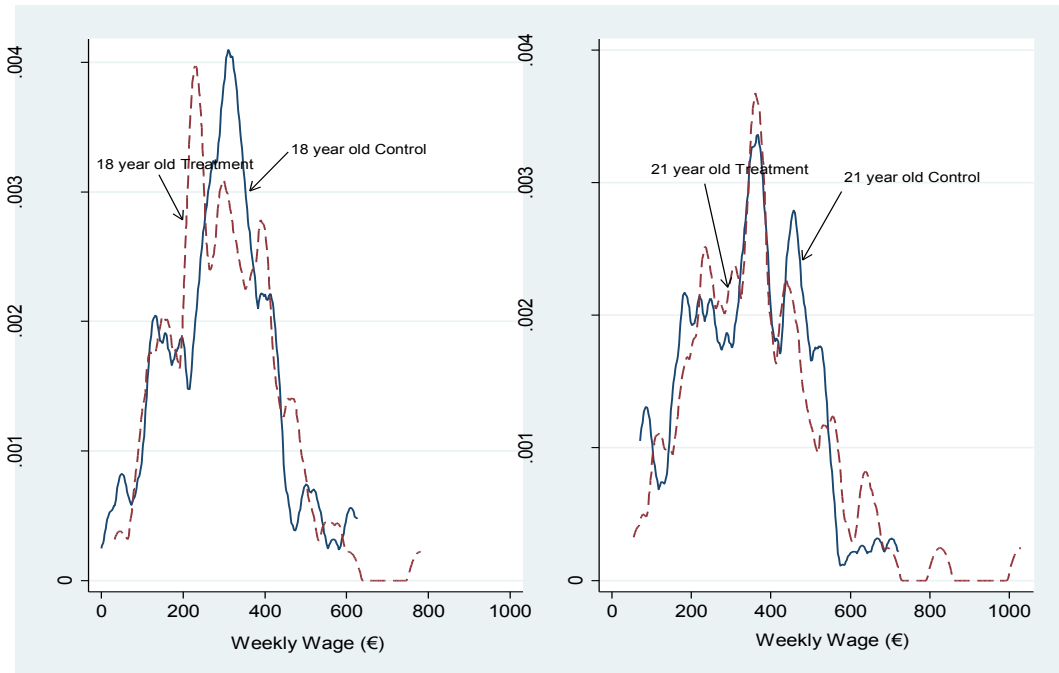


Figure 9: Regression Discontinuity Graph of Total Time Unemployed from 2009-2014, 18 Year Old Entrants to Unemployment Six Months Before and After May 1 2009

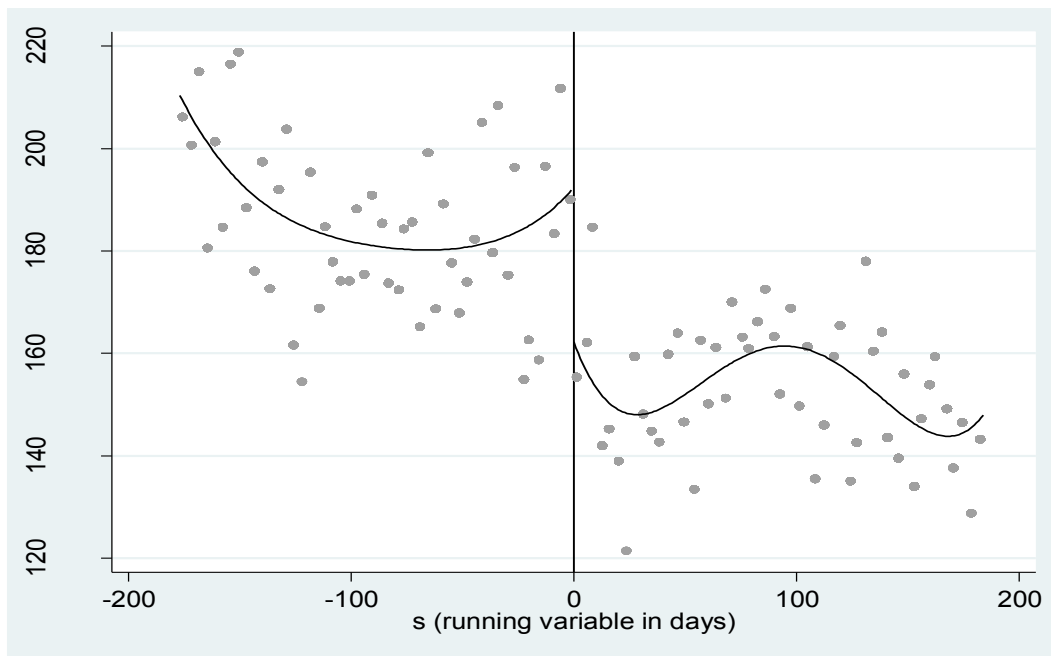


Table A1: Fuzzy Regression Discontinuity Results for Women, Four Benefit Cuts

Standard Errors in Parentheses

	Age 18 2009	Age 19 2009	Age 20 2010	Age 21 2010
First Stage: Effect on Proportion Treated	0.790 0.055	0.440 0.066	0.159 0.089	0.244 0.064
Effect of Treatment on Unemployment Duration	-51.99*** 16.98	-37.89 28.92	75.47 92.95	-46.55 79.89
<i>N</i>	3591	2867	2369	2399
<i>Elasticity Calculations</i>				
Mean Unemployment Spell Duration Before Treatment (Weeks)	103.0	88.54	69.71	57.12
Estimated Duration Change (%)	-50.47	-42.79	108.26	-81.50
Benefit Change (%)	-50.9	-50.9	-50.9	-50.9
Estimated Elasticity	0.99	0.84	-2.12	1.60

References

- ARULAMPALAM, W., and M. B. STEWART (1995): "The Determinants of Individual Unemployment Durations in an Era of High Unemployment," *The Economic Journal*, 105, 321-332.
- ATKINSON, A. B., and J. MICKLEWRIGHT (1991): "Unemployment Compensation and Labor Market Transitions: A Critical Review," *Journal of Economic Literature*, 29, 1679-1727.
- BANERJI, A, S. SAKSONOV, H. LIN, and R. BLAVY (2014): "Youth Unemployment in Advanced Economies in Europe: Searching for Solutions," IMF Staff Discussion Note SDN 14/11
- BARGAIN, O., and K. DOORLEY (2011): "Caught in the Trap? Welfare's Disincentive and the Labor Supply of Single Men," *Journal of Public Economics*, 95, 1096-1110.
- BELL, D. N. F., and D. G. BLANCHFLOWER (2011): "Young People and the Great Recession," *Oxford Review of Economic Policy*, 27, 241-267.
- BENNMARKE, H., K. CARLING, and B. HOLMLUND (2007): "Do Benefit Hikes Damage Job Finding? Evidence from Swedish Unemployment Insurance Reforms," *LABOUR*, 21, 85-120.
- CALONICO, S., M. CATTANEO, M. FARRELL, and R. TITIUNIK (2016): "Regression Discontinuity Designs Using Covariates," *LSE Working paper*.
- CALONICO, S., M. D. CATTANEO, and R. TITIUNIK (2014): "Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs," *Econometrica*, 82, 2295-2326.
- (2015): "Optimal Data-Driven Regression Discontinuity Plots," *Journal of the American Statistical Association*, 110, 1753-1769.
- CARD, D., and P. B. LEVINE (2000): "Extended Benefits and the Duration of UI Spells: Evidence from the New Jersey Extended Benefit Program," *Journal of Public Economics*, 78, 107-138.
- CARD, D., R. CHETTY, and A. WEBER (2007): "Cash-on-Hand and Competing Models of Intertemporal Behavior: New Evidence from the Labor Market," *The Quarterly Journal of Economics*, 122, 1511-1560.
- CARD, D., D. S. LEE, Z. PEI, and A. WEBER (2015a): "Inference on Causal Effects in a Generalized Regression Kink Design," *Econometrica*, 83, 2453-2483.
- CARD, D., A. JOHNSTON, P. LEUNG, A. MAS, and Z. PEI (2015b): "The Effect of Unemployment Benefits on the Duration of Unemployment Insurance Receipt: New Evidence from a Regression Kink Design in Missouri, 2003-2013," *American Economic Review*, 105, 126-30.
- CARD, D., J. KLUVE, and A. WEBER (2015c): "What Works? A Meta Analysis of Recent Active Labor Market Program Evaluations," *National Bureau of Economic Research Working Paper Series*, No. 21431.
- CARLING, K., B. HOLMLUND, and A. VEJSIU (2001): "Do Benefit Cuts Boost Job Finding? Swedish Evidence from the 1990s," *The Economic Journal*, 111, 766-790.
- CARLING, K., and L. LARSSON (2005): "Does Early Intervention Help the Unemployed Youth?," *Labour Economics*, 12, 301-319.
- CHEMIN, M and E. WASMER (2012): "Ex-Ante and Ex-Post Evaluation of the 1989 French Welfare Reform Using a Natural Experiment: The 1908 Social Laws in Alsace-Moselle," *LIEPP Working paper no. 3*.
- COVIELLO, V., and M. BOGGESS (2004): "Cumulative Incidence Estimation in the Presence of Competing Risks," *Stata Journal*, 4, 103-112.
- FORTIN, B., G. LACROIX, and S. DROLET (2004): "Welfare Benefits and the Duration of Welfare Spells: Evidence from a Natural Experiment in Canada," *Journal of Public Economics*, 88, 1495-1520.
- FOUGÈRE, D., F. KRAMARZ, and T. MAGNAC (2000): "Youth Employment Policies in France,"

- European Economic Review*, 44, 928-942.
- GELMAN, A., and G. IMBENS (2014): "Why High-Order Polynomials Should Not Be Used in Regression Discontinuity Designs," *National Bureau of Economic Research Working Paper Series*, No. 20405.
- HAHN, J., P. TODD, and W. VAN DER KLAUW (2001): "Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design," *Econometrica*, 69, 201-209.
- HAN, A., and J. A. HAUSMAN (1990): "Flexible Parametric Estimation of Duration and Competing Risk Models," *Journal of Applied Econometrics*, 5, 1-28.
- HOWELL, D. R., and B. M. AZIZOGLU (2011): "Unemployment Benefits and Work Incentives: The U.S. Labour Market in the Great Recession," *Oxford Review of Economic Policy*, 27, 221-240.
- HUNT, J. (1995): "The Effect of Unemployment Compensation on Unemployment Duration in Germany," *Journal of Labor Economics*, 13, 88-120.
- IMBENS, G., and K. KALYANARAMAN (2011): "Optimal Bandwidth Choice for the Regression Discontinuity Estimator," *The Review of Economic Studies*.
- IMBENS, G. W., and T. LEMIEUX (2008): "Regression Discontinuity Designs: A Guide to Practice," *Journal of Econometrics*, 142, 615-635.
- JENKINS, S. P. (1995): "Easy Estimation Methods for Discrete-Time Duration Models," *Oxford Bulletin of Economics and Statistics*, 57, 129-136.
- JENSEN, P., M. ROSHOLM, and M. SVARER (2003): "The Response of Youth Unemployment to Benefits, Incentives, and Sanctions," *European Journal of Political Economy*, 19, 301-316.
- JONASSEN, A. (2013): "Disincentive Effects of a Generous Social Assistance Scheme," SFI – Danish National Centre for Social Research, School of Business and Social Sciences, Aarhus University, WP-01:2013.
- KALBFLEISCH, D. and R. PRENTICE (2002): *The Statistical Analysis of Failure Time Data*. John Wiley & Sons, Inc., New York.
- KROFT, K., and M. NOTOWIDIGDO (2016): "Should Unemployment Insurance Vary with the Unemployment Rate? Theory and Evidence," *Review of Economic Studies*, 83, 1092-1124.
- KRUEGER, A. B., and B. D. MEYER (2002): "Labor Supply Effects of Social Insurance," Chapter 33 in *Handbook of Public Economics*: Elsevier, 2327-2392.
- KYYRÄ, T. and H. PESOLA (2017): "The Effect of UI Benefits on Unemployment and Subsequent Outcomes: Evidence from a Kinked Benefit Rule," *IZA Discussion Paper* no.10484.
- LALIVE, R. (2007): "Unemployment Benefits, Unemployment Duration, and Post-Unemployment Jobs: A Regression Discontinuity Approach," *The American Economic Review*, 97, 108-112.
- LAYARD, R., S. NICKELL, and R. JACKMAN (1991): *Unemployment: Macroeconomic Performance and the Labour Market*. Oxford University Press.
- LEE, D. S., and D. CARD (2008): "Regression Discontinuity Inference with Specification Error," *Journal of Econometrics*, 142, 655-674.
- LEMIEUX, T., and K. MILLIGAN (2008): "Incentive Effects of Social Assistance: A Regression Discontinuity Approach," *Journal of Econometrics*, 142, 807-828.
- LEVINE, P. B. (1993): "Spillover Effects between the Insured and Uninsured Unemployed," *Industrial and Labor Relations Review*, 47, 73-86.
- MCCRARY, J. (2008): "Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test" *Journal of Econometrics*, 142, 698-714.

- MEYER, B. D. (1990): "Unemployment Insurance and Unemployment Spells," *Econometrica*, 58, 757-782.
- MORTENSEN, D (1977) "Unemployment Insurance and Job Search Durations," *Industrial Labor Relations Review*, Vol. 30, no. 4, pp. 505-517.
- NARENDRANATHAN, W., S. NICKELL, and J. STERN (1985): "Unemployment Benefits Revisited," *The Economic Journal*, 95, 307-329.
- OECD (1994): *The OECD Jobs Study: Facts, Analysis, Strategies*: Paris OECD.
- OECD (2010): *Off to a Good Start: Jobs for Youth*: Paris OECD.
- OECD (2013): *OECD Employment Outlook 2013*: Paris OECD.
- O'NEILL, D. (2017): "A Competing Risk Decomposition of Average Treatment Effects," *mimeo*
- PUBLIC EXPENDITURE AND REFORM (2013) Expenditure Allocation 2014-2016.
- REBELLO-SANZ, Y., and N. RODRIGUEZ-PALANAS (2016): "When the Going Gets Tough..Financial Incentives, Duration of Unemployment and Job Match Quality," *IZA DP 10044*
- SCARPETTA, S. A. A. SONNET and T. MANFREDI. (2010) "Rising Youth Unemployment During the Crisis: How to Prevent Negative Long-term Consequences on a Generation?" OECD Social, Employment and Migration Working Papers No. 106
- SCHMIEDER, J. F., T. VON WACHTER, and S. BENDER (2016): "The Effect of Unemployment Benefits and Non-employment Durations on Wages," *American Economic Review*, 106, 739-777.
- TATSIRAMOS, K., and J. C. VAN OURS (2014): "Labor Market Effects of Unemployment Insurance Design," *Journal of Economic Surveys*, 28, 284-311.
- VAN OURS, J. C., and M. VODOPIVEC (2008): "Does Reducing Unemployment Insurance Generosity Reduce Job Match Quality?," *Journal of Public Economics*, 92, 684-695.
- VAN OURS, J. C. (2015): "The Great Recession Was Not So Great," *Labour Economics*, 34, 1-12.
- WALSH, F.,(2015): "Labour Market Measures in Ireland 2008–13:The Crisis and Beyond," International Labor Office Report