ESRI ECONOMIC & SOCIAL RESEARCH INSTITUTE

ESRI Special Article

Did increasing the state pension age in Ireland affect the retirement rate of 65-year-olds?

P. Redmond, S. McGuinness and E. Kelly

DID INCREASING THE STATE PENSION AGE IN IRELAND AFFECT THE RETIREMENT RATE OF 65-YEAR-OLDS?

Paul Redmond, Seamus McGuinness and Elish Kelly[1]

ABSTRACT

In January 2014, the qualifying age for the Irish contributory state pension increased from 65 to 66 years. Individuals born after 1 January 1949 could no longer qualify for the pension at age 65, while individuals born before this date could still qualify, provided they had the required social insurance contributions. In this paper, we examine whether this change in the qualifying age had a causal effect on the retirement rate of 65-year-olds in Ireland. To do this, we compare the retirement rates of two groups of 65-year-olds in 2014; one group was born just after the cut-off date, thereby making them ineligible for the state pension at age 65, while the other group was born just before the cut-off date, making them potentially eligible, subject to meeting the insurance contribution requirements. We do not find clear evidence that the change in the retirement age had a causal effect on the retirement rate or the employment and unemployment rates of 65-year-olds in 2014.

INTRODUCTION

Prior to 2014, the transition state pension in Ireland was payable at age 65 to individuals who retired from insurable employment and satisfied certain social insurance contribution conditions. The transition state pension was payable for only one year from age 65, after which the person was automatically transferred to the contributory state pension at the age of 66.[2] In 2014, the transition state pension was abolished, thereby increasing the pension qualification age to 66 years.[3] The implementation of the policy in 2014 was based on a person's date of birth, such that individuals born before January 1949 could still qualify at age 65, whereas those born on or after 1 January 1949 had to wait until age 66. As a result, a 65-year-old in 2014 who was born in December 1948 and had the required social insurance contributions could receive the state pension, while a

1 The work carried out in this article was funded by the Department of Social Protection and we would like to thank all individuals in the Department who provided assistance during the project. We would also like to thank the Central Statistics Office (CSO) for providing us with access to the Quarterly National Household Survey (QNHS) with the additional month of birth information required to conduct this study. Finally, we are grateful to Alan Barrett and an anonymous referee for providing valuable comments on an earlier draft.

2 There also exists a non-contributory state pension in Ireland which is a means-tested payment to individuals who are aged 66 and over and do not qualify for the contributory state pension.

3 These changes were outlined by the Irish government in March 2010, thereby giving individuals affected by the policy change almost four years advance notice of the pension age change.

65-year-old individual born in January 1949 (one month younger) with the same contributions had to wait until age 66 to receive the state pension.

We exploit this sharp cut-off in the pension qualification age to analyse whether the policy change had a causal effect on the retirement rate of 65-year-olds in 2014 using a regression discontinuity design. This is undertaken using data from the Quarterly National Household Survey (QNHS), Ireland's labour force survey, which is conducted by the Central Statistics Office (CSO). The identification strategy can be explained in a straightforward way. Consider a scenario where we have data in 2014 on two groups of individuals; one group was born on 31 December 1948 and the other on 1 January 1949. Both groups are 65, with one group being only one day older. However, the older group qualifies for the pension at 65 while the younger group does not. Comparing the retirement outcomes of both groups allows us to assess whether the policy change had a causal effect on the retirement rate. In addition to being virtually the same age, there is no reason to suspect that these two groups will be systematically different with respect to other characteristics (both observable and unobservable) that might impact on the retirement decision. Therefore, any difference in retirement outcomes can be attributed to a causal effect of the policy.

While the policy change lends itself to a regression discontinuity analysis, it is important to point out some data limitations which impact the study. The finest level of information available to us on an individual's birth date was their month of birth. As such, individuals are grouped into discrete monthly bins and identification relies on comparing the outcomes of those born in the month prior to the cut-off date, i.e., just qualified (based on age), to those born in the month after the cut-off date, i.e., barely missed out. As we are focusing on a narrow subset of the population, there is a relatively small number of observations. For example taking a three-month period on either side of the cut-off provides us with 431 observations; 238 individuals who just missed out on qualifying, i.e., were born between January and March 1949, and 193 individuals who just barely qualified, i.e., were born between October and December 1948.[4] Small sample sizes are common in RD studies due to the narrow focus on a subset of the population. Expanding the analysis to include individuals further away from the threshold increases the sample size. However, given that identification relies on comparing people to the immediate left and right of the cut-off, the estimates should not be overly reliant on these observations. In this paper, we address this

[4] These are the unweighted sample sizes. There is a weighting factor in the data which is designed to ensure it is representative of the population. We verify the results do not change when we estimate discontinuities using the weighting factors.

issue by verifying the robustness of our results to a restrictive specification which focuses on a very narrow bin width on either side of the month of birth cut-off.

An additional consideration relates to the 'bite' of the policy among the sample of individuals studied, i.e., the number of 65-year-olds in 2014 that were actually affected by the policy change. Given the transition state pension was based on satisfying certain social insurance contributions, not everybody in the sample will be affected by the policy change. Conditioning only on 65-year-olds who were entitled to the state pension would require a larger and more detailed dataset that would allow us to identify people with the required insurance contributions and to ensure there was a large enough sample of these individuals for meaningful estimates. Ideally, this would be a linked administrative dataset between Revenue and the Department of Social Protection, which, in addition to providing details on the types of benefits received, would show an individual's employment history. Should the data be made available, future research could examine the causal effect of the pension age change on the retirement decision, conditioning on individuals who had the required social insurance contributions.[5] However, in this study, we examine the retirement rates of all 65-year-olds in 2014. Therefore our analysis provides evidence of the causal effect of the policy change on the overall retirement rate of 65-year-olds. In light of the data issues highlighted above, an additional contribution of this paper is to highlight potential avenues for further work and to make suggestions as to the type of data that would facilitate research which would be useful to inform the policy debate in this area.

While the data do not allow us to identify a person's eligibility based on their social insurance contributions, we make use of the QNHS question which captures information on whether a person has ever been in employment.[6] If a person has never been employed then they will not have the required contributions to qualify for the transition state pension. Therefore, in addition to reporting estimates of the causal effect of the policy change on the overall retirement rate of 65-year-olds, we also report results where we condition only on individuals who have some previous employment experience. While this does not fully address the data issues outlined above, it goes some way towards narrowing the analysis to individuals who are likely to qualify for the transition state pension by focusing only on those who have some social insurance contributions. However, only 7 per cent of all 65-year-olds in our analysis report having no previous employment experience and excluding these individuals from

5 Such information (social insurance contributions) is not collected by the QNHS data, which are the data used to conduct this analysis.

6 There is also information on when an individual started work for their current employer. However, we cannot use this as it does not give information on their full employment history.

the analysis has very little impact on the estimates and does not change the overall results of the paper.

Relatively few studies examine how changes in retirement age rates affect retirement decisions. The research which does address this question utilises large administrative datasets (Staubli and Zweimuller, 2013; de Grip et al., 2013; Puur et al., 2015; Sanchez-Martin et al., 2014; Vestad, 2013). Staubli and Zweimuller (2013) examine the labour market effects of increasing the early retirement age in Austria and find that raising the retirement age increased employment of affected men by 9.75 percentage points and affected women by 11 percentage points. They also find that a large number of affected individuals bridge the gap to retirement by drawing on unemployment benefits; specifically, there was a 12.51 percentage point increase in registered unemployment among men and 11.77 percentage points among women. Other work has shown that increasing the pension age increases the actual retirement age (Puur et al., 2015) as well as the expected retirement age of employees (de Grip et al., 2013) and can lead to increases in labour supply (Vestad, 2013; Sanchez-Martin et al., 2014) and lower pension costs (Sanchez-Martin et al., 2014).

The paper proceeds as follows. In the next section we describe the data and present some descriptive statistics relating to the sample of 65-year-old individuals in 2014. The following section outlines the methodology and the Results section presents the main results as well as various robustness and sensitivity checks. The final section concludes and discusses some potential explanations for the lack of any clear evidence of a causal effect of the policy change on retirement rates.

DATA

We use data from the 2014 Quarterly National Household Survey (QNHS). Given that we are studying the effect of a change to the pension qualification age, and that identification in the RD design relies on people close to the qualification threshold, we focus our analysis on an older segment of the population. Individuals who were not close to 65 years at the time the policy was implemented will not be affected by the policy change, and therefore these individuals should not influence our estimates.

There is a trade-off when it comes to choosing the sample of individuals to include in an RD study. Focusing on observations that lie very close to the assignment threshold is beneficial as it is these types of individuals upon which identification hinges, however this can often lead to a very small sample size. Expanding the analysis to include individuals further from the threshold can increase the sample size, however, the results should not be overly reliant on these individuals. The age threshold in this study is whether a person was born before or after 1 January 1949. An alternative way of defining the threshold is a person's age (in months) at the time the policy was implemented. People who were at least 780 months old at January 2014 were born before 1 January 1949 and qualify, while those aged 779 months or younger were born after. In this paper we take two approaches. We begin our analysis with a broad age range by including individuals aged between 681 months (56.75 years) and 879 months (73.25 years) at the time the policy was implemented (on 1 January 2014). This gives us a relatively large sample size of 14,911 individuals. We then verify whether the results from our baseline specification are robust to an alternative age-restricted specification, which focuses only on individuals who were aged 65 in 2014, i.e., people who were aged between 768 and 791 months at the time the policy was introduced. This reduces our sample size from 14,911 to 1,829; there were 940 individuals aged 65 in 2014 whose date of birth meant they missed out by between one and 12 months on qualifying for the pension at age 65. There were 889 individuals aged 65 in the older 12 month age range who could, based on their age, qualify for the pension at age 65.

Our dependent variable is the probability of being in retirement and is based on an individual's self-reported main labour status. Table 1 shows the distribution of 65-year-olds in 2014 by main labour status. Almost 40 per cent of 65-year-olds were retired from employment in 2014, making this the largest category. A large number were also working for payment or profit (27 per cent) and engaged in home duties (21 per cent). Approximately 10 per cent report being unable to work due to permanent sickness or disability while almost 4 per cent are unemployed, having lost or given up their previous job.

Table 2 splits 65-year-olds in 2014 into two groups; those whose date of birth meant they could not qualify for the state pension at age 65 (born after 1 January 1949) and those whose date of birth meant they could still potentially qualify for the pension (born before 1 January 1949). Descriptive statistics are presented for both groups as well as statistics on their main labour status. The retirement rate of the group born before 1 January 1949 is 43 per cent compared to 36 per cent for the group born after 1 January 1949. However the difference between the two rates cannot be taken as a causal effect of the change in qualification age on retirement rates. The group born before 1 January 1949 are, on average, eight months older than the group born after 1 January 1949 and, as such, we would expect their retirement rates to be higher. This serves to motivate the benefits of using the regression discontinuity design, which overcomes this issue by comparing the average retirement rates of individuals just to the left and right of the qualification threshold who are closer in age.

TABLE 1 MAIN LABOUR STATUS OF 65-YEAR-OLDS IN 2014

Main Labour Status	Frequency	Per Cent
Working for payment or profit	538	26.58
Looking for first regular job	*	*
Unemployed, having lost or given up previous job	75	3.71
Actively looking for work after voluntary interruption of working life for personal or domestic reasons	*	*
Student or pupil	*	*
Engaged in home duties	433	21.39
Retired from employment	786	38.83
Unable to work due to permanent sickness or disability	180	8.89
Other	*	*

Source: CSO Quarterly National Household Survey.
Note: Estimates for numbers of persons or averages where there are less than 30 persons in a cell are not produced as estimates are too small to be considered reliable.

TABLE 2 AVERAGE CHARACTERISTICS OF 65-YEAR-OLDS IN 2014

	Born before 1 Jan 1949	Born after 1 Jan 1949
Gender (% male)	51.6%	49.4%
Highest educational attainment (ISCED 11)	2.61	2.73
Married	70.3%	73.9%
Widowed	9.3%	7.3%
Age at 1 Jan 2014 (months)	784	776
Main Labour Status		
Working for payment or profit	25.1%	26.5%
Looking for first regular job		*
Unemployed, having lost or given up previous job	*	*
Actively looking for work after voluntary interruption of working life for personal or domestic reasons		*
Student or pupil	*	*
Engaged in home duties	21.4%	22.1%
Retired from employment	43.0%	35.6%
Unable to work due to permanent sickness or disability	7.7%	10.3%
Other	*	*

Source: CSO Quarterly National Household Survey.
Note: Estimates for numbers of persons or averages where there are less than 30 persons in a cell are not produced as estimates are too small to be considered reliable.

METHODOLOGY

Regression discontinuity is a quasi-experimental research design that allows for the causal analysis of a treatment, when assignment to that treatment changes discontinuously at a pre-defined threshold.[7] Regression discontinuity analysis can be implemented either parametrically or non-parametrically (Lee and Lemieux, 2010). The parametric approach involves fitting a conditional mean function to the data on either side of the cut-off that determines treatment, using polynomials of various orders. The non-parametric approach is based on estimating a regression function in a neighbourhood of the cut-off, using a specified bandwidth and kernel. With a large number of observations and continuous data, the non-parametric estimator is desirable as it focuses the analysis on observations close to the cut-off (Skovron and Titiunik 2015; Gelman and Imbens 2014). However with discrete data that are reported in coarse intervals, such as the month of birth data used in this study, non-parametric methods may be of limited use. As noted by Lee and Card (2008), with coarse data an irreducible gap exists between the treatment group just above the cut-off and the control group just below, and as such it may not be possible to estimate a causal effect in the absence of a parametric assumption.

The treatment under investigation in this study is whether or not an individual could qualify, provided they had the required social insurance contributions, for the transition state pension in Ireland at age 65. The variable which determines treatment assignment is known as the forcing variable, which in this case is date of birth. We were unable to source day of birth data, however we have data on month of birth. We can set up the forcing variable as an individual's age, in months, at January 2014 (the implementation date of the policy). Those born in December 1948 are 780 months old in January 2014 and those born in January 1949 are 779 months. Therefore, *Ti* is defined as a treatment dummy which indicates whether an individual can qualify, based on their month of birth, for the pension at age 65 in 2014, such that,

$$Ti = \begin{cases} 1 \text{ if age of individual } i \text{ at Jan 2014} \geq 780 \text{ months} \\ 0 \text{ if age of individual } i \text{ at Jan 2014} < 780 \text{ months} \end{cases}$$

As such, the treated group are those individuals who can potentially qualify for the state pension at age 65, while the untreated are those who do not qualify. Any discontinuity in outcomes that exist as a result of the treatment should be evident at the 780 month threshold, i.e. a sharp jump in the probability of retirement for the 780 month group compared to the 779 month group. We focus our analysis on data from Quarters 2 and 3 of 2014. This ensures that the 780 month and 779 month groups, upon which identification hinges, are both

7 For a detailed exposition of the regression discontinuity estimator and its properties, see Lee and Lemieux (2010).

aged 65 at the time of survey. If we were to include data for Quarter 4, 2014, then it could be the case that some of the 780 month group would be 66 at the time of survey, whereas all of the 779 month group would be 65. Likewise, including Quarter 1 data from 2014 would mean that some of the 779 group would still be 64 at the time of survey.[8]

We employ a parametric regression discontinuity specification, which involves running separate regressions on both sides of the threshold of the outcome of interest on the forcing variable. We estimate the following regressions for those to the left of the 780 month threshold, i.e. the individuals who do not qualify for the state pension until age 66, and those to the right of the threshold, i.e. the individuals who could qualify, based on their month of birth, for the pension at age 65 (the treated),

$$RetiredL_i = \alpha_L + \beta . f_L(AgeatJan14_i - 780) + \varepsilon_i \tag{1}$$

$$RetiredR_i = \alpha_R + \beta . f_R(AgeatJan14_i - 780) + \varepsilon_i \tag{2}$$

where f_L and f_R are polynomials in the forcing variable. The inclusion of the polynomial terms underlies the importance of getting the specification correct in order to avoid mistaking a non-linearity in the conditional expectation function for a discontinuity. The threshold value of 780 is subtracted from the forcing variable for convenience; the estimated discontinuity is then simply the difference between the intercepts, $\alpha_R - \alpha_L$. Instead of estimating two separate regressions, it is straightforward to estimate a single pooled regression which gives identical results and has the advantage of yielding a direct estimate of the discontinuity and standard errors. The pooled regression is

$$Retired_i = \alpha + \beta . T_i + \rho . f(Age_i - 780) + \gamma . T_i . f(Age_i - 780) + \varepsilon_i \tag{3}$$

where *Ti* is the treatment dummy defined above and β is an estimate of the causal effect of the pension age change on the probability of retirement. An interaction term between the treatment dummy and the polynomial is also included.

RESULTS

As a first step, before running the regression model (Equation 3), we calculate average outcomes for each of the discrete age points in the forcing variable. This is shown in Figure 1 for individuals aged between 681 months (56.75 years) and 879 months (73.25 years) at the time the policy was implemented (on 1 January

8 We verify that the results are robust to including all four quarters of data.

2014). The average retirement probability by age (in months) at January 2014 is calculated based on the main labour status of respondents (see Table 1). A dummy variable is created which equals one if the main labour status of respondents is 'retired from employment' and zero otherwise.[9] Averages of this dummy variable are calculated for each age category and are plotted in Figure 1. A vertical line is inserted in the graph to show the month-of-birth threshold. The average retirement probabilities are increasing with age, which explains the upward sloping averages. However, from this graph it is not immediately apparent whether a discontinuity exists at the threshold. Close inspection reveals that the groups who just barely missed out on treatment, upon which identification relies heavily, have similar retirement probabilities to those who just qualified. This becomes clearer when we look at Table 3 which lists the retirement probabilities for individuals depending on their age (in months) at January 2014. The retirement probability for the 778 group is roughly the same as the 780 group, and is higher than the 781-784 groups. Moreover, if we were to expand our bin widths to include two months of data, then the difference in retirement probabilities between those who barely qualified (the 780 and 781 group) would barely differ from those who just missed out (the 778 and 779 group). As such, there is no clear evidence that the change in the retirement age had a causal effect on the retirement rate.

FIGURE 1 AVERAGE RETIREMENT PROBABILITIES

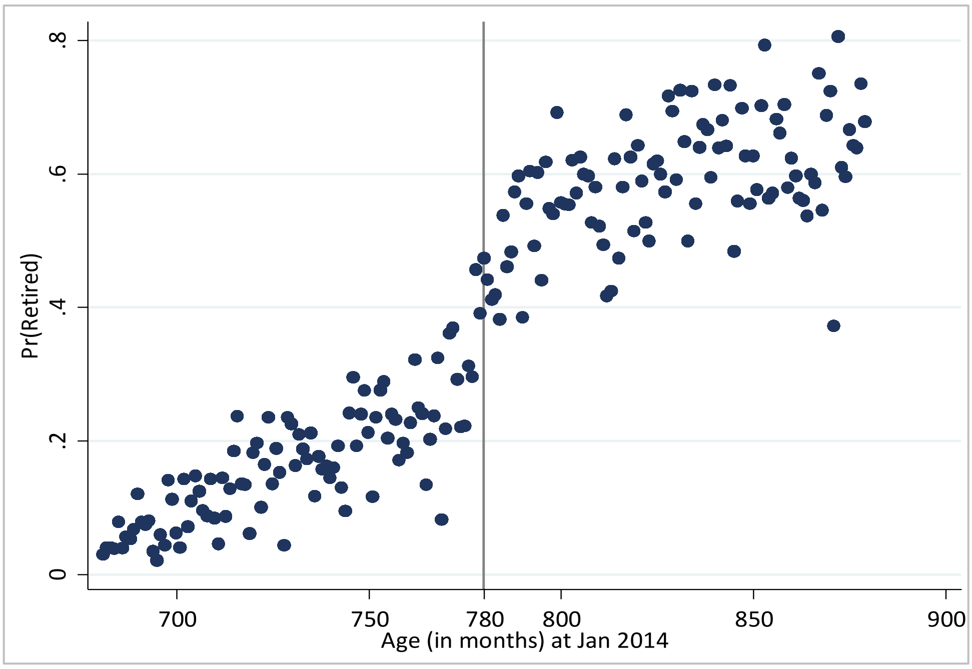

Source: CSO Quarterly National Household Survey.

9 We verify the results are robust to an alternative setup whereby the zeros in our retirement dummy are just those who are either employed or unemployed. This reduces our sample size and produces noisier estimates, however the main results remain unchanged.

TABLE 3 AVERAGE RETIREMENT PROBABILITIES BY AGE AT JANUARY 2014

Age at Jan 2014	Month and Year of Birth	Mean	Obs
771	Sept 1949	0.36	72
772	Aug 1949	0.37	73
773	July 1949	0.29	89
774	June 1949	0.22	77
775	May 1949	0.22	90
776	April 1949	0.31	80
777	March 1949	0.30	81
778	Feb 1949	0.46	70
779	Jan 1949	0.39	87
780	Dec 1948	0.47	57
781	Nov 1948	0.44	68
782	Oct 1948	0.41	68
783	Sept 1948	0.42	74
784	Aug 1948	0.38	89
785	July 1948	0.54	78
786	June 1948	0.46	89
787	May 1948	0.48	89
788	April 1948	0.57	82
789	March 1948	0.60	62

Source: CSO Quarterly National Household Survey.

We employ a parametric specification (Equation 3) using polynomials of different orders, beginning with a linear specification, i.e., a first order polynomial. The conditional expectation function and the local averages are shown in Figure 2. The estimate of the discontinuity from this linear specification is quite large and statistically significant and taken in isolation this would indicate that the individuals who barely missed out on the treatment (those just below the threshold) are 17.5 percentage points less likely to be in retirement in 2014 than those who just barely qualified for treatment. However, in parametric regression discontinuity specifications, a linear model is often not the most suitable specification and in order to be confident in the estimates, they should be robust to more flexible, higher order polynomial specifications. There is no fixed rule for choosing which order of polynomial is most suitable and often the local averages provide a guide as to which conditional mean function appears to be the best fit to the data. However, Lee and Card (2008) provide some guidance to evaluate whether a low order polynomial, such as the first order specification in Figure 2, may be too restrictive, thereby calling into question the reliability of such results. They suggest using a goodness of fit statistic;

$$G = \frac{(ESS_R - ESS_{UR})/(J-K)}{ESS_{UR}/(N-J)},$$

where ESS_R is the error sum of squares from the polynomial (restricted) model and ESS_{UR} is the error sum of squares from an unrestricted model which regresses the outcome variable on a full set of dummy variables for each of the discrete forcing variable bins. *J* denotes the number of bins, *N* the number of observations and *K* the number of parameters estimated in the restricted model. The statistic is distributed as *F(J-K, N-J)*. If the statistic exceeds the critical value, this implies that our low order polynomial may be too restrictive.

The goodness of fit statistic indicates that the linear specification used in Figure 2 is restrictive.[10] As such, the linear specification may be mistaking a non-linearity for a discontinuity. In Table A1 we present results using different polynomial specifications, ranging from a linear specification to a sixth order polynomial. The estimated discontinuity decreases quite quickly as we use more flexible functional forms, eventually disappearing in the fifth order specification. This relates to our previous discussion of Table 3, which shows that the retirement rates of individuals in the two bins just to the left of the threshold are very similar, or even larger, than those in the five bins just to the right of the threshold. This influences the conditional expectation function in the more flexible parametric functional forms, which can be seen in Figure 3, which plots the conditional expectation function of the fifth order polynomial. In this graph, no discontinuity exists at the threshold.[11]

10 The test statistic is 1.8 and the critical value is 1.17.

11 The fifth order polynomial is clearly a more flexible specification, yet the goodness of fit test statistic (1.6) is still above the critical value (1.18). As mentioned, the Lee and Card (2008) test is not a precise test as to which order polynomial to use, but rather should be applied when a significant result is detected to give reassurance of the model's suitability. The fact that the discontinuity disappears at the more flexible fifth order polynomial, which is still restrictive as per the Lee and Card (2008) test is conclusive. We can keep moving to more flexible, higher order polynomials, which would produce lower test statistics, but the results are still not statistically significant.

FIGURE 2 FIRST ORDER POLYNOMIAL SPECIFICATION

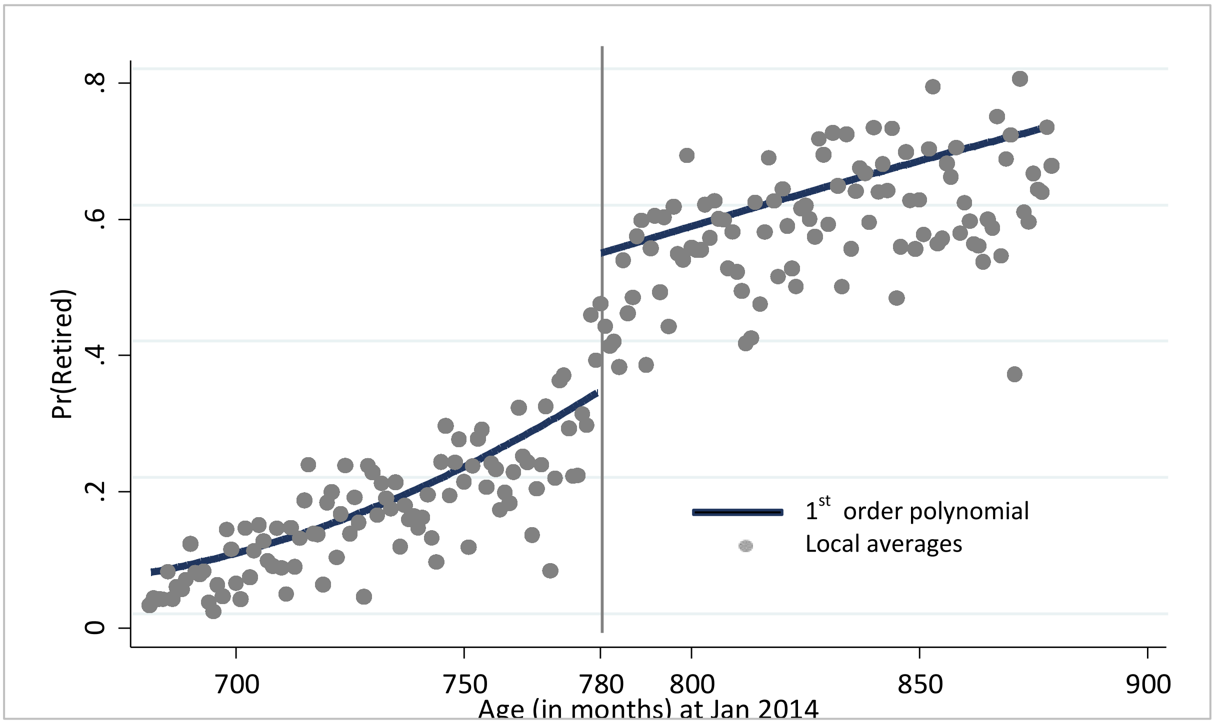

Source: CSO Quarterly National Household Survey.

FIGURE 3 FIFTH ORDER POLYNOMIAL SPECIFICATION

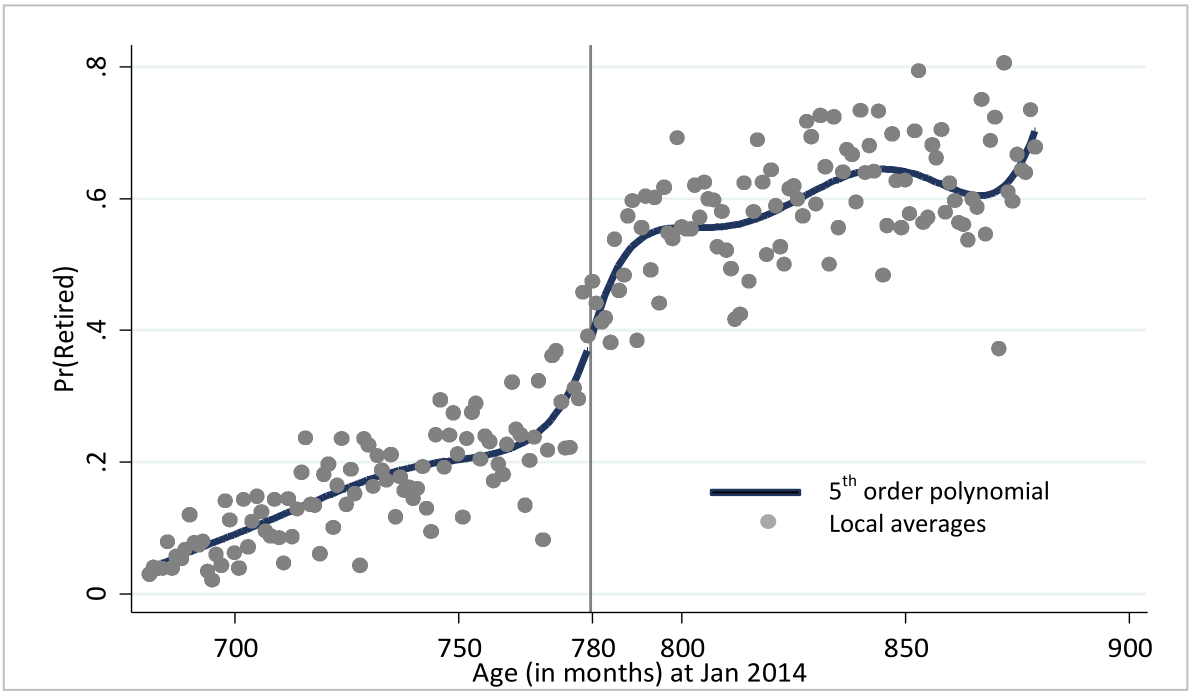

Source: CSO Quarterly National Household Survey.

As mentioned, the data are not ideally suited to non-parametric estimates due to the coarsely distributed forcing variable. Nonetheless, as an additional robustness check we also report estimates using a non-parametric local linear regression, as in Calonico et al. (2014). While the method used is local linear regression, our data are in monthly bins to the left and right of the threshold, and therefore even this non-parametric method relies on observations away from the cut-off, although these are given a lower weight.[12] The results of the non-parametric method indicate that there is no statistically significant effect of the policy change on retirement rates (see the last column of Table A1).

In the last row of Table A1, we report results of a model which conditions only on individuals with some previous employment experience. The coefficients remain relatively consistent with the baseline model, both in terms of their magnitude and statistical significance. As with the baseline model, the estimates are not statistically significant for higher order polynomials and the non-parametric specification.

Age-restricted specification

Identification in the RD design relies on comparing individuals close to the threshold. This is particularly relevant in our study, as individuals far from the threshold may not be affected by the policy change. For example, the pension age changing from 65 to 66 may have less of an impact on the retirement decision of a 60-year-old compared to a 65-year-old.[13] Likewise, a person who is already 66 years old (792 months) at January 2014 will not be affected by the policy change. There is an additional consideration with our data which relates to the age at which a person's outcome is observed. We use outcome data from Quarters 2 and 3 of 2014, which ensures that all individuals in a three-month bin to the left of the cut-off are 65 years of age. However, beyond this, there are some individuals who are 64 when surveyed.[14] This is shown in Table 4, which shows the average age (in years) at the time of survey (reference month) for each of the forcing variable months. For example, this indicates that all of the individuals who were 777 months old at January 2014 were 65 years old at the reference month, whereas the average age of individuals in the 776 group was 64.91 years, meaning some were still 64 years old when their labour status was recorded.

12 The local linear regression uses a triangular kernel and an optimal sized bandwidth as per Imbens and Kalyanaraman (2012).

13 It may affect their decision to retire as it will impact the amount of time they have to wait upon retirement before receiving the state pension.

14 For example, a person who was 774 months old at January 2014 and who was surveyed in April 2014 will still be 64 at the time of survey.

TABLE 4 AVERAGE AGE DURING REFERENCE MONTH

Age (in months) at Jan 2014	Average Age (in years) at Reference Month	N
771	64.08	72
772	64.30	73
773	64.43	89
774	64.61	77
775	64.74	90
776	64.91	80
777	65.00	81
778	65.00	70
779	65.00	87
780	65.00	57
781	65.00	68
782	65.00	68
783	65.07	74
784	65.29	89
785	65.44	78
786	65.57	89
787	65.75	89
788	65.94	82
789	66.00	62

Source: CSO Quarterly National Household Survey.

While the inclusion of a large number of bins either side of the threshold is useful for analysing the continuity of the conditional mean function and increasing the sample size, the main results should not be overly reliant on these observations. We examine whether our main results are consistent with an alternative specification in which we use all available data from 2014 and condition on individuals who are 65 when their labour status is recorded. This automatically discards individuals whose age is +/- 12 months of the 780 month threshold at January 2014, reducing the total sample size from 14,911 to 1,829. However, while our sample size is smaller, we are limiting our analysis to 65-year-olds and therefore we can be confident that our analysis is focusing on individuals whose age implies that they were potentially affected by the policy change. Table 5 shows the average retirement probabilities in 2014 of all individuals who were aged 65 in that year and Figure 4 displays the results graphically. Visual inspection of the average retirement probabilities suffices in this instance as we are including only 12 bins on either side of the cut-off, meaning parametric RD estimates of the conditional mean function provide little additional useful information. The slightly older 65-year-olds to the right of the threshold were potentially entitled to the state pension at age 65, whereas the slightly younger 65-year-olds to the left of the threshold were not. Again, this analysis does not provide convincing evidence of a causal effect of the policy change on the retirement rate of 65-year-olds in 2014. The average retirement rate of 65-year-olds who missed out on qualifying by at most two months was 41.6 per cent, compared to 45 per cent for the slightly older 65-year-olds who just barely qualified by at most two months. As before, the retirement rates of the 778 group, who just missed out, were equal to or higher than the rates for the 780-784 groups, which again suggests the absence of a policy effect.

TABLE 5 AVERAGE RETIREMENT PROBABILITIES OF 65-YEAR-OLDS IN 2014

Age (in months) at Jan 2014	Pr(Retired)	N
771	[0.45]	40
772	0.39	64
773	0.32	79
774	0.28	90
775	0.31	108
776	0.34	110
777	0.30	115
778	0.45	126
779	0.39	160
780	0.45	111
781	0.45	121
782	0.43	113
783	0.40	106
784	0.36	104
785	0.48	86
786	0.41	78
787	0.39	67
788	[0.48]	46

Source: CSO Quarterly National Household Survey.
Notes: Estimates for numbers of persons or averages where there are less than 30 persons in a cell are not produced as estimates are too small to be considered reliable. Parentheses [ ] indicate where there are 30-50 persons in a cell. Such estimates are considered to have a wider margin of error and should be treated with caution.

FIGURE 4 AVERAGE RETIREMENT PROBABILITIES OF 65-YEAR-OLDS IN 2014

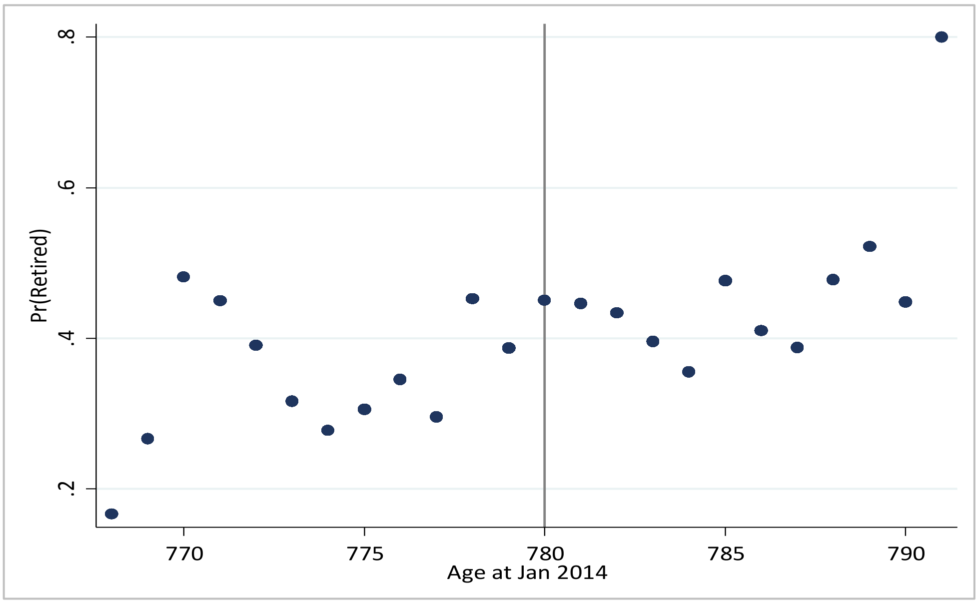

Source: CSO Quarterly National Household Survey.

Robustness, validity and sensitivity analysis

Identification in the regression discontinuity design is based on the assumption that individuals just to the left of the threshold possess similar observable and unobservable characteristics to those just to the right of the threshold, with the only difference between the two groups being that one barely qualified for treatment and one just missed out. While we cannot test whether the two groups are similar in their unobservable characteristics, we can test for comparability in observed characteristics. As mentioned above, of particular importance in this study is the age of the individual (in months) at the time of survey in 2014. Retirement is a function of age and if a systematic difference existed between the 780 group and the 779 group in terms of their age (in months) at the time of survey, then this could bias the results. To see this, consider an example. The data captures outcomes of individuals surveyed in Quarters 2 and 3 of 2014. If the 779 group were all surveyed in September 2014 and the 780 group were all surveyed in April 2014, that would mean that the 779 group at the time of survey, when the retirement outcome is captured, would be six months older than the 780 group. Therefore, this could bias our results, as our estimate would reflect the higher probability of retirement for the 779 group and the lower probability of the 780 group which is due to differences in age, as opposed to the causal effect of treatment. Table 6 confirms that this is not an issue with our data. There is no systematic difference between the ages of the two groups at the time of survey.

We also carry out a sensitivity analysis to see if the results change when a person's age (in months) at the time of survey is included as an additional explanatory variable in Equation 3. The results are not sensitive to the inclusion of this variable, as shown in the second row of results in Appendix Table A1.

TABLE 6 AGE AT TIME OF SURVEY (IN MONTHS)

Age at Jan 2014	Average Age at Time of Survey	Obs
780	785.5088	57
779	784.6207	87

Source: CSO Quarterly National Household Survey.

Additional covariates are examined to investigate whether differences exist between the 780 and 779 groups, including; gender, highest educational attainment, the probability of being married and the probability of being widowed. The average scores for both groups on each of these characteristics are shown in Table 7 and, as we can see, both groups are comparable. These covariates are then added into the parametric specification, along with age at time of survey, and the results are presented in the third row of Table A1. The estimates are not sensitive to the inclusion of the additional covariates, which indicates that the estimate is not being influenced by systematic differences in characteristics between the groups to the right and the left of the threshold.

TABLE 7 COMPARISON OF CHARACTERISTICS

	Mean	
Characteristic	**779 group**	**780 group**
Highest educational attainment (ISCED 11)	2.7	2.2
Probability of being male	0.49	0.53
Probability of being married	0.72	0.74
Probability of being widowed	0.11	0.09

Source: CSO Quarterly National Household Survey.

The graphical plots of the local average retirement rates (Figures 1 and 4) do not provide clear evidence of a discontinuity at the treatment threshold (780 months). Moreover, any significant discontinuity at the threshold for the low order polynomial specifications vanishes when we introduce a more flexible functional form, as in Figure 3. However, a notable feature that emerges from the graphs and the table of local averages (both Tables 1 and 3), is an apparent jump in the retirement probabilities at the 778 group. In both the baseline and age-restricted models, moving from the 777 to the 778 group sees a 16 percentage point increase in the retirement rate, from 30 per cent to 46 per cent. While the relatively low sample sizes can generate noisy estimates, the increase in retirement rates at 778 appears large and is statistically significant.

It is common practice in RD designs to carry out placebo tests which test for discontinuities away from the treatment threshold. If significant discontinuities are found at placebo points without any theoretical justification, this calls into question the reliability of the results at the threshold. We carry out a placebo test by designating the 778 group as a false cut-off and testing for discontinuities. The apparent jump at the 778 group is also of interest as it raises questions as to whether there was treatment contamination for those who just missed out on qualifying for the pension at age 65, i.e., did some individuals from the 778 and 779 group still manage to avail of the state pension at age 65? If so, the regression discontinuity design would be invalidated.[15] The results for each of the polynomial specifications are shown in the fourth row of results in Table A1. The estimates of the discontinuities at this point are larger in magnitude and show greater statistical significance at higher order polynomials and are statistically significant in the non-parametric specification. Therefore, the results using the 778 month cut-off are more consistent with a causal effect of the policy compared to the actual 780 month cut-off. While this is potentially attributable to noisy estimates as a result of relatively small sample sizes, given the estimator used in the analysis, the result also raises questions as to whether treatment spill-over occurred for individuals just to the left of the threshold.

Further investigation of this apparent jump in the retirement outcomes of the 778 group suggests that it may relate to individuals' labour outcomes changing from 'engaged in home duties' to 'retired'. Compared to the 777 group, the 778 group has approximately 15 per cent less people engaged in home duties and 15 per cent more who are retired. The questions relating to potential treatment spill-over remain if people engaged in home duties had the required insurance contributions and were granted the transition state pension despite barely missing out on the age threshold. However, the sample sizes are small, especially when we condition on individual categories, such as people engaged in home duties, making it difficult to draw concrete conclusions. Again, further analysis of this issue would require a richer dataset, ideally linking Revenue data to the Department of Social Protection data.

Employment

We carry out the same type of descriptive analysis as in Figure 1, but instead of using retirement as the outcome variable, we use employment. A dummy variable is created which indicates whether a person is 'working for payment or

15 While we have no direct evidence that this occurred, the possibility of treatment spill-over should always be considered in studies of this nature.

profit' during the reference month. We calculate averages of this outcome for each age category and Figure 5 shows the average employment in each of the monthly forcing variable bins. There is no clear evidence of a discontinuity at the threshold. As with the retirement outcome, when we run the RD regression (Equation 3), the discontinuity in the employment outcome is not robust to a flexible functional form, nor is it statistically significant when estimated using local linear regression (see Appendix Table A2).

FIGURE 5 AVERAGE EMPLOYMENT

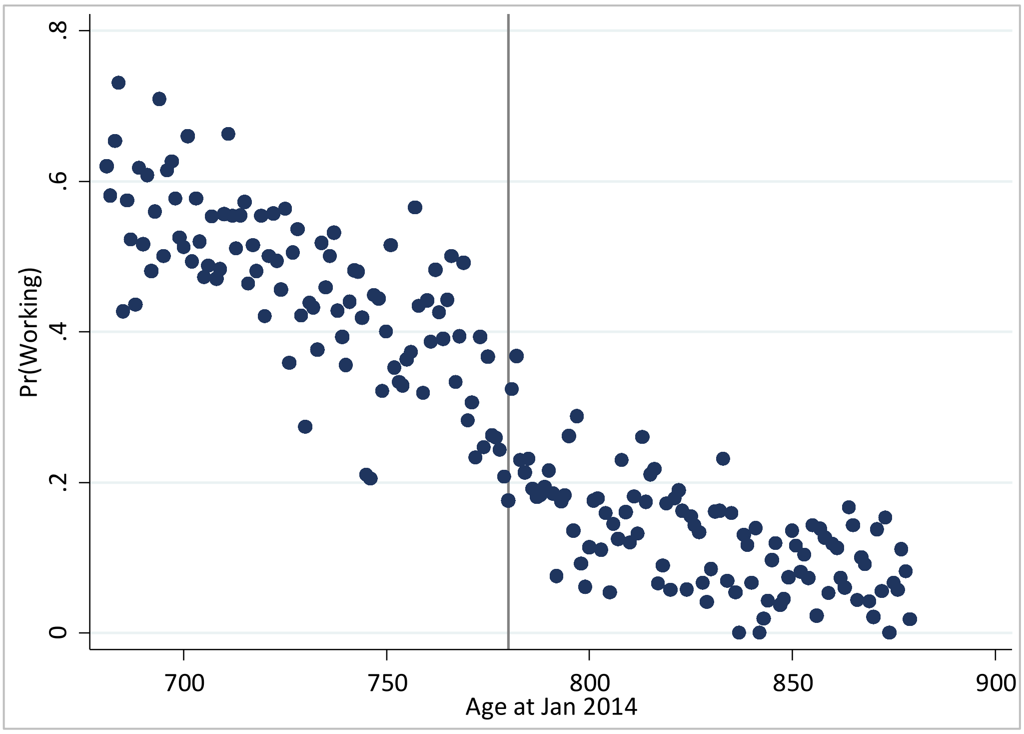

Source: CSO Quarterly National Household Survey.

Unemployment

We carry out the same descriptive analysis using unemployment as our outcome variable. A dummy is created which indicates whether a person's labour status is 'unemployed, having lost or given up previous job'. Figure 6 shows the average unemployment in each of the monthly forcing variable bins. Again, there is no clear evidence of a discontinuity. When we run the RD regression (Equation 3), the estimates are not statistically significant when flexible function forms are used in the parametric estimation, nor are they significant when estimated using local linear regressions (see Appendix Table A3).

FIGURE 6 AVERAGE UNEMPLOYMENT

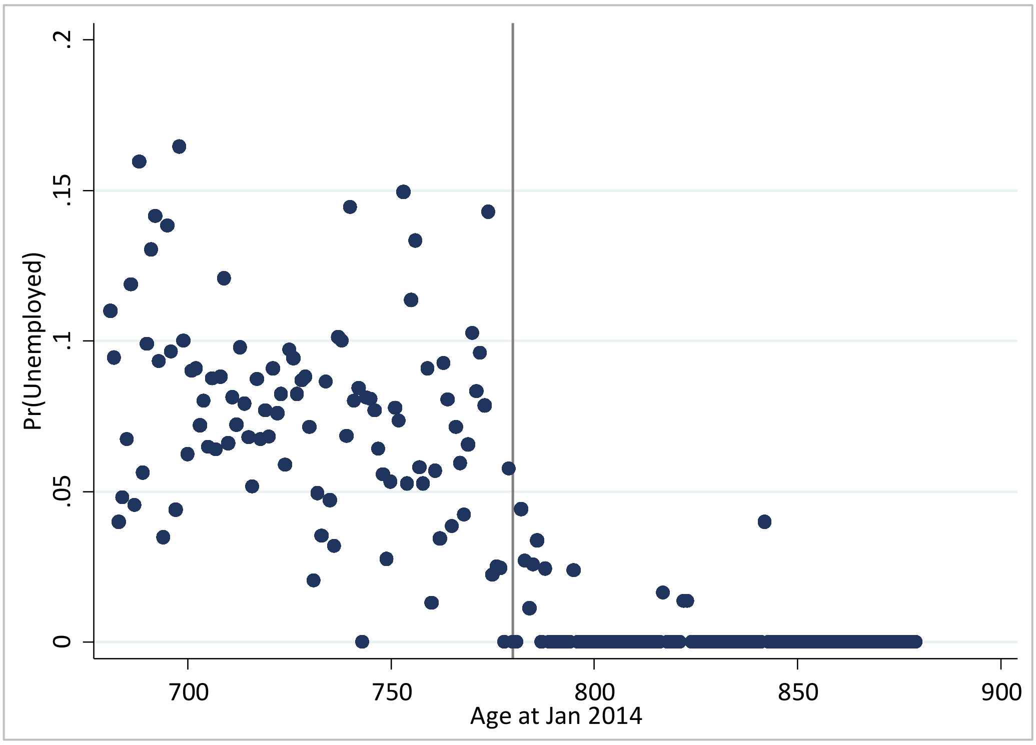

Source: CSO Quarterly National Household Survey.

CONCLUSION

In 2014, the qualification age for the transition state pension in Ireland increased from 65 years to 66 years. A sharp qualification threshold was implemented, such that individuals born before January 1949 could still qualify for the transition pension at age 65, whereas individuals born on or after January 1949 had to wait until age 66. By exploiting this sharp threshold using a regression discontinuity design, we were able to estimate, using data from the QNHS, the causal effect of the policy change on the retirement rate of 65-year-olds in 2014. Our analysis does not provide clear evidence of a causal effect of the policy on retirement rates. There are several potential explanations for this. To qualify for the transition state pension, an individual must have the required social insurance contributions. Not all 65-year-olds will meet this requirement and therefore the 'bite' of the policy may be limited as not all 65-year-olds are impacted by the change. Therefore, this may limit the effect of the policy change on the overall retirement rates of 65-year-olds. For example, if only a small percentage of the 65-year-olds in our sample were affected by the change, it is possible that the retirement rates of this subsample could be impacted and yet this would not show up as a strong impact on the overall retirement rate.

In addition, people's contracts may specify a retirement age of 65, or even where none is specified, there may be an expectation that people will retire at this age. As such, these individuals may still have to retire despite not qualifying for the transition pension. Moreover, the age at which an individual's occupational pension begins may remain at 65. Therefore, the one-year income gap created by not qualifying for the state pension may not be too severe for some individuals. It has also been the case that some 65-year-olds who did not qualify for the transition state pension were receiving Jobseeker's Benefit as a temporary payment until they reached the age of 66. The Department of Social Protection were aware that this was a temporary stop-gap measure to bridge people's retirement income with little expectation that these individuals would find work. In this scenario, whereby the transition state pension is unavailable, but Jobseeker's Benefit becomes a type of de facto pension payment which takes the place of the transition state pension, it is perhaps not surprising that the retirement rates are unaffected. This relates to how the outcome variable, i.e. the retirement rate, is constructed. This is based on an individual's self-reported main labour status, which could give rise to a number of complications when evaluating the causal effect of the policy change. For one, it is unclear how individuals who could not receive the state pension at age 65, due to being born after 1 January 1949, but received jobseeker's benefit as a type of de facto pension are categorised. Some of these individuals may report themselves as being unemployed, but others may report themselves as being in retirement. More detailed data would allow further investigation of this issue.

Finally, in studies of this nature, the possibility of treatment spill-over should be considered. The retirement rates of those who barely missed out on qualifying for the pension at age 65, namely the 778 and 779 group, are in line with the retirement rates of the people who qualified. If some individuals from these groups still managed to avail of the state pension at age 65, this could help explain the lack of any clear treatment effect. However, we have no direct evidence that this occurred and our data does not allow for further investigation of this issue.

We conclude with suggestions surrounding future work and improved data availability. Our analysis has focused on the retirement rate of 65-year-olds in 2014. An avenue for future research would be to condition the analysis on individuals who had the required social insurance contributions, thereby ensuring that the policy change affected all individuals being studied. This would overcome concerns surrounding the limited bite of the policy among the full sample of 65-year-olds. However, this would require a larger, more detailed dataset, which would provide data on an individual's employment and social insurance contribution history and ensure that enough observations existed to produce meaningful estimates. A linked administrative dataset between Revenue and the Department of Social Protection may be useful in this regard. In addition, while we use month of birth data in this analysis, day of birth data would be more desirable, especially in a dataset with larger sample sizes.

REFERENCES

Calonico, S., M. Cattaneo and R. Titiunik (2014). 'Robust Data Driven Inference in the Regression Discontinuity Design', *Stata Journal*, 14 (4): 909-946.

de Grip, A., D. Fourage and R. Montizaan (2013). 'How Sensitive are Individual Retirement Expectations to Raising the Retirement Age?', *De Economist*, 161 (3): 225-251.

Gelman, A. and G. Imbens (2014). 'Why High-order Polynomials Should not be Used in Regression Discontinuity Designs', NBER Working Paper No. 20405.

Imbens, G. and K. Kalyanaraman (2012). 'Optimal Bandwidth Choice for the Regression Discontinuity Estimator', *Review of Economic Studies*, 79 (3): 933-959.

Lee, D. and D. Card (2008). 'Regression Discontinuity Inference with Specification Error', *Journal of Econometrics*, 142 (2): 655-674.

Lee, D. and T. Lemieux (2010). 'Regression Discontinuity Designs in Economics', *Journal of Economic Literature*, 48 (2): 281-355.

Puur, A., L. Leppik and M. Klesment (2015). 'Changes in Pension Take-Up and Retirement in the Context of Increasing the Pension Age: The Case of Estonia in the 2000s', *Post-Communist Economies*, 27 (4): 497-516.

Sanchez-Martin, A.R., J. Ignacio Garcia-Perez and S. Jimenez-Martin (2014). 'Delaying the Normal and Early Retirement Ages in Spain: Behavioural and Welfare Consequences for Employed and Unemployed Workers', *De Economist,* 162 (4): 341-375.

Skovron, C. and R. Titiunik (2015). 'A Practical Guide to Regression Discontinuity Designs in Political Science', University of Michigan Working Paper.

Staubli, S. and J. Zweimuller (2013). 'Does Raising the Early Retirement Age Increase Employment of Older Workers?', *Journal of Public Economics*, 108: 17-32.

Vestad, O.L. (2013). 'Labour Supply Effects of Early Retirement Provision', *Labour Economics,* 25: 98-109.

TABLE A1 THE EFFECT OF THE CHANGE IN THE STATE PENSION AGE ON THE PROBABILITY OF RETIREMENT

	Order of Polynomial						
VARIABLES	1st	2nd	3rd	4th	5th	6th	Local Linear Regression
T	0.175***	0.167***	0.105***	0.070*	-0.0016	-0.0207	0.057
	(0.0154)	(0.0230)	(0.0305)	(0.038)	(0.0456)	(0.0536)	(0.0379)
Additional controls	No	No	No	No	No	No	No
Observations	14,911	14,911	14,911	14,911	14,911	14,911	3,217
T	0.175***	0.167***	0.105***	0.071*	-0.00164	-0.0207	0.057
	(0.0154)	(0.0230)	(0.0305)	(0.038)	(0.0456)	(0.0536)	(0.0379)
Additional controls	Age at Interview	Age at Interview	Age at Interview	Age at Interview	Age at Interview	Age at Interview	Age at Interview
Observations	14,911	14,911	14,911	14,911	14,911	14,911	3,217
T	0.181***	0.168***	0.096***	0.067*	-0.00852	-0.0051	0.057
	(0.0159)	(0.0238)	(0.0313)	(0.039)	(0.0461)	(0.0544)	(0.0379)
Additional controls	All	All	All	All	All	All	All
Observations	14,696	14,696	14,696	14,696	14,696	14,696	3,217
Placebo Test							
T (778 months)	0.179***	0.168***	0.111***	0.0816**	0.0026	-0.0039	0.117***
	(0.0153)	(0.0229)	(0.0303)	(0.0378)	(0.0505)	(0.0531)	(0.0375)
Additional controls	No	No	No	No	No	No	No
Observations	14,911	14,911	14,911	14,911	14,911	14,911	3,063
Individuals with employment experience							
T	0.189***	0.176***	0.102***	0.054	-0.023	-0.046	0.025
	(0.0164)	(0.0245)	(0.0325)	(0.040)	(0.048)	(0.057)	(0.044)
Additional controls	No						
Observations	13,774	13,774	13,774	13,774	13,774	13,774	2,514

Source: CSO Quarterly National Household Survey.
Note: The first three rows of results show the estimates from the baseline model with and without covariates. The fourth row estimates the placebo model and the fifth row conditions only on individuals with previous employment experience.

TABLE A2 THE EFFECT OF THE CHANGE IN THE STATE PENSION AGE ON THE PROBABILITY OF EMPLOYMENT

	Order of Polynomial				
VARIABLES	1st	2nd	3rd	4th	Local Linear Regression
T	-0.109***	-0.076***	-0.050	0.013	0.064
	(0.0153)	(0.0227)	(0.038)	(0.0387)	(0.0418)
Additional controls	No	No	No	No	No
Observations	14,911	14,911	14,911	14,911	2,208
Individuals with employment experience					
T	-0.119***	-0.088***	-0.065*	-0.004	0.048
	(0.0163)	(0.0242)	(0.0323)	(0.0412)	(0.044)
Additional controls	No	No	No	No	No
Observations	13,774	13,774	13,774	13,774	2,021

Source: CSO Quarterly National Household Survey.
Note: The first row of results shows the estimates from the baseline model. The second row conditions only on individuals with previous employment experience.

TABLE A3 THE EFFECT OF THE CHANGE IN THE STATE PENSION AGE ON THE PROBABILITY OF UNEMPLOYMENT

	Order of Polynomials				
VARIABLES	1st	2nd	3rd	4th	Local Linear Regression
T	-0.028***	-0.024***	-0.009	-0.003	-0.018
	(0.008)	(0.008)	(0.0062)	(0.0023)	(0.0127)
Additional controls	No	No	No	No	No
Observations	14,911	14,911	14,911	14,911	3,061
Individuals with employment experience					
T	-0.033***	-0.028***	-0.011*	-0.004	-0.025*
	(0.0097)	(0.0095)	(0.0067)	(0.0033)	(0.0138)
Additional controls					
Observations	13,774	13,774	13,774	13,774	2,942

Source: CSO Quarterly National Household Survey.
Note: The first row of results shows the estimates from the baseline model. The second row conditions only on individuals with previous employment experience.