Does Reducing Unemployment Benefits During a Recession Reduce Youth Unemployment? Evidence from a 50% Cut in Unemployment Assistance[1]

Aedín Doris[*], Donal O'Neill[**] & Olive Sweetman[*]

June 2017

Abstract

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 spells, claimants whose unemployment start dates differed by one day received very different benefits; we exploit this feature in our Regression Discontinuity and Difference-in-Difference analyses. We find that the benefit cut significantly reduced unemployment duration for young claimants, with an elasticity close to one. Exits to training and work account for the majority of this effect. Our analysis provides only weak evidence that the cuts had beneficial long-run effects.

[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, Chris Jepsen , John Kennan, Paul Redmond and participants at the Labour Market Council conference on program evaluation (Dublin), the 31st Irish Economic Association Conference (Dublin), the 22nd Society of Labor Economists Conference (North Carolina), the 5th NERI Annual Labour Market Conference (Maynooth), the 31st European Society for Population Economics Annual conference (Glasgow) and seminar participants at Maynooth University for helpful comments on an earlier version of this paper.

[**]Corresponding author; Maynooth University and IZA, Bonn. E-mail: donal.oneill@nuim.ie; tel.: 353-1-7083555; fax: 353-1-7083934; address: Rhetoric House, Maynooth University, Maynooth, Co. Kildare, Ireland.
[*]Maynooth University.

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 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 job search and training requirements, enforced by the threat of benefit sanctions (OECD 2013)[2].

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. Moreover, little is known about the impact of unemployment benefit reform during severe economic downturns. In this paper, we address both these issues by examining the labour market responses of 18 and 19 year olds to a cut in unemployment benefit introduced in Ireland during the Great Recession. For those affected, weekly benefits fell from €204.30 to €100. To evaluate this benefit cut, we use a quasi-experimental approach that exploits the fact that only new claimants were subject to the cut. 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 a substantial intervention with rich administrative data on the entire population of claimants

[2] For a recent evaluation of the latter policy see van den Berg et al. 2017.

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

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. When considering destination states, we find evidence that the duration effect is largely driven by exits to training and work. Examining post-unemployment wages, we find that claimants are typically moving to low-paid jobs but that the benefit cut had no significant effect on the wage rates accepted. In addition, our analysis provides only weak evidence that the cuts had beneficial long-run effects.

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 the relevant features of the Irish welfare system and describes the changes made by the government 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. We also consider the wages of those exiting to work. 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 (e.g. Jensen et al. 2003, Carling and Larson 2005, Bell and Blanchflower 2011, Banerji et al. 2014), our focus in this paper 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 their reservation wage, 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. In this model, a cut in benefits will reduce unemployment duration both by reducing the reservation wage and by increasing job search intensity. Since both the reservation wage and search intensity also depend on the state of the economy, identifying the effect of a benefit cut may be difficult if macroeconomic conditions are changing at the same time.

There is a large body of empirical work that aims to identify the effect of unemployment benefits 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 ranges 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 tend 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, both Landais (2015) and Kroft and Notowidigdo (2016) find that increases in the state unemployment rate in the U.S. are associated with decreases in benefit elasticities. Kroft and Notowidigdo (2016) estimate a benefit elasticity of 0.99 when the state unemployment rate is below the U.S. national average, but 0.28 when it is above the average, while Landais (2015) finds that the estimated duration elasticity is weakly pro-cyclical, varying between 0.25 and 0.38.

To date, only a small number of papers have considered the impact of benefit cuts during the Great Recession. 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) examine the responsiveness of unemployment duration to benefit changes in the state of Missouri over the period 2003-2013. In contrast to the work of Landais (2015) and Kroft and Notowidigdo (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 differ in their time profile; UA payments 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 and 19 year olds. A 19 year old would have expected his payment to increase once he reached his 20th birthday. This may weaken the effect of the benefit cut.

Secondly, 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. On the one hand, low skilled workers have little scope to add lower segments of the labour market to their job search possibilities as the spell progresses, so 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, which can lead to larger benefit responses (Narendranathan et al. 1985); any fall in the reservation wage will tend to have a larger if situated in a dense part of a wage offer distribution because it brings in many more potential job offers.

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.[3] 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 points[4]. 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 latter group is complicated by other reforms implemented at the same time.

[3] For a related analysis of this programme see Chemin and Wasmer (2012).
[4] A related study on Denmark by Jonassen (2013) reports findings in line with Lemieux and Milligan (2007) and Bargain and Doorley (2011).

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 cannot directly identify those claimants who were eligible for benefits, nor can he identify which claimants were subject to the cuts, making it difficult to identify the treatment effect of interest. Although Walsh finds no evidence of a higher rate of transition from unemployment to employment for those affected by the benefit cuts, he is careful to say that the results are suggestive rather than definitive, given the limitations of the data available at that time.

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. Furthermore, 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 the unemployment rate 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 led to the introduction of draconian austerity measures.[5]

Of particular focus in this paper is the substantial reduction 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."[6] 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.

Prior to the benefit cuts examined in this paper, all JA claimants were paid a basic rate of €204.30 a week. On April 29 2009 claimants aged 18 and 19 had their weekly rate cut to €100.[7] 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. We exploit this feature of the legislation in order to identify the impact of the benefit cuts.

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

[5] 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.
[6] http://www.welfare.ie/en/pressoffice/Pages/pr231013.aspx
[7] Between 2010 and 2013, there was also a series of other cuts for those aged 20-25. However, many of the eligible pool in these age groups were exempt from the benefit cut, making it difficult to identify an effect. In addition, all the later cuts came into effect at the beginning of the year so seasonal effects specific to the Christmas period further complicate identification. Therefore, we do not analyse these cuts here.

conditions, the proportion of eligible claimants exempted from the benefit cuts differed between the 18 and 19 year olds. We account for this in our econometric analysis.

The changes to benefits for 18-19 year olds were first announced as part of an emergency budget introduced on April 7 2009 and the legislation putting them into effect was passed three weeks later on April 29. Because the legislation was enacted so soon after its announcement, there was little opportunity for strategic behaviour. Nevertheless, we examine this formally in Section 6c.

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. The restriction of the cuts to new entrants provides the 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.[8] The idea behind RD is that assignment to the treatment is determined

[8] For a review of the RD approach, see Imbens and Lemieux (2008).

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 unemployment durations associated with and without treatment respectively. 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:

$$\alpha = E[Y_i(1) - Y_i(0) | S_i = s_o] \qquad (1)$$

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

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

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

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. Firstly, we have access to the population of claimants, resulting in a large number of observations. Secondly, we know the exact start date of every claim, allowing specification of the running variable in days rather than in weeks or months. We discuss these features 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 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 fact that the running variable is measured in days means that RD approach can identify 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 appropriate bandwidth and in this way contribute to the final estimates. To estimate the RD parameter we use data six months before and six months after the reform.[9]

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

[9] Our main results are robust to using different timeframes.

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:

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

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

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

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

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 April 29, 2009. We regard 18 year olds who commenced a spell in the month following April 29 as the treatment group and 18 year olds who commenced a spell in the month prior to April 29 as the control group. To account for any seasonal effects that may cause durations for those entering prior to April 29 to differ from those entering after this date, we adopt a DiD specification, which also includes spells from the same months in 2008. Specifically, we estimate:

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

$\boldsymbol{Z}_i(t)$ is a vector of covariates including nationality, education and previous employment; T_i is a dummy variable indicating entry after April 29; and $D_{2009,i}$ is a dummy variable indicating entry into unemployment in 2009. The parameter of interest is ϕ, which measures the change in the hazard resulting from the cut in benefit payments.

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.

Several advantages of these data are worth noting. 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 individuals 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, nationality and work history. 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, we also have administrative data on annual earnings and weeks worked for every year in which the individual worked. These data are taken from tax returns submitted by employers. Since it is an offence to misreport earnings to the tax authorities, these data are likely to be free of measurement error. We use these earnings data to examine changes in earnings in response to the benefit cuts.

In total, between 2007 and 2014, there were 136,183 claims by individuals aged 18 or 19, of which 94% 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 parameters. 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 1. For both age groups, about 95% of claimants have Irish nationality. The variable labelled "Low Education" denotes that the claimant did not complete 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 1 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. As anticipated, the proportion with no prior work experience is higher for 18 year olds than for 19 year olds. The fact that those who enter unemployment later in 2009 are somewhat better educated and are less likely to have had a prior job reflects seasonal effects and the deterioration in the labour market during that year. However, as we will see later, these differences are not evident at the RD threshold, and therefore do not affect the validity of the RD design.

The treatment status variable indicates whether the claimant was subject to the legislated benefit cut. For 18 year-old claimants, a substantial majority was subject to the cut. The proportion affected is lower for 19 year olds because, as discussed earlier, older claimants are more likely to 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 both groups. It is worth noting the length of unemployment spells during this period, with average unemployment duration of over a year. This reflects the depressed nature of the Irish labour market at this time. It also noteworthy that for both age groups, 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

Initial results for the RD analysis are shown in Figures 1a and 1b. 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 both age groups, where the running variable is days before or after the introduction of the benefit cut. Each figure consists of two graphs. The graphs on the left 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 legislation, April 29 2009. Examination of these graphs allows us to explore the bite of the legislation, which provides the basis for the denominator of our fuzzy RD estimator. The graphs on the right 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 both cases. As discussed above, the legislation had less bite for 19 year olds. The graphs suggest that the likelihood of treatment increased by between 60 and 70 percentage points for 18 year olds and by about 40 percentage points for 19 year olds. The differential bite of the treatments is taken into account in the fuzzy RD analysis so that the estimates of all treatment effects are consistent. However, it should be noted that the smaller bite of cuts for 19 year olds makes precise estimation of the effects for this groups more difficult.

Turning to the RD plots for unemployment duration, we see that for both 18 and 19 year olds, unemployment durations fell substantially when benefits were cut, with the effect particularly pronounced for 18 year olds. For this group, the graph suggests 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 2.[10] 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).[11] The first row provides the estimated effect of the legislation on the likelihood of receiving a benefit. The second row provides the fuzzy RD estimates of the causal effects of the benefit cuts on unemployment duration and confirm the results of the RD graphs. There is evidence of a strong negative effect for both 18 and 19 year olds, with unemployment durations falling by 61 weeks and 50 weeks respectively. However, only the 18 year-old effect is statistically significant. The lack of significance for the 19 year olds may reflect the weaker bite of the treatment for this group.

We can use the RD results to estimate a benefit duration elasticity for both groups and these are reported in the last row of Table 2. For both age groups, the fall in duration combined with the reduction in benefits imply an elasticity of just over one. This estimate is consistent with the range of estimates reported in the previous literature. Despite the depressed nature of the labour market young youngers reacted substantially to the imposed benefit cuts.

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

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

6b: Robustness Checks

We carry out a number of robustness checks 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. A third check uses the fact that JB claimants were not subject to the benefit cuts, and therefore we should see no significant effect at the threshold for this group. Finally, we examine the impact of covariates for our findings, both by repeating the analysis conditioning on covariates and also by using covariates themselves as pseudo-outcomes.

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 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 (March 29) and a month later (May 29). 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 0.13 and -4.87 respectively, compared to the point estimate of -60.96 obtained with the correct threshold. The RD plot for the population of 18 year old JB claimants in 2009 is given in Figure 4. We see no evidence of a discontinuity in JB durations at the threshold.[12] These robustness checks all support the identifying assumptions underlying our RD estimation.[13]

Covariates can also play a useful role in assessing the plausibility of any RD design (Athey and Imbens (2017)). For the RD identification strategy to be valid, the covariates should be

[12] This also suggests that the benefit cut for JA claimants did not have a spillover effect on JB claimants, who were not subject to the legislation 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.

[13] 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.

uncorrelated with the treatment when the running variable is near the threshold. To check this we follow previous work and repeat the RD analysis using covariates as pseudo-outcomes. A discontinuity in a covariate at the threshold would cast doubt on the validity of the RD approach. The results from this analysis are given in Table 3. We see no evidence of a discontinuity in education, nationality or previous employment for either 18 or 19 year olds. The point estimates are small and statistically insignificant.

Finally, following Calonico et al. (2016), we have also considered the impact of including these covariates in our RD analysis to account for any compositional changes around the threshold. Their inclusion 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 cut had an effect at the extensive margin, reducing the numbers entering unemployment by encouraging young people to stay in school. Such effects would not be picked up in the earlier duration analysis. An additional concern affecting the extensive margin is the possibility of anticipation effects; these occur when individuals initiate claims earlier than otherwise to avoid announced benefit cuts that have not yet taken effect. Given the short time period between announcement and enactment of the legislation discussed in this paper, we think the scope for anticipation effects is limited, but nevertheless requires examination.

To consider effects at the extensive margin, we adjust the RD design used above and check for discontinuities in the density of the running variable itself.[14] If 18 year olds remained in education longer following the reduction in benefits or changed behaviour in anticipation of the benefit cut, we would expect to see a discontinuous fall in the density of entries to unemployment at the threshold. The estimated density is given in Figure 5. The points represent the proportion of all claimants entering unemployment in each bin. The estimated density is continuous at the

[14] McCrary (2008) 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.

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. By ruling out anticipation effects[15], this result also supports the validity of the RD analysis in identifying causal effects.

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 the month after April 29. The control groups are those entering one month earlier. The intervention year is 2009. The hazards for 18 year olds are given in Figure 6 and those for 19 year olds are given in 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 after April 29, the date of the legislated benefit cut, is consistently higher. 18 year olds subject to the benefit cut were more likely to leave unemployment in almost every week following the commencement of their spell. The graph indicates a similar pattern for 19 year olds but the evidence in this case appears to be weaker.

To examine these changes more formally, we estimate the hazard DiD model given by equations (2) and (3) and present the results 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 and those with no previous job were less likely to exit. The key parameter is the coefficient on the interaction term between year and month of entry. We see a significant effect of the legislation for 18 year olds, while the effect is positive but not significant for 19 year olds. The results from the estimated hazard imply that 18 years olds

[15] As an additional check for announcement effects, we conduct an RD analysis using April 7 as the threshold and found no effect.

entering after the legislation were 37%[16] more likely to exit their JA spell than those in receipt of the higher benefit. These hazard results, when combined with the earlier RD results, provide convincing evidence that the benefit cut substantially reduced unemployment duration for 18 year olds.

7. Analysis of Exit States

7a. Competing Risk Decomposition

In Section 6, we reported robust evidence of a substantial and significant effect on unemployment durations for 18 year old claimants, with weaker evidence for 19 year olds. Given the depressed nature of the labour market in 2009, it may have been easier for claimants to exit to training or inactivity than to find employment. Since alternative exit states will have different policy implications, we now extend the previous analysis of unemployment duration by considering the state to which claimants exited.

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

$$\Delta Y = \overline{Y_T} - \overline{Y_C}$$

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 fi_1, fi_2 and fi_3, we can write the overall difference as

$$\Delta Y = (f_{T1}\overline{Y_{T1}} + f_{T2}\overline{Y_{T2}} + f_{T3}\overline{Y_{T3}}) - (f_{C1}\overline{Y_{C1}} + f_{C2}\overline{Y_{C2}} + f_{C3}\overline{Y_{C3}})$$
$$\equiv \left(\frac{N_{T1}}{N_T}\overline{Y_{T1}} + \frac{N_{T2}}{N_T}\overline{Y_{T2}} + \frac{N_{T3}}{N_T}\overline{Y_{T3}}\right) - \left(\frac{N_{C1}}{N_C}\overline{Y_{C1}} + \frac{N_{C2}}{N_C}\overline{Y_{C2}} + \frac{N_{C3}}{N_C}\overline{Y_{C3}}\right)$$

[16] This is calculated as (exp(0.317)-1)*100.

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 $\overline{Y_{ik}}$ is the average duration for those in group *i* who exit to state *k*.

Suppose we observe spells over a period of *D* weeks. Then we can write $\overline{Y_{Tk}}$ as $\sum_{d=1}^{D} \frac{N_{Tk}^d}{N_{Tk}}.d$, 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

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

The terms inside the curly brackets represent the contributions of each of the exit states to the overall difference in duration.[17] In this way, the decomposition allows us to 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".[18] The results of the decompositions for both 18 and 19 year olds 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, rescaled using the first stage treatment effects reported in Table 2. The sizes of the effects are similar to those reported in Section 7 using the RD approach. 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, are not the dominant

---

[17] 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).

[18] 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.

determinant of the overall effect. The same three states are important when considering 19 year olds, although all of the effects are somewhat smaller than the corresponding effects for 18 year olds. As discussed earlier, the government's stated motivation for the benefit cuts was to ensure that education, employment or training were preferable to unemployment. Our results confirm that education, training and work all contributed to the overall reduction in unemployment durations for 18 year olds, and to a lesser extent for 19 year old claimants.

7b: Wage Analysis

Given the importance of exits to work, we examine these exits in more detail in the remainder of this section. As noted earlier, in a simple job search model, faster exits to work following a benefit cut arise as a result of increased search intensity and/or lower reservation wages. While we have no data on search intensity, we do have information on annual earnings and weeks worked for every year in which the individual worked. This allows us to calculate weekly earnings in the year the claimant exited unemployment. These data have the advantage that they are taken from tax records and are therefore free of measurement error. However, because the earnings data refer to the entire calendar year, we cannot identify the earnings actually received on exiting unemployment for a minority of our sample who have multiple employment spells in a given year.[19] Provided such observations appear as multiple records in our sample we delete these multiple entries from the wage analysis.

Figure 8 plots the density of accepted wages for the 18 and 19 year olds entering unemployment during the six months before (control) and the six months after (treatment) the benefit cut.[20] For context, we also include lines at €270 and €304, which correspond to youth

[19] For example, consider a person who became unemployed in June 2009, exited this spell in January 2011, worked until March, was unemployed from April to July and then worked for the remainder of 2011. For this individual, annual earnings in the exit year will be a combination of the earnings from their two jobs, rather than earnings received on exiting their 2009 unemployment spell.

[20] Given that we restrict our attention to those who exit to work, the sample sizes are relatively small when analysing wages. It is for this reason that we use the six months before and after the legislation when considering the accepted wage distributions.

subminimum wage rates for 18 and 19 year olds respectively, based on a 39 hour working week.[21] We see that the average wages accepted by these workers typically correspond to low paid minimum wage level jobs, as might be expected given their characteristics. There is some evidence of a shift to the left in the distributions of those in the treatment groups. However, for both ages, the densities for the treatment and control groups are quite similar and suggest only a limited role for lower reservation wages in explaining the faster exits to work in response to the benefit cut.

While the wage densities provide a useful summary of accepted wages, one must be careful in using them to infer the impact of the benefit cut on wages. As noted by Schmieder et al. (2016), changes to the benefit system change post-unemployment wages through two channels. Firstly, a benefit cut may shift the post-unemployment wage path down; the accepted wage at a given duration falls. Secondly, the benefit cut may change the distribution of claimants along the post-unemployment wage path; those subject to the cut may have shorter durations. The densities given in Figure 8 combine both effects, which may offset each other in aggregate. To identify the shift in the path of post-unemployment wages, we therefore follow Schmieder et al. (2016) and estimate expected post-unemployment wages conditional on duration. To allow for possible seasonal effects we estimate a DiD model using the corresponding period from 2008 as the control year. To identify shifts in the post-unemployment wage path, we condition on the duration of the unemployment spell prior to the job. Formally, we estimate

$$W_i = \boldsymbol{Z}_i'\boldsymbol{\theta} + \alpha_1 T_i + \delta D_{2009,i} + \phi T_i D_{2009,i} + \beta_1 Dur_i + \beta_2 {Dur_i}^2 + \varepsilon_i \tag{4}$$

W_i is the weekly post-unemployment wage. As was the case in the DiD hazard model, $\boldsymbol{Z}_i$ is a vector of covariates including dummies for nationality, education and previous employment; T_i is a dummy variable indicating entry to unemployment in the six months following the introduction of the treatment; $D_{2009,i}$ is a dummy variable indicating entry into unemployment in the treatment year. Dur_i measures the duration (in months) of the relevant unemployment spell and we include a quadratic in duration to allow for a nonlinear post-unemployment wage path. If individuals reduce reservation wages in response to longer spells of unemployment, we would expect this

[21] In 2009/10 the Irish national minimum wage was €8.65 per hour. However, workers in their first year of employment after turning 18 are entitled to only 80% of the full rate, while those in the second year of employment are entitled to 90%.

effect to be negative. As before the key parameter of interest is ϕ, which measures the shift in the post-unemployment wage path resulting from the cut in benefit payments.

The results of this model are given in Table 6. The coefficients are similar across age groups. The coefficients on unemployment duration are of interest as they show the change in re-employment wages associated with longer unemployment durations. As expected, longer spells of unemployment reduce post-unemployment wages. However, the effect sizes are relatively modest. An additional year of unemployment duration reduces post-unemployment wages for 18 year olds by approximately €8.28, or 2.8%. The key parameter of interest is ϕ, which measures the shift in the post-unemployment wage path from the cut in benefit payments. For both age groups, ϕ is small and statistically insignificant. This is consistent with the findings reported in Figure 8. Together these results indicate a limited role for lower reservation wages in explaining the faster exits to work in response to the benefit cut. This is plausible since the accepted wages of this group are already close to the minimum wage rate, so there is limited scope for reducing reservation wages.

8. Long-Run Effects

We conclude our analysis by considering 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 will not only find jobs more quickly but will remain in those jobs for longer. 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.[22] This could lead to substantial churning between states,

[22] 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.

weakening 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 carry out our long-run analysis we repeat our earlier RD analysis for a number of long-run outcomes. For those exiting to work, we first examine the duration of the subsequent employment spell. 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 7 and indicate a positive but statistically insignificant effect of the benefit cut on subsequent employment duration. We also examine the impact of the 2009 benefit cut on outcomes in 2014, the last year for which new unemployment spells are recorded in our data. We consider three outcomes: 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 given in the last three columns of Table 7, 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 almost five weeks; and increased average weekly wages in 2014 by €26.45. However, none of these long-run effects is statistically significant. It may be that the severity of the recession meant that the negative signal associated with unemployment was weaker during this period, so that longer initial spells were not heavily penalised.

9. Conclusion

This paper evaluates 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 evidence on the benefit responsiveness of very young labour market participants, a group that is of particular policy interest. Our analysis is facilitated by access to high quality administrative data on the population claimants and by the quasi-experimental of the benefit cut, which resulted in claimants whose unemployment start dates differed by one day receiving very different benefits.

We find that the benefit cut substantially reduced unemployment duration for 18 year olds. For JA claimants in 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 young claimants, even during a recession. To examine the effects of the benefit cuts in more detail, we decompose the overall effect into the components due to different exit states. Despite the depressed nature of the labour market, a significant proportion of the treatment effect is accounted for by earlier exits to work. Exits to training are similarly important, but there is no evidence that claimants were forced out of the labour market by the cuts.

When we examine the wages earned by those who exit to work, we find that these young claimants are typically taking low-paid, minimum wage jobs. Furthermore, we find that the benefit cut had no significant effect on the wage rates accepted, implying a limited role for lower reservation wages in explaining the faster exits to work. This is plausible given that the scope for reducing reservation wages is limited by the fact that the accepted wages of this group were already close to the minimum wage rate even before the benefit cut.

Despite the fact that there was a substantial immediate effect on unemployment duration of the benefit cut, there are some reasons for caution. We find little evidence that this short-run effect translated into a significant long-run effect in terms of incidence of unemployment and weekly wages five years after the introduction of the cuts. However, as noted earlier, the long-run effects may have been stronger had the benefit cuts been introduced in better labour market conditions.

As well as finding limited long-run effects, our analysis takes no account of other potential negative effects of the cuts. It is plausible that a benefit cut of this magnitude had negative consequences for consumption, reducing the ability to consumption smooth and increasing claimants' dependence on family members. This in turn may have led to increased pressure on low-income families. For those without family support, there is anecdotal evidence of an increase in homelessness affecting those whose benefits was cut. Further research is needed on these negative effects, which, when combined with the large positive incentive effects found in our paper, would provide the evidence needed to determine the appropriate benefit rates for young people.

Table 1: Variable Means for New Claimants by Age and Date of Entry to Unemployment

	Age 18		Age 19	
	Six Months *Before* Threshold	Six Months *After* Threshold	Six Months *Before* Threshold	Six Months *After* Threshold
Nationality Irish	0.95	0.94	0.94	0.94
Low Education	0.49	0.40	0.35	0.30
No Previous Employment Spell	0.40	0.52	0.19	0.29
Affected by Benefit Cut	0.02	0.87	0.01	0.56
Unemployment Spell Duration (Weeks)	115	80	90	74
N	3430	2893	2418	2180

Notes: Threshold date is April 29, 2009. Low education indicates not having completed secondary schooling; education data are available for approximately 90% of claimants

Table 2: Fuzzy Regression Discontinuity Results for Benefit Cuts
Standard Errors in Parentheses

	Age 18	Age 19
First Stage: Effect on Proportion Treated	0.65*** (0.05)	0.37*** (0.07)
Effect of Treatment on Unemployment Duration	**-60.96** (23.11)**	**-49.81 (37.92)**
N	6323	4598
Elasticity Calculations		
Mean Unemployment Spell Duration Before Treatment (Weeks)	114.77	90.42
Estimated Duration Change (%)	-53.11	-55.09
Benefit Change (%)	-50.9	-50.9
Estimated Elasticity	1.04	1.08

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

Table 3: Fuzzy Regression Discontinuity Results for Differences in Claimant Characteristics
Standard Errors in Parentheses

	Age 18	Age 19
Low Education	-0.05 (.12)	0.04 (.17)
Nationality Irish	-0.04 (.05)	-0.05 (.13)
No Previous Employment Spell	0.17 (0.11)	-0.06 (0.18)

Table 4: Difference-in-Difference Hazard Function Results
Standard Errors in Parentheses

	Age 18	Age 19
Treatment Month	-0.002 (0.082)	-0.007 (0.095)
Treatment Year	0.016 (0.084)	0.013 (0.098)
Treatment Month x Treatment Year	**0.317***** **(0.110)**	**0.105** **(0.129)**
Nationality Irish	-0.178 (0.118)	-0.065 (0.142)
Low Education	-0.522*** (0.057)	-0.283*** (0.071)
No Previous Employment Spell	-0.367*** (0.055)	-0.229*** (0.079)
t	-0.012*** (0.001)	-0.014*** (0.001)
$t^2/100$	0.004*** (0.0002)	0.004*** (0.0003)
Constant	-3.622*** (0.137)	-3.662*** (0.154)
N	1388	986

Notes: Reference year for Difference-in-Difference estimation is one year earlier in each case. *** Denotes significant the 1% level. ** Denotes significant at the 5% level. * Denotes significant at the 10% level.

Table 5: Competing Risks Decompositions

	Age 18	Age 19
Overall Treatment Effect	-58.29	-40.16
Decomposition		
Training & Education	-18.74	-15.95
Work	-20.17	-11.70
Inactivity	-4.49	-2.32
Other	-14.89	-10.19

Table 6: Difference-in-Difference Post-Unemployment Wage Model
Standard Errors in Parentheses

	Dependent Variable Weekly Wage	Dependent Variable Weekly Wage
Treatment Month	-28.42*** (8.52)	-18.04 (11.17)
Treatment Year	-14.64 (8.37)	-14.86 (10.88)
Treatment Month x Treatment Year	**14.34 (11.78)**	**-4.16 (14.92)**
Nationality Irish	25.32* (14.37)	26.54 (18.00)
Low Education	11.04* (6.05)	9.14 (8.22)
No Previous Employment Spell	-12.28* (6.82)	-6.04 (11.81)
Dur	-0.81* (.46)	-2.29** (0.62)
Dur^2	0.01 (.007)	0.025*** (.009)
Constant	298.88*** (15.72)	325.75*** (19.78)
N	2267	1958

Notes: Reference year for Difference-in-Difference estimation is one year earlier in each case. *** Denotes significant the 1% level. ** Denotes significant at the 5% level. * Denotes significant at the 10% level.

Table 7: Fuzzy Regression Discontinuity Results on Long-Run Outcomes, 18 Year Olds
Standard Errors in Parentheses

Duration of Subsequent Employment Spell	Incidence of Unemployment, October 2014	Total Unemployment, 2014 (Weeks)	Weekly Wages, 2014
8.81 (40.93)	-0.06 (0.08)	-4.67 (3.77)	26.45 (38.83)

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 April 29 2009

(a) 18 Year Olds

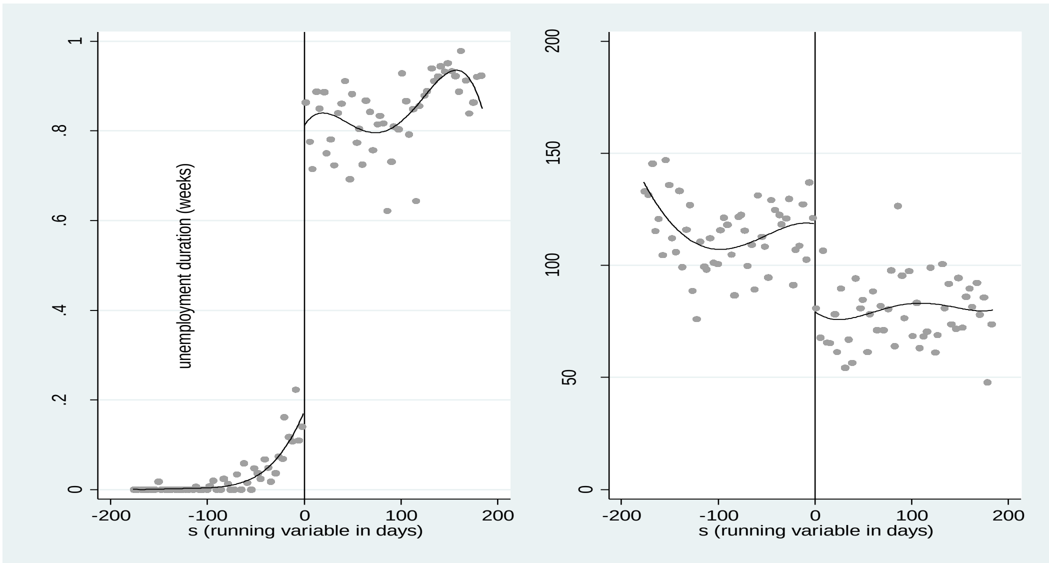

(b): 19 Year Olds

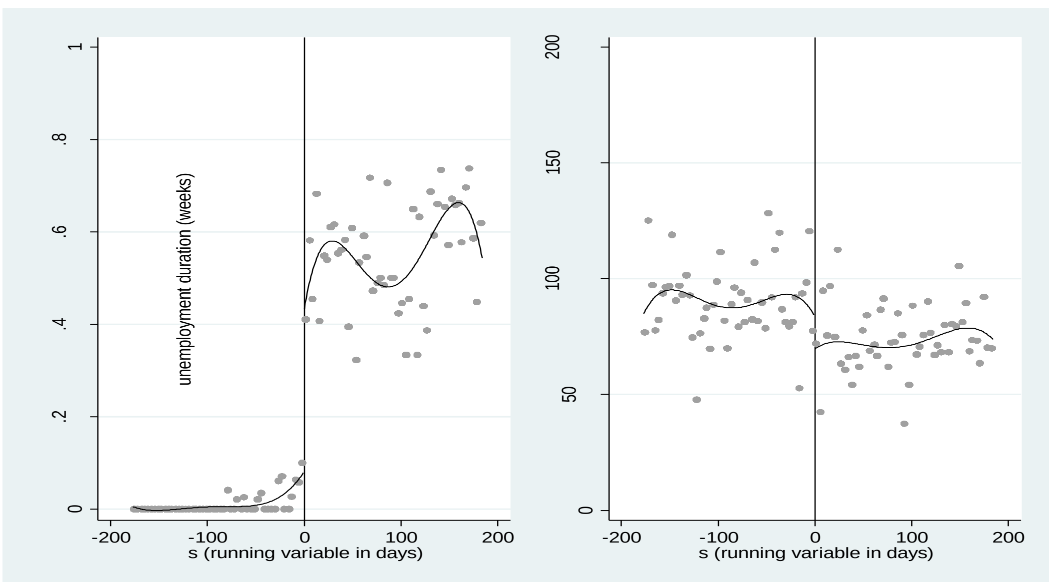

Figure 2: Regression Discontinuity Graphs of Average Unemployment Duration, 18 Year Old Entrants to Unemployment Six Months Before and After April 29 2008 (left panel) and April 29 2010 (right panel)

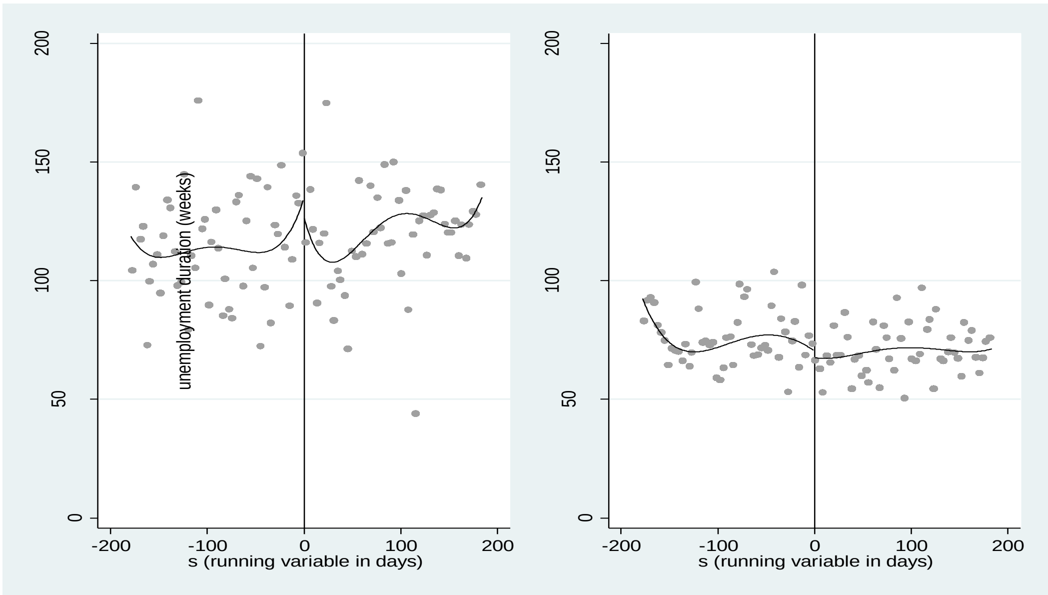

Figure 3: Regression Discontinuity Graphs of Average Unemployment Duration, 18 Year Old Entrants to Unemployment Six Months Before and After March 29 2009 (left panel) and May 29 2009 (right panel)

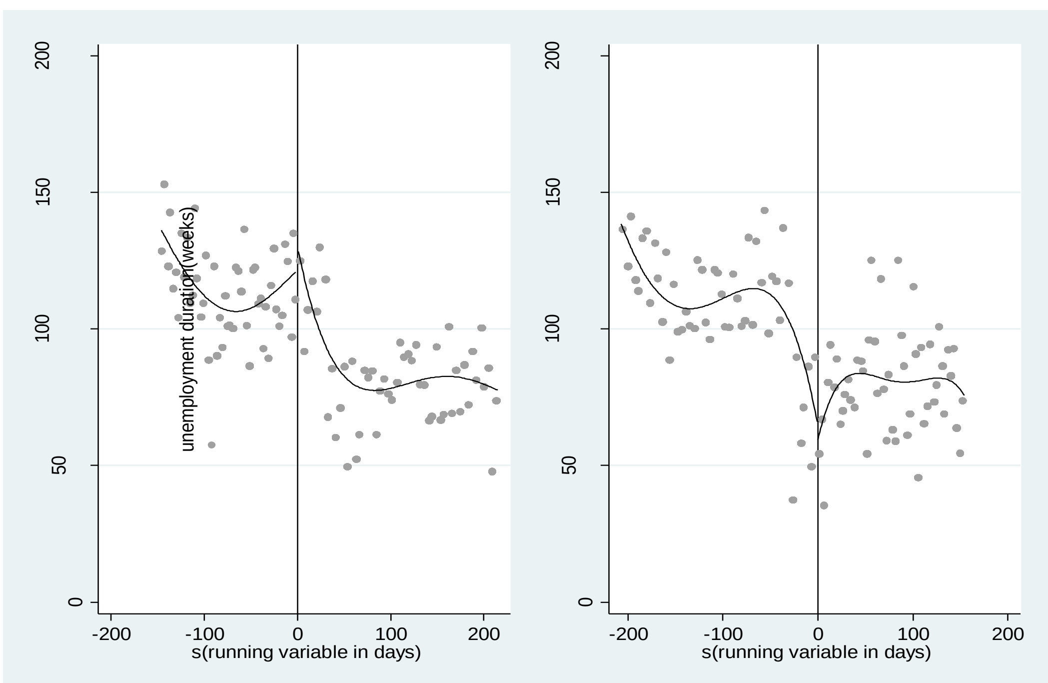

Figure 4: Regression Discontinuity Graph of Average Unemployment Duration, 18 Year Old Entrants to Unemployment Six Months Before and After April 29 2009 for JB Claimants

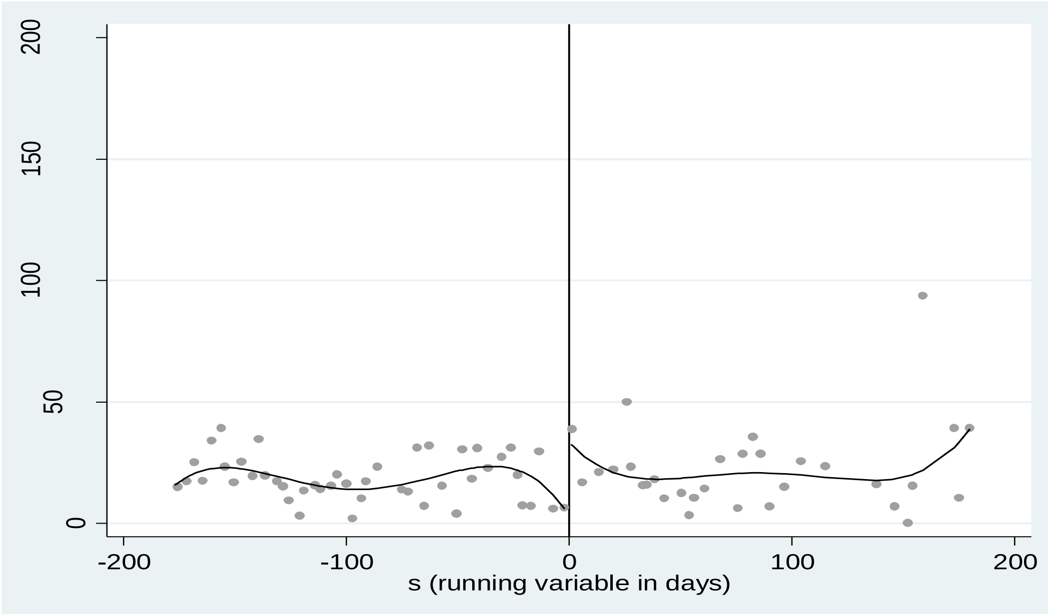

Figure 5: Regression Discontinuity Graph of Density of Entries to Unemployment, 18 Year Old Entrants to Unemployment Six Months Before and After April 29 2009

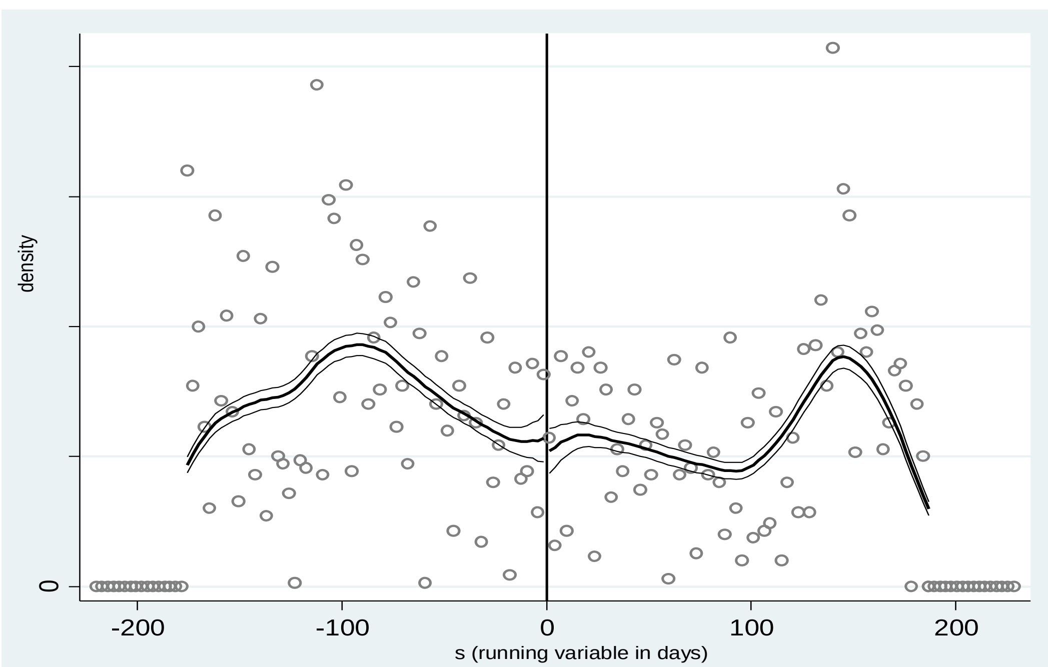

Figure 6: Kaplan-Meier Unemployment Exit Hazard Functions, 18 Year Old Entrants to Unemployment One Month Before and After April 29, 2008 (left panel) and 2009 (right panel)

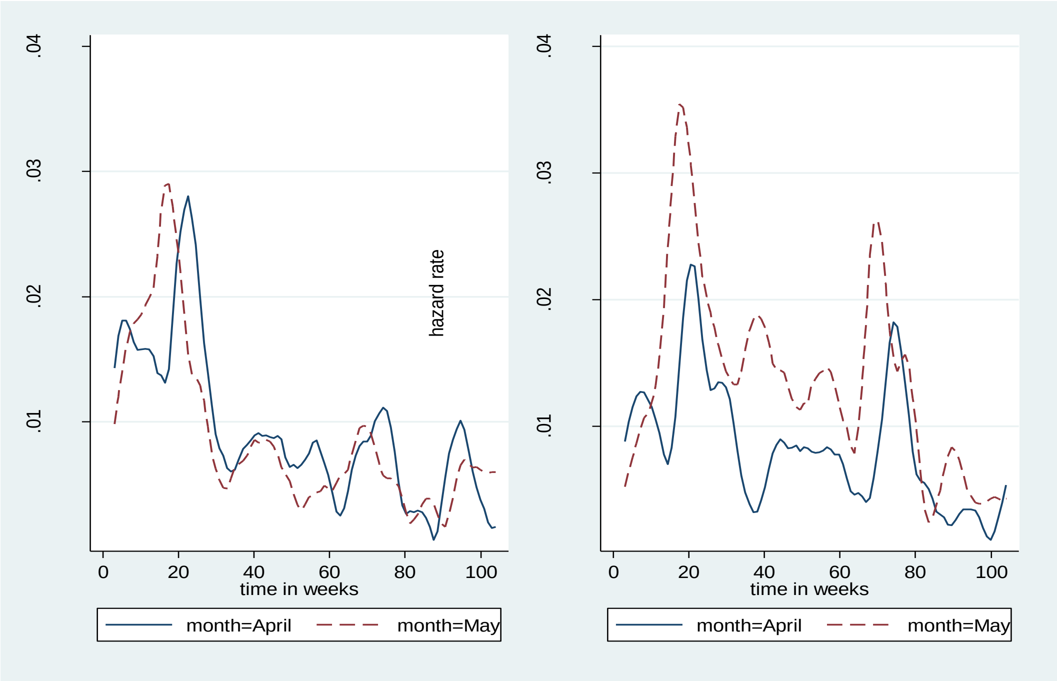

Figure 7: Kaplan-Meier Unemployment Exit Hazard Functions, 19 Year Old Entrants to Unemployment One Month Before and After April 29, 2008 (left panel) and 2009 (right panel)

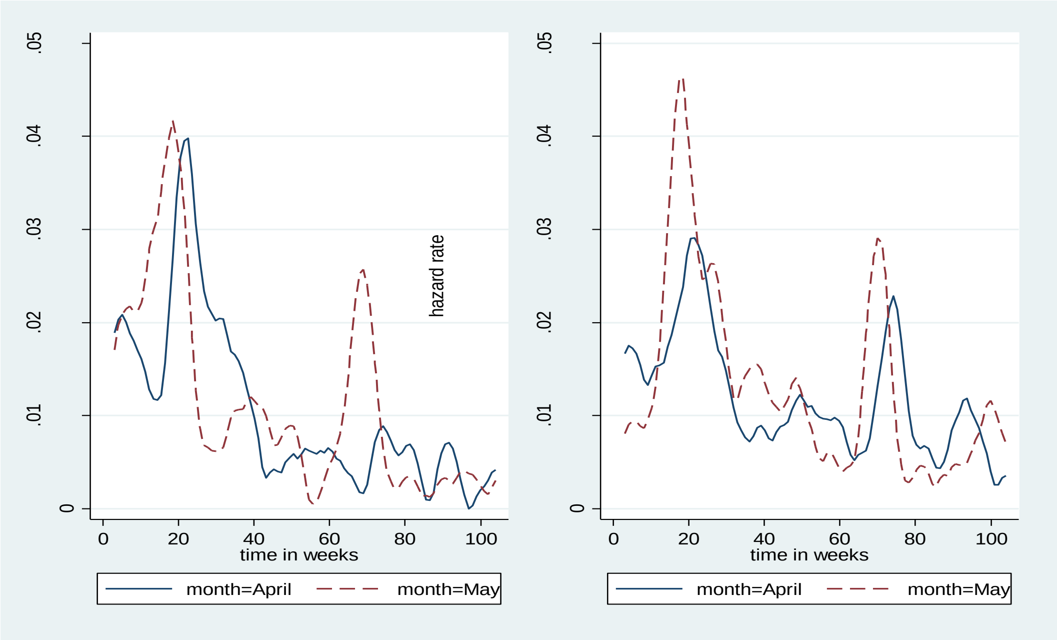

Figure 8: Kernel Densities of Weekly Wages in Year of Exit from Unemployment by Treatment and Control Groups, 18 and 19 Year Olds

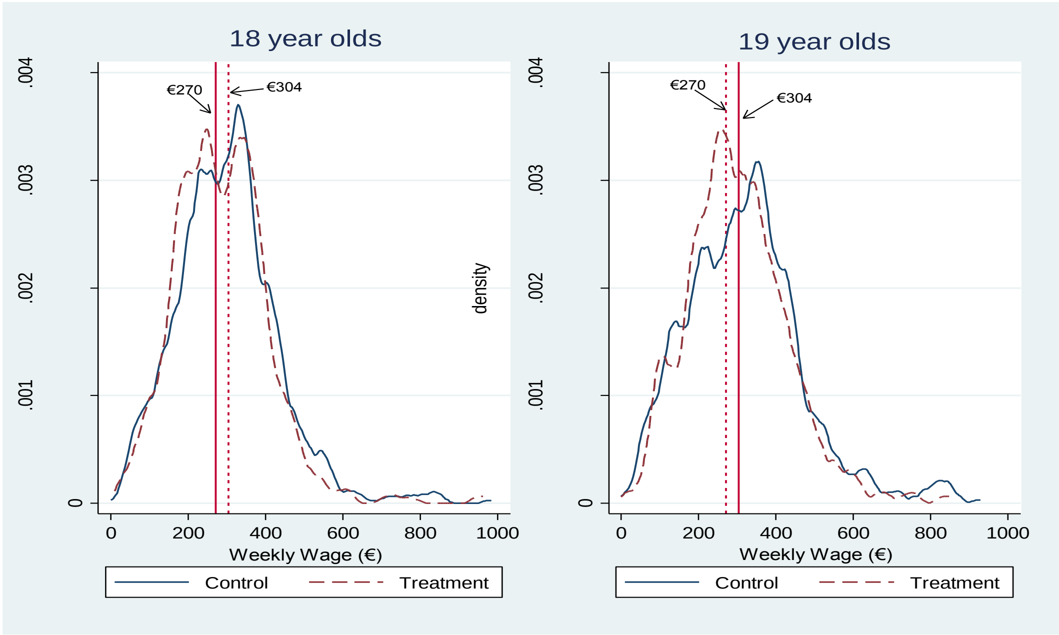

Table A1: Fuzzy Regression Discontinuity Results for Women
Standard Errors in Parentheses

	Age 18	Age 19
First Stage: Effect on Proportion Treated	0.790 0.055	0.440 0.066
Effect of Treatment on Unemployment Duration	**-51.99***** **16.98**	**-37.89** **28.92**
N	3591	2867
Elasticity Calculations		
Mean Unemployment Spell Duration Before Treatment (Weeks)	103.0	88.54
Estimated Duration Change (%)	-50.47	-42.79
Benefit Change (%)	-50.9	-50.9
Estimated Elasticity	0.99	0.84

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.

ATHEY, S. and G. IMBENS (2017): "The State of Applied Econometrics: Causality and Policy Evaluation," *Journal of Economic Perspectives*, 31, no. 2 Spring, 3-32.

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. SAKSONOVS, 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.

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.

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.

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 KLAAUW (2001): "Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design," *Econometrica*, 69, 201-209.

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*, 79, 933-959.

IMBENS, G. W., and T. LEMIEUX (2008): "Regression Discontinuity Designs: A Guide to Practice," *Journal of Econometrics*, 142, 615-635.

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.

LANDAIS, C. (2015): "Assessing the Welfare Effects of Unemployment Benefits Using the Regression Kink Design." *American Economic Journal: Economic Policy*, 7(4): 243-78.

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): 05885"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*.

REBELLO-SANZ, Y., and N. RODRIGUEZ-PLANAS (2016): "When the Going Gets Tough… Financial

Incentives, Duration of Unemployment and Job Match Quality," *IZA DP 10044*
SCARPETTA, S.,. 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 DEN BERG. G., A. UHLENDORFF and J. WOLFF (2017): "Under Heavy Pressure: Intense Monitoring and Accumulation of Sanctions for Young Welfare Recipients in Germany," *IZA DP 10730*.
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