IFS

Institute for Fiscal Studies

Courtney Coile
Jonathan Cribb
Carl Emmerson
Laurence O'Brien

Working paper

25/01

Why does raising the early retirement age affect employment?

Economic and Social Research Council

Why Does Raising the Early Retirement Age Affect Employment?

Courtney Coile[1], Jonathan Cribb[2], Carl Emmerson[2] and Laurence O'Brien[3]

[1] Wellesley College and NBER

[2] Institute for Fiscal Studies

[3] Institute for Fiscal Studies and University College London

Abstract:

Why do workers delay retirement when their early retirement age (ERA) rises? Understanding this is essential for predicting the effects of ERA increases in different contexts and assessing distributional impacts of such reforms. We study ten years of reforms in the United Kingdom which increased the ERA for women from 60 to 66. The UK is an attractive setting for this question because there is no financial incentive to retire at the ERA. Credit constraints are found to be important: credit constrained groups experience around a 13-percentage-point larger increase in employment rates as a result of the reform than non-constrained groups. In contrast, groups with greater loss of pension wealth do not see larger increases in employment rates and employers do not delay dismissing employees as a result of the reform. There are smaller, but still large and statistically significant, increases in employment rates of around 11 percentage points for wealthier groups for whom other explanations are shown to be unimportant. This points towards behavioural reasons such as signalling, framing, or social norms also being important.

Keywords: Early retirement age; retirement; credit constraints

JEL codes: H55, J21, J26

---

* We are grateful for funding from the National Bureau of Economic Research's Working Longer program (funded by the Alfred P. Sloan foundation under the grant "Economics of Working Longer and an Aging Workforce", grant number G-2019-12400), the Economic and Social Research Council (ESRC) Joint Programming Initiative More Years Better Lives (Project "Unequal ageing: life-expectancy, care needs and reforms to the welfare state", grant reference ES/W001594/1) and the ESRC Centre for the Microeconomic Analysis of Public Policy (ES/T014334/1). We thank Stuart Adam, James Banks, Richard Blundell, Antoine Bozio, Eric French, Lisa Laun, Andreas Haller, Peter Haan, Rachel Griffith, participants in the JPI project "PENSINEQ", and seminar participants at the IFS, Paris School of Economics, Leiden University, Royal Economic Society Conference, and Netspar International Pension Workshop. Contact details: Coile: ccoile@wellesley.edu; Cribb: jonathan_c@ifs.org.uk ; Emmerson: carl_e@ifs.org.uk; O'Brien (corresponding author): Laurence_o@ifs.org.uk.

I. Introduction

In response to the public finance pressures from ageing populations, many governments have increased the minimum age at which someone can claim a public pension (the early retirement age, ERA) and/or the age at which individuals become eligible for a full unreduced public pension (the normal retirement age, NRA). These policies strengthen the public finances both by reducing pension payments and by delaying retirement, thereby boosting tax revenue. A growing scientific literature has studied these reforms, and typically finds that such increases in retirement ages lead to substantial increases in employment at older ages.

Importantly, however, the mechanisms underlying these employment responses are still not well understood. Atalay and Barrett (2015) argue that, as a higher pension age reduces public pension wealth, it leads to later retirement through an income effect. Rabaté et al. (2024) instead find evidence that credit constraints and employer responses play an important role; Manoli and Weber (2016), on the other hand, argue that credit constraints are not an important driver. Behavioural effects have also been shown to have an effect: both Seibold (2021) and Gruber et al. (2022) find significantly larger changes in retirement behaviour at statutory retirement ages than at points with similar financial incentives to retire, but which are not labelled as a retirement age, pointing to strong impacts of norms around retirement. Contributing to the lack of consensus is the fact that no previous study has examined the various potential mechanisms together in a single setting, making it difficult to determine their relative role in the employment response.

Better understanding these mechanisms would shed light on how people make retirement decisions and is essential for predicting the effects of future pension eligibility age increases in different contexts, where wealth effects or credit constraints may loom larger or smaller due to differences in financial circumstances or interactions with other policies. They are also key to understanding the distributional consequences of increases in pension eligibility ages, which is particularly important given these reforms are a common, but sometimes controversial, way of addressing public finance challenges caused by an ageing population.

In this paper, we measure the importance of several different mechanisms driving employment responses to increases in ERAs in a single setting. In particular, we use large-scale survey datasets from the UK to analyse the effect of a large increase in the ERA for women from 60 to 66 that occurred over just ten years between 2010 and 2020. We estimate employment responses to the reform during women's early 60s using difference-in-differences models that effectively compare the employment outcomes of women born a few months apart who face different ERAs, and identify the mechanisms underlying the overall employment response by comparing precisely estimated effects for different groups of women.

The UK setting is particularly advantageous for identifying the mechanisms underlying the employment responses to increases in the ERA. First, it provides us with a large increase in the ERA over a relatively small period of time. Second, it represents a particularly clean setting to tease apart the importance of credit constraints, wealth effects, behavioural reasons, and employer responses. This is because, unlike in many other countries, there is no financial incentive to retire at the ERA.

We show empirically that employer responses (such as those identified in Rabaté et al. (2024) in the Netherlands) are not an important driver of employment responses to increases in the ERA in the UK, as there is no spike in dismissals or redundancies at the ERA. This is consistent with strong employment protections that prevent employers from terminating employment once their employees reach the ERA.

By comparing how the employment response varies for women who face different degrees of credit constraints in the years approaching the ERA, and women who experienced very different wealth losses due to the reform, we can disentangle the extent to which the overall employment effect we observe is driven by credit constraints and wealth effects. This then leaves behavioural effects such as signalling, framing, or social norms to retire at a focal age in the public pension system as the most likely explanation for any residual effect.

Our results focus on women who were in paid work at age 59 (i.e. prior to the pre-reform ERA of 60). This is because – consistent with previous work (Staubli and Zweimuller, 2013; Amin-Smith and Crawford, 2018) – we find that increases in the ERA only lead to those in paid work delaying their retirement, rather than women re-entering employment or "unretiring". On average amongst those in paid work at age 59, we estimate that being under the ERA increases the probability of a woman being in paid work in her early 60s by 15 percentage points.

We find considerably larger employment impacts for groups of women who have very low levels of financial wealth, implying an important role for credit constraints. Women in the bottom quartile of the liquid financial wealth distribution experience an employment increase of 24 percentage points as a result of increasing the ERA. This compares to an employment response of 11 percentage points for women in the top quartile of the liquid financial wealth distribution.

In contrast to the important role played by credit constraints, we find no evidence that birth cohorts with (much) larger public pension wealth losses as a result of the reform are more likely to delay their retirement than those with much lower wealth losses, leading us to conclude that wealth effects are not an important driver of the employment response.

Finally, our finding of a large and statistically significant positive employment response among groups with substantial financial wealth, and the lack of importance of wealth effects or employer responses leaves a residual employment response, which could be explained by behavioural factors, such as signalling, framing, and social norms to retire at the ERA.

The rest of this paper proceeds as follows. Section II sets out the potential mechanisms that can drive the employment response to the increase in the ERA and highlights where evidence for these mechanisms exists in previous studies. Section III describes the institutional setting and the policy reforms we exploit, while Section IV introduces the data used for our analysis. Section V describes our empirical strategy. Section VI contains our results and Section VII concludes.

II. Potential mechanisms

There has been a burgeoning empirical literature examining how increases in retirement ages affect employment and retirement behaviour. Large employment increases have consistently been found in various different countries in response to increases in both ERAs and NRAs, including in Australia (Atalay and Barrett, 2015; Morris, 2022); Austria (Staubli and Zweimüller, 2013; Manoli and Weber, 2016), Denmark (Sæverud, 2024); Estonia (Soosaar et al., 2020), France (Rabaté and Rochut, 2019), Germany (Geyer and Welteke, 2019; Seibold, 2021), the Netherlands (Rabaté et al., 2024), Norway (Johnsen et al., 2021), the US (Mastrobuoni, 2009; Deshpande et al., 2024), as well as in the UK (see Cribb et al. (2016) who analyse the initial stage of the reform we consider in this paper – the increase in the ERA from 60 to 62 that occurred between 2010 and 2014).

The mechanisms driving these employment responses are, however, still unclear, and may differ across countries depending on the institutional details and the circumstances of those approaching typical retirement ages. There are five main reasons why an increase in a retirement age might increase employment, and the importance of these reasons might differ for ERAs and NRAs. In this section, we briefly set these out, highlight where evidence for these mechanisms exists, and how our work builds upon the existing literature.

The first potential mechanism is that statutory retirement ages can affect the marginal financial incentive to stay in paid work. In many countries, public pension income is subject to an earnings-test, or is means-tested in some other way, after the ERA and/or the NRA. When combined with a less than actuarially fair increase for delayed claiming, raising the retirement age can reduce the financial incentive to stay in paid work at particular ages. This is particularly important in some European countries; for example, Staubli and Zweimüller (2013) analysed the Austrian system where pension benefits were withdrawn entirely if the claimant earned above €380 per month. Börsch-Supan and Coile (2021) draw together evidence from a range of high-income countries and find that financial incentives to work are important in driving employment decisions, particularly in people's early 60s (the age group we study). However, in the UK, state pension income is not means-tested or earnings-tested, and there is no change in the financial incentive to retire at the ERA. This means that the implicit tax on delaying retirement as a result of the pension system is close to zero, in contrast with many other high-income countries (see Banks et al., 2023; Börsch-Supan and Coile, 2023).

A second mechanism through which statutory retirement ages can affect employment is through employers' responses in terms of labour demand for older workers. Rabaté et al. (2024) highlight the importance of automatic job termination for explaining employment responses to the NRA in the Netherlands, and Deshpande et al. (2024) find suggestive evidence of a similar mechanism in the US. In the UK, in contrast, employers cannot legally force someone to retire except in particularly rare circumstances (such as being a commercial pilot where a retirement age of 65 exists). In addition, under the Equality Act 2010, employers cannot legally discriminate against people based on their age. Therefore, while employees can be dismissed without a reason in the first two years of their job in the UK, it is illegal to do so because of their age. The period we analyse in the UK, the 2010s, is also a period with generally low levels of unemployment, including among women in their sixties. Consistent with this, we find no evidence of increased redundancies or other dismissals at the ERA.

The UK therefore provides a clean setting to examine the relative importance of the three other mechanisms: credit constraints, wealth effects, and behavioural effects or norms that could induce people to retire at the ERA.

Credit constraints may play an important role in leading people to delay retirement because the increased pension eligibility age significantly reduces incomes (Cribb and Emmerson, 2019) and, in absence of savings to draw down upon, individuals could respond by delaying their retirement to help fund their expenditure. Rabaté et al. (2024) find that employment responses to the NRA are around 50% larger for those in the lowest quartile of the wealth distribution than for those in the highest quartile in the Netherlands. It is likely, though, that credit constraints are more important for explaining retirement at ERAs than NRAs, as people have the option of claiming public pension income before the NRA (as long as they have reached the ERA). Consistent with this, in our setting the bottom quartile of the liquid financial wealth distribution has an employment response to increases in the ERA twice the size of that of the top quartile. This implies that credit constraints play an especially important role in driving higher employment in response to increases in the ERA.

Wealth effects are an alternative explanation. In the UK the increased public pension claiming age leads to lower public pension wealth, as the ERA is also the same as the age you can claim an unreduced pension (NRA). With leisure a normal good, this leads to later retirement. Atalay and Barrett (2015) analyse an increase in the ERA in Australia and show that the employment responses are larger for cohorts of women who lost more public pension entitlement as a result of the reform, which they interpret as evidence of wealth effects. However, Morris (2022) questions Atalay and Barrett (2015)'s use of men as a control group for women, which could potentially bias their results. We use a similar method to Atalay and Barrett (2015) (but without using different generations of men to control for the trend in employment for women) and find no evidence of stronger employment responses for those women who lost more wealth as a result of the reform, indicating that wealth effects are not a driver of the overall employment response we find.

The final set of mechanisms through which the ERA leads to individuals retiring at the ERA are behavioural ones. There is a range of potential behavioural reasons, such as the ERA being a signal or social norm about an appropriate age to retire (Gruber et al., 2022), or there being peer effects from others retiring at the ERA. These types of channels are emphasised by Seibold (2021). The fact that we find no evidence of wealth effects or employer responses, and yet still find significant positive employment responses for women who are not credit constrained, indicates that this mechanism may be important in driving some, but not all, of the increase in employment in response to higher pension eligibility ages.

III. Institutional background

The state pension age is the earliest age at which an individual in the UK can claim a state pension. It is therefore similar to the ERA in other countries such as the United States. Unlike in the US, there is no other focal age, such as a NRA, in the UK state pension system. While it is possible to delay receiving a state pension, and to receive an actuarial adjustment for doing so,[1] deferral is rare with around 95% of individuals receiving the state pension within two months of reaching their ERA (Crawford and Tetlow, 2010).

Prior to April 2016, the UK state pension consisted of two parts. The first was the basic state pension, based on the number of years of contributions made, but not the level of these contributions. The second part was the second-tier pension, which is positively related to earnings across the whole of working life. However, the majority of employees opted out of building up entitlement to this second-tier pension in return for a reduction in payroll taxes and built up a private pension instead.[2] The basic state pension was worth approximately £6,000 to £6,500 per year ($7,800 to $8,300) between 2010–11 and 2015–16 for someone with 30 qualifying years (all monetary figures are expressed in 2020–21 prices, and 30 years being the amount needed to quality for a full basic state pension during this period).

Individuals reaching their ERA after 5 April 2016 can instead qualify for the new state pension. This is a flat-rate state pension, worth around £9,600 ($12,200) per year in 2022 (around 30% of median full-time employee earnings) for someone with 35 years of contributions (which is now the number of years required to qualify for a full award).[3] The UK no longer has an earnings-related state pension.

Although the state pension in the UK is less valuable than in many other European countries,[4] it still makes up a large proportion of income for many pensioners. The state pension constitutes half of income for middle-income pensioners aged 66-74 and 75% of income for the poorest fifth. Private pensions are also important for middle and high income

[1] For those reaching ERA before 6 April 2016, deferring state pension receipt by one year led to an increase in its value by around 10.4%. For those reaching ERA from 6 April 2016, this adjustment was reduced to 5.8% for each full year deferred.

[2] According to Pensions Commission (2004), in the early 2000s around 60% of the working-age population were building entitlement to a salary related pension scheme: for around 25% this was the state second-tier pension while for around 35% this was contributions into a private pension arrangement.

[3] In early 2022, the average receipt of a new state pension among women was £165.05 per week, compared to the maximum of £179.60 per week at that point, implying they received on average 92% of the maximum. See https://www.gov.uk/government/statistics/dwp-benefits-statistics-august-2022

[4] See https://appsso.eurostat.ec.europa.eu/nui/show.do?dataset=ilc_pnp3&lang=en

pensioners making up 25% of income for the middle fifth, and half of income for the highest-income fifth (Cribb, 2023).[5]

In 1948, the ERA was lowered to 60 for women and remained at this age until April 2010. Then between 6th April 2010 and 6th October 2020 the female ERA increased gradually, but swiftly, to 66. This was done by two sets of legislation. The 1995 Pensions Act increased the ERA for women from 60 to 65 from 2010 to 2020 (in order to equalise the ERA with that of men). The 2011 Pensions Act subsequently accelerated the rise from 63 to 65 between 2016 and 2018, and increased it (and the ERA for men) to 66 by the end of 2020. Figure 1 shows the resulting ERA by date of birth for women born in the early 1950s. During the increase, each month the date at which a woman could claim their state pension increased by either two or four months, leading to the "sawtooth" pattern in Figure 1.[6]

Importantly, the age at which individuals can withdraw funds from private pension schemes, the normal minimum pension age, was not affected by the increase in the ERA, and remained at 55 during our analysis period. Similarly, Normal Pension Ages for public- and private-sector defined benefit occupational pension schemes were also not affected by the rise in the ERA for the birth cohorts we analyse. In addition, average private pension wealth varies little between women in the different (but close) birth cohorts we compare in our analysis. All our analysis therefore isolates the effect of a change in the age at which women can first claim the UK state pension, with private pension claiming ages unchanged.

[5] From age 55, individuals can withdraw 25% of defined contribution pension wealth in a tax-free cash lump sum. Prior to 2015, the remaining 75% had to be annuitized in the majority of cases; however, compulsory annuitisation was abolished in 2015, allowing people more flexibility over how to draw their private pension.

[6] Unlike in other settings such as the Netherlands (see Rabaté et al. (2024)) or in Germany (see Geyer et al. (2020)), there is no large discontinuity in ERA by date of birth for women in the UK. As a result, we use a difference-in-differences methodology for our analysis as opposed to a regression discontinuity design.

Figure 1 UK early retirement age for women by date of birth

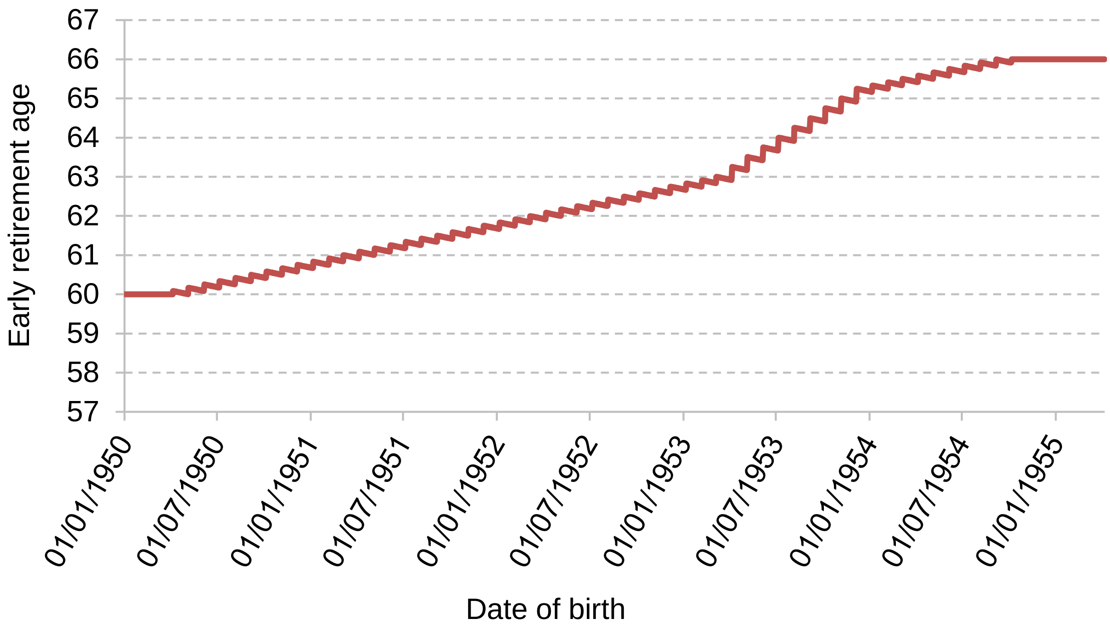

Source: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/310231/ERA-timetable.pdf

Notes: This figure shows the UK ERA for women by date of birth, for those born between 1 January 1950 and 6 April 1955. The reason that the ERA increases in a "sawtooth" pattern, rather than a smooth line or a "step" pattern, is that women born in a given month are allocated a single "state pension date" at which they are eligible for a state pension. Therefore, those born later in the month can have a slightly lower ERA than those born earlier in the month.

There is no increase in the marginal financial incentive to retire at the ERA, in sharp contrast to in many other countries. One of the main reasons for this is that there is no earnings test on state pension income. In addition, the tax and benefit system changes upon reaching the ERA in a way that actually increases the incentive to remain in paid employment. First, employees above the ERA are no longer liable for employee National Insurance contributions (a payroll tax), increasing their take-home pay for a given gross wage. Second, the state benefit/transfer system is means-tested for those below ERA, which provides an incentive not to engage in paid work. This incentive disappears for those above ERA since the state pension is not means-tested. Appendix Table 1 sets out this calculation in detail for an example individual. Furthermore, Banks et al. (2023) show that the implicit tax on continuing to work at older ages for women in their early 60s is close to zero in the UK.

This is in contrast to many other high-income European countries such as France, Germany and the Netherlands – where similar reforms have been studied – where it is considerably higher (Börsch-Supan and Coile, 2021; 2023).

In the early part of our sample period, employers were legally allowed to force workers to retire at the ‘Default Retirement Age’; however, this was set at age 65 and so significantly above the contemporaneous female ERA. In April 2011, the Default Retirement Age was abolished in the UK, meaning that employers can only force workers to retire at a certain age if they can provide a good reason, such as if the job requires certain physical abilities or in rare cases where the job has an age limit set by law, such as commercial airline pilots. The vast majority of women in their 50s and 60s do not work in such jobs, meaning that if they wish to work longer they cannot legally be discriminated against. This implies that forced retirement is unlikely to be a large driver of the employment effects we find of reaching the ERA and indeed we find evidence that dismissals do not contribute to the employment response.

IV. Data

For most of our analysis, we combine information from two datasets: the Labour Force Survey (LFS) and the Local Labour Force Survey (LLFS). We also use the English Longitudinal Study of Ageing (ELSA), which contains detailed wealth variables, to test some additional hypotheses.

The LFS is a quarterly household survey that is representative of the UK population, and which aims to capture information on individuals’ labour market activities. We supplement our LFS dataset with Local Labour Force Survey data, which is also a UK-wide dataset using approximately the same questionnaire as the LFS, boosting the sample size for our analysis. As well as information on labour market activities, the surveys contain background information on individuals’ age, sex, region, marital status, education, broad housing tenure, and some self-reported measures of health. We can also observe date of birth in the data (month of birth in the LLFS), allowing us to calculate whether an individual is above or below the ERA on their interview date.

Throughout our analysis, we use the International Labour Organization definitions of economic activity. An individual is defined as employed if they do any paid work (either as an employee or self-employed) in the week of their interview, if they are temporarily away from work (for example due to holiday or illness, or – as we describe in Footnote 7 – if they are on furlough during the Covid-19 pandemic), or if they are on a government training scheme (but this last category is uncommon for women of the ages we are studying). We define full-time work as working 30 hours or more in a usual week.[7]

Our full sample covers the period 2009Q2 (one year before the ERA started to rise) to 2021Q2 (two quarters after the ERA had reached age 66) and contains women between the ages of 59 and 67 (inclusive). Table 1 shows the background characteristics of women in this sample. Over two thirds are in a couple, more commonly with an older partner. Less than one third of our sample has a partner in paid work. Over 80% of our sample are homeowners, the vast majority of whom have already paid off their mortgage. Around 30% of our sample have a higher education qualification, such as a degree, while just over a fifth have no formal educational qualifications at all. Finally, approximately half the women in our sample report having a health problem that has lasted, or is expected to last, for at least one year.

[7] The data we use cover – in part – the Covid-19 pandemic in which many people did not undertake their normal activity due to lockdowns or social distancing. In response to the pandemic, the UK government rapidly created a furlough scheme, by which the government would pay 80% of February 2020 salaries up to a cap of £2,500 per month for furloughed private-sector employees. At its peak in Spring 2020, 8.9 million employees (30%) were furloughed (see Cribb and Johnson 2022). However, as these people are counted as being temporarily away from their work, they are still counted as “employed” in our data. Our results are robust to excluding data after February 2020 and also to re-categorising furloughed people as not employed.

Table 1 Sample summary statistics of women aged 59-67; 2009 to 2021

Characteristic	All women	Those who were in paid work at age 59
Single	30%	30%
In a couple	70%	70%
Younger partner	18%	19%
Older partner	52%	50%
Partner in paid work	27%	33%
Not white	4%	3%
Homeowner	81%	86%
Renter	19%	14%
Higher education	30%	33%
Secondary education	47%	50%
No formal qualifications	23%	16%
Longstanding health problem	50%	43%
Number of observations	345,808	214,177

Notes*:* This table shows the summary statistics of our two analysis samples. The 'All women' sample includes women between the ages of 59 and 67 for the sample period 2009Q2 to 2021Q2. The 'In paid work at 59' sample restricts the main sample to 59- to 67-year-olds who were in paid work on their 59th birthday. Data come from the Labour Force Survey and the Local Labour Force Survey. 'Single' refers to people who are not currently in a couple, that is, it includes those who are divorced, separated or widowed.

Past studies on the effect of higher retirement ages on employment emphasise that the vast majority of the employment response comes from workers staying in their existing job for longer, rather than people joining the labour force having been out of work previously (Amin-Smith and Crawford, 2018; Staubli and Zweimüller, 2013). This implies that the pool of women who are realistically likely to respond to the ERA increase from 60 to 66 is women who are still in paid work at age 59. We restrict our main analysis sample to women who are in paid work in the month of their 59th birthday.[8]

[8] For those not in paid work we construct this by using information on the month and year they report last being in paid work. For those in paid work we assume that they were also in paid work in the month of their 59th birthday. Because unretirement is rare in the UK (Kanabar, 2015), this approximation should not introduce much measurement error. Appendix Figure 1 shows the share of those born in 1952 who we estimate were in

Our main sample of 59- to 67-year-old women who were in paid work at age 59 contains 214,177 observations – this large sample size allows for precision in our estimates and for us to test meaningfully for differences in the effects between different groups. Table 1 shows that those who were in paid work at age 59 are slightly better off socio-economically on a number of dimensions, including homeownership, education, and health, than the sample of all women of the same age.

To measure how credit constraints affect the employment response to the increase in the ERA, we need to identify credit constrained groups. Since the LFS and LLFS do not contain any information on assets (apart from a measure of housing tenure), we make use of ELSA (see Steptoe et al., 2013). This is a longitudinal, biennial household survey of people aged 50 and over in England, which contains detailed information on respondents' assets and background characteristics.[9] Using ELSA data on the same age groups of women as described in Table 1, for women in paid work at age 59, we have a sample of 5,698 observations of 2,048 individuals in waves 4 (2008-09) to 9 (2018-19) of ELSA.

We identify credit constrained groups using a measure of liquid financial wealth that includes the partner/spouse's financial wealth where appropriate, and includes current accounts, cash savings accounts, and investment accounts (such as equities and bonds held in accessible (non-pension) accounts). It is measured net of financial debts such as credit card and personal loans.[10] We split women into quartiles based on levels of this measure of wealth when in their mid 50s (before the old ERA and where the position in the liquid wealth distribution is plausibly unaffected, or at most only marginally affected by the reform of

paid work in the month of their 59th birthday in 2011, by their current age. Reassuringly, the estimated share in paid work at this point is essentially constant across ages.

[9] We do not use ELSA for our main analysis because it has a much lower sample size than the LFS and LLFS, leading to less precise estimates.

[10] A full description of the measurement of this financial wealth is provided in the ELSA documentation. The variable name is netfw_bu_s. See here for details: https://www.ucl.ac.uk/epidemiology-health-care/sites/epidemiology_health_care/files/financial_derived_variables_elsa_waves_1_-_7.pdf

interest), who will face differing degrees of credit constraints, and test whether we see significantly different employment responses.[11]

V. Empirical methodology

We use a difference-in-differences methodology to estimate the impact of increasing the female ERA from 60 to 66 on employment. Similar to Staubli and Zweimuller (2013), Cribb et al. (2016) and others, our 'treatment' variable is being under the ERA, which is a deterministic function of an individual's date of birth and the interview date. This treatment is administered to all women, but is administered to longer for women born more recently as the ERA increases. Equation (1) below sets out the difference-in-differences model we employ to estimate the employment effect of increasing the ERA.

$$(1) \quad y_{iact} = \alpha + \beta(underERA)_{iact} + \gamma_t + \lambda_c + \theta_a + X_{iact}\delta + \varepsilon_{iact}$$

We want to estimate, for an individual i, the effect of being under the ERA on an outcome y, such as employment. Therefore, our key coefficient of interest is β. Since $underERA$ is a deterministic interaction between their age (a) and interview date (t), we control flexibly for age and calendar time using age and time-period fixed effects, θ_a and γ_t, (both measured in years and quarters). This is therefore the traditional two-way fixed-effects (TWFE) difference-in-difference estimator.

It is also important to control for birth cohort (c), since women born more recently tend to have greater labour market attachment than those born earlier. To do this, we include birth

[11] To be precise, we aim to split women into quartiles based on levels of this measure of wealth when aged 55. However, we do not observe wealth of all women in our data at age 55 for three reasons. First, women are surveyed only once every two years in ELSA. Second, the wealth variable is missing for a small number of observations. Third, some women join the survey after age 55. As a result, we use the level of wealth at the closest age to 55 we observe in the data for each woman. Since we would expect wealth to vary systematically with age, we split women into quartiles conditional on the age at which wealth is observed.

cohort fixed effects for each financial year of birth (which in the UK starts each year on April 6), λ_c, as in Cribb, Emmerson and Tetlow (2016). We cannot control for birth cohort measured in years and quarters since this would be perfectly collinear with our age and time fixed effects. The inclusion of financial year of birth fixed effects allows us to take account of the different work histories of different cohorts of women approaching pensionable age. This does mean that any impact of the increase in the ERA that generates time-invariant changes in employment before the age of 60 (the pre-reform ERA) will be subsumed into these cohort effects. As a result, we estimate the employment impact of increasing the ERA for women aged at least 60, up to age 65 (as that is highest age that we observe someone to be under the ERA).[12]

We also control for a set of individual characteristics, X. These include relationship status, education level, ethnicity, housing tenure and geography, as well as (for those with a partner) their partner's age and partner's education.[13] We estimate the models using ordinary least squares and we cluster standard errors at the year-and-month-of-birth level to allow for serially correlated birth cohort shocks.

β is identified as the causal effect of being under the ERA under the 'common trends' assumption that the employment rates of different age groups would have changed similarly over time if the female ERA had remained at 60. As we include age, time, and birth cohort effects, our model is identified by comparing the behaviour of people born a few months

[12] The increase in the ERA that we study could potentially also change employment rates at younger ages. Lower lifetime wealth could lead to increases in employment rates during people's 50s. Artmann et al. (2023) document evidence of changes in women's labour supply during their 50s to a change in future pension wealth in Germany. As discussed, we focus on labour supply responses during women's early 60s and identify the importance of different mechanisms for explaining this response.

[13] Specifically, we include 4 dummies for relationship status, 5 dummies for education level, a dummy for whether the respondent is white or not, a dummy for whether the respondent rents their home or not, 17 geographical area dummies, a dummy for the dataset (LFS or LLFS), 5 dummies for the respondent's partner's education level, a dummy for whether the respondent's partner is under their ERA and partner's age (linear and quadratic), as well as the 49 dummies for quarter and 39 dummies for age in quarters mentioned earlier in the paragraph.

apart who face a different ERA; we assume that, conditional on these other factors, the only reason for employment rates to differ between the people facing different ERAs is the effect of the higher ERA.

Figure 2 presents evidence that the standard common trends assumption is plausible, and also offers an initial indication of the effect of increasing the ERA on employment. This figure shows female employment rates by age between 2004 and 2021, with the dashed lines indicating the periods when women of each age are going from over the ERA to under the ERA (e.g. it is dashed for 61-year-olds as the ERA rises from 61 to 62). Outside of these periods, the employment rates of different ages evolve in a broadly parallel manner. During the periods where the ERA is rising for a particular age, we see a marked increase in the employment rate for women of this age of approximately 10 percentage points.

We can also estimate the dynamic effects of the early retirement age on our outcome y using an "event-study" specification as set out in Equation (2) below.

(2) $$y_{iact} = \alpha + \sum_{j \neq -1} \beta_j \mathbb{1}(distERA_{iact} = j) + \gamma_t + \lambda_c + \theta_a + X_{iact}\delta + \varepsilon_{iact}$$

$distERA_{iact}$ denotes the number of quarters before individual i reaches their early retirement age (e.g. $distERA_{iact} = 0$ if individual i has reached their early retirement age in the past three months at interview date t).

Recent years have brought into question using TWFE as a way of implementing a staggered difference-in-differences methodology such as that set out here (see Baker et al., 2022). We therefore re-run key results using an alternative estimator – that proposed by Borusyak et al. (2024; henceforth "BJS") – and show the results in the Appendix. Full details on how we perform BJS estimation are available in Appendix Section B, but briefly this estimator can be represented by an imputation procedure. In the first step, the set of untreated observations is used to estimate the relationship between the outcome variable and age, time and a variety of control variables. In the second step, the parameters estimated in the first step are used to predict the counterfactual outcome for each treated observation. The treatment effect for each treated observation is then the difference between the actual outcome and the simulated counterfactual outcome, and the overall treatment effect is the average of all observation-level treatment effects. This estimation procedure does not use already-treated observations to calculate the counterfactual, and does not use negative weights when calculating the overall average treatment effect, thereby avoiding some of the potential issues that can arise with TWFE estimation (Baker et al., 2022).

The results obtained using the TWFE estimator and the BJS estimator are very similar for all outcomes and subgroups. Figure 2 provides some intuition for why the TWFE estimator is correct in our setting: there are similar rises in employment generated by each successive increase in the ERA, from the increase affecting 60-year olds to the increase affecting 65-year olds, as shown by the successive sets of dashed lines. As Baker et al. (2022) set out, TWFE is valid even in staggered difference-in-differences designs as long as there are constant treatment effects as the treatment is rolled out – as is suggested by Figure 2. Appendix Figure 2 shows a similar pattern for the set of women who were in paid work at age 59 (our main sample of interest) with slightly larger increases in employment amongst that group occurring as the ERA rises.

To disentangle the mechanisms through which the ERA affects employment, we use three further extensions. First, to identify the role played by credit constraints, we compare the employment effects for groups who have different levels of liquid financial wealth. We do this using the ELSA, running separate regressions for women in different quartiles of the liquid financial wealth distribution.

Figure 2 Employment rates of women over time, by single year of age

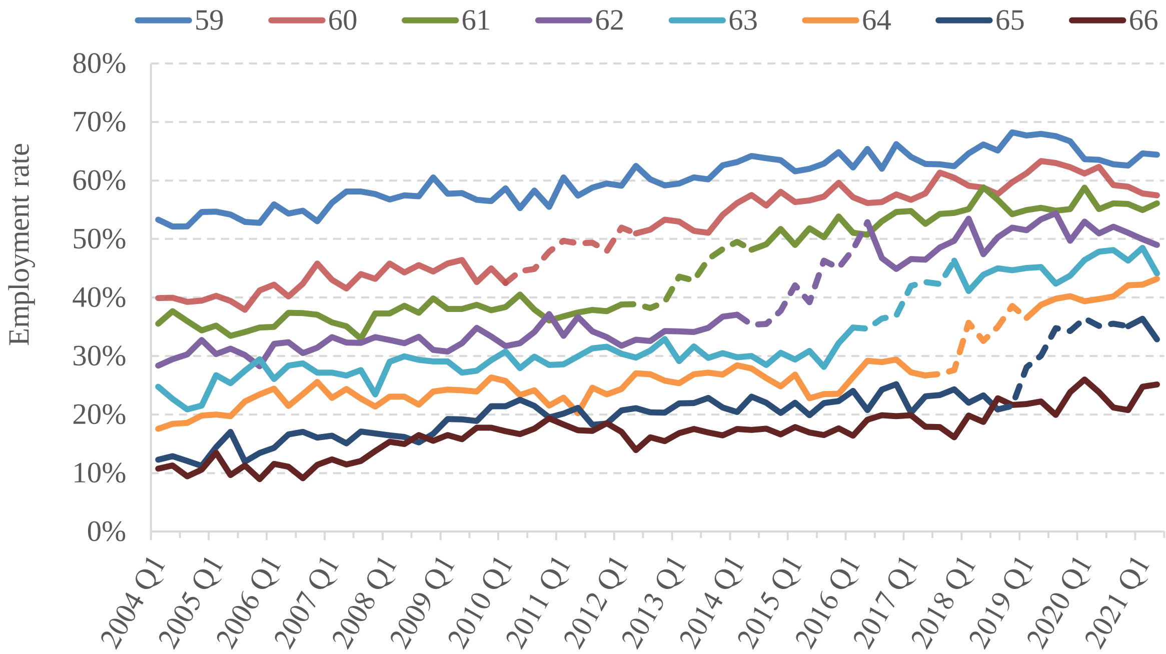

Notes: This figure shows the employment rates of 59- to 66-year-old women in the UK by quarter and single year of age between 2004 Q1 and 2021 Q2. The lines are dashed from the last quarter in which all women of that age were over the ERA until the first quarter in which all women of that age were under the ERA.

Second, to ascertain the extent to which wealth effects are driving the employment response, we compare the employment effect of being under the ERA for different birth cohorts who have lost different levels of wealth from the increase in the ERA, using a similar method to Atalay and Barrett (2015). This relies on the fact that those women reaching their ERA in their mid-60s have lost much more state pension income due to the reform than those women who reach their ERA just a few months after their 60th birthday, as calculated explicitly in Section VI.C. Given this, we can estimate the impact of wealth effects by interacting our birth cohort dummies with the treatment indicator, as in the following specification. Our coefficients of interest here are therefore φ_c.

$$(3) \quad y_{iact} = \alpha + \varphi_c * (underERA)_{iact} + \gamma_t + \lambda_c + \theta_a + X_{iact}\beta + \varepsilon_{iact}$$

To increase the interpretability of these results, and to make them more comparable with Atalay and Barrett (2015), we focus on six birth cohort groups for this analysis, defined based on a woman's ERA (age 60-61, 61-62, up to 65-66).

We can again estimate the dynamic treatment effects separately for different cohort groups by interacting the birth cohort dummies with the indicator variables denoting the number of quarters until the ERA as in Equation (4). To avoid overcrowding the graphs, we only focus on three birth cohort groups for this analysis (cohorts whose ERA is 60-61; 61-64; and 64-66).

(4) $$y_{iact} = \alpha + \sum_{j \neq -1} \varphi_{jc} \mathbb{1}(distERA_{iact} = j) + \gamma_t + \lambda_c + \theta_a + X_{iact}\delta + \varepsilon_{iact}$$

Third, we provide evidence that the UK economic and institutional environment, as described above, means that dismissals by employers of workers at or around the ERA are uncommon, and therefore that the increase in employment due to the rise in the ERA does not come from employers delaying such dismissals until the new ERA. To examine this, we run separate difference-in-difference models similar to Equation (1), but using as the outcome whether the individual left work within the last 3 months, broken down into four mutually exclusive and exhaustive potential reasons, as (self-)reported by the individual in the LFS: 1) Retirement; 2) Redundancy or other dismissal; 3) Health reasons; 4) Other. Increasing the ERA will reduce the number of people leaving work (the corollary of higher employment); but the test of whether dismissals are important is whether increasing the ERA reduces the number of people made redundant or dismissed.

VI. Results

A. Overall effect of increasing the ERA on employment

Table 2 reports the results from estimating Equation (1) on the sample of all women and on the main sample of those who were in paid work at age 59. Across all women, we estimate that being under the ERA increases the probability of being in paid work by 9.3 percentage points, which is precisely estimated (the 99% confidence interval is 7.8-10.8 percentage points). Conditioning only on those in paid work at age 59, the effect of being under the ERA is closer to 15 percentage points, again with a small standard error. Appendix Table 2 shows that our results are essentially identical when instead estimated using the BJS estimator.

Despite larger financial incentives to retire at the ERA in other countries (for example because their public pension is subject to an earnings test), our estimated employment response of 9.3 percentage points among all women is of a similar magnitude, albeit a little smaller, to the results from Staubli and Zweimüller (2013), who find an employment rate response of 11 percentage points for women in Austria, and only slightly lower than those from Geyer and Welteke (2021), who find a 13.5 percentage point employment rate response for women in Germany.

The remainder of the tables show the increases in different types of employment generated by the increase in the ERA from 60 to 66. The second row shows that being under the ERA increases the share of people who are working in a job they started in the last year by 0.3 percentage points. This is statistically significant at the 10% level, but the magnitude is tiny in comparison to the overall effect. This indicates that the vast majority of the employment effect of increasing the ERA comes from workers staying in their original job for longer, rather than inducing those who are not in paid work to return to employment, or those in paid work to move to a new job (rather than to leave paid work).

Table 2 Effect of increasing the early retirement age for women from 60 to 66 on different types of employment

Outcome	*Effect of being under ERA, all women*	*Mean for 59-year-olds pre-reform*	*Effect of being under ERA, conditional on being in paid work at 59*
In paid work	0.093*** [0.006]	56.7%	0.147*** [0.007]
Job tenure < 1 year	0.003* [0.001]	3.0%	0.004** [0.002]
Full-time work	0.055*** [0.005]	31.0%	0.088*** [0.007]
Part-time work	0.038*** [0.005]	25.7%	0.059*** [0.007]
Employed	0.089*** [0.005]	50.7%	0.141*** [0.007]
Self-employed	0.004 [0.003]	6.0%	0.006 [0.005]
Public sector	0.045*** [0.004]	28.2%	0.071*** [0.007]
Private sector	0.048*** [0.005]	28.5%	0.076*** [0.007]

Note: This table shows the effect of being under the ERA on different types of employment. Regression results are obtained by estimating Equation (1) separately for each row, changing the outcome variable to the outcome in the left-hand column. ***, ** and * denote that the effect is significantly different from zero at the 1%, 5% and 10% level respectively. Robust standard errors clustered at the year-and-month-of-birth level in parentheses. Pre-reform mean calculated on 59-year-olds during the period 2009Q2 to 2010Q1. Regressions unconditional on work status at age 59 contain 345,808 observations, while those conditional on being in paid work at 59 contain 214,177 observations. All regressions include as independent variables: quarter, age (in quarters), marital status, education level, ethnicity, housing tenure, region, financial year of birth, the quintile of an indicator of local area deprivation, dataset (either LFS or LLFS), partner's qualifications, whether partner under ERA, partner's age in quarters (in quadratic) and also the five-year band of partner's age.

This allows us to focus confidently on the results for those in paid work at age 59. The results in the remaining rows show that there was a slightly larger response from full-time work than from part-time work and that the almost all of the employment response comes from people working as employees rather than self-employment. The final two rows show that around half of the increase in employment comes from public-sector workers. This is in line with slightly less than half of 59-year-old women being employed in the public sector. The results in this table therefore demonstrate the importance of women's economic activity just before they reach ERA for explaining the estimated employment responses from reaching ERA.

Figure 3 shows the effect of distance to ERA on the probability of being in paid work for the set of women who were in paid work at age 59 from estimating equation (2). We estimate slightly positive coefficients of around 3 percentage points of being more than two quarters under ERA relative to being at most one quarter under ERA. This indicates that some women choose to retire just before their ERA. This is unsurprising as women might not retire exactly on the day they reach their ERA, but might retire at an appropriate point close to their ERA, (for example, teachers might retire at the end of the nearest school year). There is then a big fall in the probability of being in paid work in the quarter after reaching ERA of around 9 percentage points, with this effect increasing to around 14 percentage points over the following year or so, again indicating that many women retire close to, but not exactly at, their ERA. Appendix Figure 3 presents very similar estimates estimated using the Borusyak, Jaravel and Spiess (2024) estimator.[14]

We now turn to disentangling different mechanisms underlying the employment response in the rest of this section.

[14] The Appendix figure looks slightly different, with more negative coefficients both before and after the ERA in Appendix Figure 3 than in Figure 3. However, this can be explained by the difference in the base periods for the two estimators. In the dynamic TWFE specification, the coefficients present the estimated difference in the probability of being in paid work relative to period -1 (i.e. the quarter before reaching ERA). In BJS, the pre-treatment coefficients display the estimated difference in the probability of being in paid work relative to the omitted pre-treatment periods i.e. periods before -12, while the post-treatment coefficients display the estimated difference relative to all pre-treatment periods. Both plots imply that being in the quarter before ERA has a more negative effect on the probability of being in paid work than either being more than 12 quarters under ERA, or the average effect of being under ERA, which is why the difference in base periods makes a difference to the estimated coefficients. See Roth (2024) for more details on the differences in interpreting event study plots for these two methods.

Figure 3 Estimated effect of distance to early retirement age on the probability of being in paid work

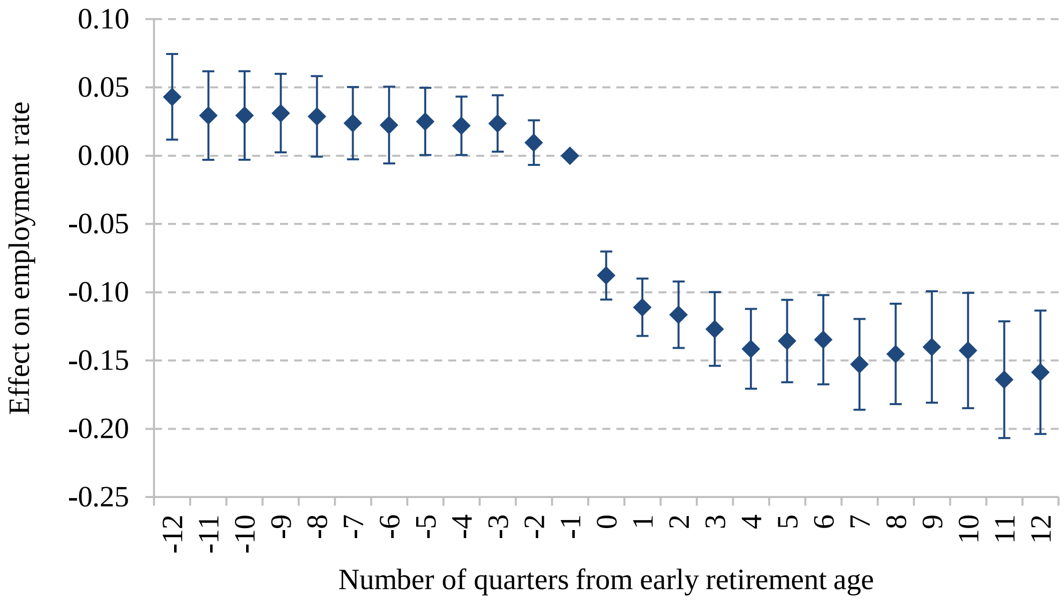

Note: This figure shows the effect of distance to the ERA on the probability of being in paid work. This is estimated using Equation (2) for the sample of women aged 59 to 67 who were in paid work at age 59 for the sample period 2009Q2 to 2021Q2. We only plot the coefficients and 95% confidence intervals for 12 quarters each side of the early retirement age for clarity. The regression includes as independent variables: quarter, age (in quarters), marital status, education level, ethnicity, housing tenure, region, financial year of birth, the quintile of an indicator of local area deprivation, dataset (either LFS or LLFS), partner's qualifications, whether partner under ERA, partner's age in quarters (in quadratic) and also the five-year band of partner's age.

B. The contribution of credit constraints to the employment response to the ERA

To separate out the effect of credit constraints on how employment changes in response to the ERA, we use ELSA data to compare the employment response for groups of women who are relatively likely to be credit constrained with the employment response for groups extremely unlikely to be credit constrained. Specifically, we estimate Equation (1) separately for women in each quartile of the distribution of liquid financial wealth in their mid 50s, as well as for the full ELSA sample.

The results are shown in Table 3. The overall employment response is 15.7 percentage points, similar to the 14.7 percentage point employment response we estimated using the LFS/LLFS as reported in Table 2, though with a slightly larger standard error. Table 3 shows that women with lower levels of liquid financial wealth in their mid-50s have a significantly larger employment response than wealthier women. This is especially pronounced for those in the lowest quartile of the wealth distribution, for whom being under the ERA increases their employment rate by almost 24 percentage points. This group of women is particularly likely to be credit constrained. In contrast, the employment rate for women in the top quartile of the liquid financial wealth distribution increases by 11 percentage points. The employment responses for the lowest and highest wealth quartile are statistically significantly different below the 5% significance level. Considerably smaller sample sizes in ELSA relative to the LFS/LLFS data preclude us from estimating dynamic treatment effects (as is shown in Figure 3 for the whole sample), separately for different wealth groups.

The fact that we find larger employment responses for women with lower levels of liquid wealth before reaching affected ages implies that credit constraints are a key mechanism underlying the overall employment response to the increase in the ERA. However, there are also positive and significant impacts among the top quartile of the liquid financial wealth distribution, who are extremely unlikely to be credit constrained, suggesting that credit constraints are not the only explanation. We examine other mechanisms behind this employment response in the following sections.

Table 3 Percentage point effect of increasing the early retirement age for women from 60 to 66 on employment rates for women by quartile of financial wealth during mid-50s

Group	*Effect of being under ERA*	*Std. error*	*Observations*	*P-value for difference in coefficients compared to lowest quartile*
All	0.157***	[0.021]	5,698	
Lowest wealth quartile	0.237***	[0.040]	1,429	
Second wealth quartile	0.158***	[0.045]	1,428	0.168
Third wealth quartile	0.143***	[0.048]	1,421	0.103
Highest wealth quartile	0.106**	[0.050]	1,415	0.037

Note: This table shows the effect of being under the ERA on employment, estimated using ELSA, and how this varies by initial wealth quartile. We restrict our sample to women aged 59-67 who were in paid work at age 59, and split women into quartiles based on their level of liquid net financial wealth in the wave when they are closest to age 55. Any individuals who never respond to the wealth questions are dropped from the regressions by wealth quartile. Regression results are then obtained by estimating Equation (1) both on the whole sample as well as separately for the set of women in each wealth quartile. ***, ** and * denote that the effect is significantly different from zero at the 1%, 5% and 10% level respectively. Robust standard errors clustered at the year-and-month-of-birth level in parentheses. All regressions include as independent variables: quarter, age (in quarters), marital status, education level, ethnicity, housing tenure, financial year of birth, partner's qualifications, whether partner under their ERA, partner's age in quarters (in quadratic).

C. The contribution of wealth effects to the employment response to the ERA

We now consider whether the wealth effect is driving any of the employment response we observe. To do this, we identify groups with different total wealth losses from the reform due to their date of birth and compare their employment responses. Women born later will have a higher ERA due to the reform, and so will have lost more years of state pension income than those whose ERA was unchanged or increased only slightly.

The third column of Table 4 illustrates how this loss in state pension income differs by birth cohort. We define six groups of women based on their ERA (which is a deterministic function of their date of birth). Then, for each cohort, we calculate the average (inflation adjusted) loss in state pension income due to the reform by summing all the basic state pension income they would have accumulated (assuming they have a full set of 35 qualifying years, as will be the case for the vast majority of these women) between the age of 60 (previously the ERA) and when they actually reached the ERA. There are large differences in the average loss in state pension income between these groups: those whose ERA ended up between 60 and 61 lost only on average around £3,000 ($3,900) of state pension income, while those whose ERA was 65 or 66 lost more than 10 times that – almost £40,000 ($52,000) on average.

We use Equation (3) to estimate how the employment effect of the ERA varies across these different birth cohorts, with the key coefficients being on the interactions between each cohort and the $underERA$ dummy (φ_c). If wealth effects were important for the employment response, we would expect to see larger employment effects for cohorts who had higher wealth losses (which is what Atalay and Barrett (2015) find). Instead, as is shown in Table 4, we find no evidence of a larger employment effect for cohorts with greater losses in state pension wealth.[15] If anything, there is some evidence of a *smaller* employment effect for the cohort born between August and December 1953 (who have an ERA of between 64 and 65), though the estimated employment response for those with an ERA of 65+ is much more similar to the younger cohorts. We plot the dynamic treatment effects by cohort, estimated using both TWFE and BJS, in Appendix Figure 4. This figure also shows no evidence of a larger employment effect for the later cohorts who faced a larger loss in state pension wealth.

[15] As noted earlier, after accounting for the trend of increasing female labour market participation, Morris (2022) finds an estimated employment effect that is much smaller than in Atalay and Barrett (2015) and not statistically different from zero. We do not use men as a control group for women. Although Morris (2022) does not repeat the specification where the treatment effect is allowed to vary by birth cohort, this could be one of the reasons for the discrepancy between the results in this paper and those in Atalay and Barrett (2015).

Table 4 Effect of increasing the early retirement age for women from 60 to 66 on employment for different birth cohorts, for those in paid work at age 59

Birth cohort	*ERA*	*Average loss in basic state pension income due to reform (£)*	*Effect of being under ERA*
Before 6 Apr 1950	60	0	0.135 [0.014]
6 Apr 1950 – 5 Apr 1951	60-61	3,061	0.160 [0.013]
6 Apr 1951 – 5 Apr 1952	61-62	9,364	0.143 [0.010]
6 Apr 1952 – 5 Apr 1953	62-63	15,860	0.162 [0.010]
6 Apr 1953 – 5 Aug 1953	63-64	22,675	0.141 [0.024]
6 Aug 1953 – 5 Dec 1953	64-65	29,457	0.087[††] [0.018]
6 Dec 1953 +	65+	39,636	0.118 [0.013]

Note: This table shows, for different birth cohorts of women, their average loss in basic state pension income due to the increase in the ERA, and their estimated effect of being under the ERA on employment. To calculate the average loss in state pension income, we take the average date of birth for each birth cohort and calculate the date a person with this date of birth would have reached their ERA if it had remained at 60, and the date they reach it under the reformed system. Then, we sum the amount of basic state pension income (adjusted for inflation) the individual would have received between these two dates. The final column shows the estimated coefficients and standard errors on the interaction between birth cohort and $underERA$ from Equation (3). Robust standard errors clustered at the year-and-month-of-birth level in parentheses. [†], [††] and [†††] denote effects are statistically significantly different from the estimated effect for the cohort born before 1950 (in the top row) at the 10%, 5% and 1% significance level, respectively. Regression sample contains 141,397 observations of women who were in paid work at age 59.

In Appendix Figure 5 and Appendix Table 4 we present additional evidence supporting our conclusion that credit constraints are an important mechanism while wealth effects are not. First, one concern with comparing employment responses for different birth cohorts to identify wealth effects is that there could be differences other than the loss in state pension wealth between birth cohorts that affect their employment response. Specifically, we might expect later born cohorts to have higher levels of wealth, and therefore to be less affected by credit constraints. However, Appendix Figure 3 shows that the distribution of wealth is very similar for earlier and later born cohorts in our analysis, suggesting that there are no large differences in how likely different birth cohorts are to be credit constrained in our analysis, and that we are indeed testing for the presence of wealth effects by comparing different cohort's responses.

A second concern could be that the larger employment responses we found for groups with lower levels of liquid wealth in the previous subsection might potentially reflect wealth effects rather than credit constraints. The change in public pension wealth due to the reform would represent a much larger proportional drop in wealth for those with lower levels of lifetime wealth. The analysis in this sub-section suggests that wealth effects are not an important mechanism, somewhat allaying this concern. However, it could still be possible that wealth effects, while not important for explaining the overall employment response, might be important in driving the employment response of subgroups with low levels of wealth. In Appendix Table 4, we show that this is not the case, as there are similar employment responses by birth cohort for women in the bottom two quartiles of the financial wealth distribution. This implies that the higher employment responses we found for subgroups with lower levels of liquid wealth indeed reflect the importance of credit constraints, not wealth effects.

In summary, the fact we find no relationship between the employment effects and the size of wealth loss caused by the reform points towards wealth effects not being an important driver of the employment response to the rising ERA.

D. The contribution of employer responses to the employment response to a higher ERA

Employer responses are unlikely to play a large role in people working longer in response to increases in the ERA in the UK due to the institutional environment, as argued in Section

II. To test this, we now formally examine whether we see any evidence that women working longer due to the ERA increase is brought about by employers delaying dismissing workers until they reach the new ERA.

Figure 4 shows that the proportion of women who were recently made redundant or otherwise dismissed from their job does not spike at the ERA during the period we examine. In the left panel, we take all observations of women aged between 55 and 65 between 2010Q2 and 2021Q1, and plot the proportion of them who report having left work in the previous three months for four (mutually exclusive and exhaustive) reasons, by how close they are to their ERA. There is a large spike in the proportion of women who leave work due to retirement at the ERA; however, the proportion who report having being dismissed or made redundant actually drops at this age. Part of this drop may be due to the fact that fewer women are in paid work to start with above the ERA. In the right panel, we therefore restrict the denominator to be women who were in paid work three months prior to their interview. Again, this figure shows that the large spike in the proportion of women leaving work at the ERA is due to an increase in retirements at this point – the proportion of female workers who report leaving work due to dismissals or redundancies remains low and, if anything, falls among workers who have recently reached their ERA.

Figure 4 Percentage of working women who leave paid work for different reasons over three months, by number of years to ERA (2010Q2 to 2021Q1)

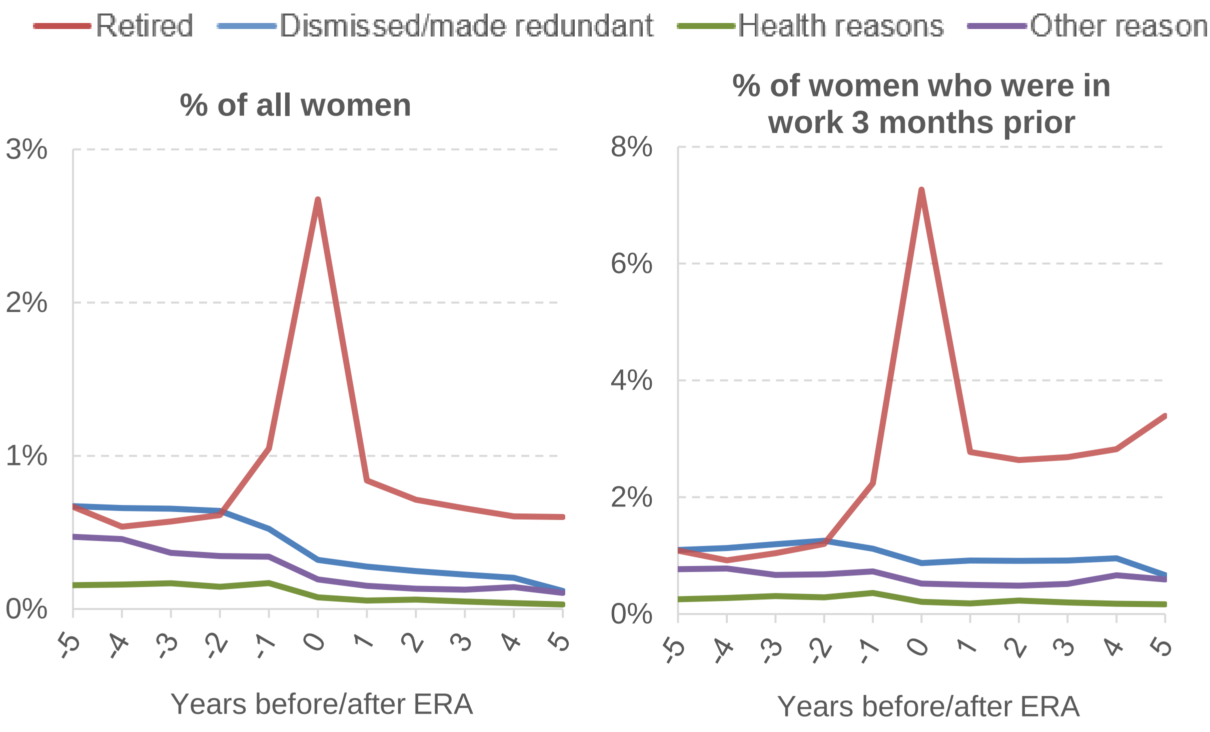

Notes: This figure shows the proportion of women who had left paid work for different reasons over the last three months, by how far they are from reaching their ERA. Panel A shows this across all women, while Panel B restricts the denominator to be women who were in paid work three months prior. The reasons for leaving paid work are self-reported by respondents to the LFS/LLFS.

In Table 5, we complement this graphical analysis by repeating our main regression for the sample of 59- to 67-year-old women who were in paid work on their 59th birthday, but changing the outcome variables to be whether the woman had left work in the last three months (for a specific reason). This table shows that being under the ERA reduces the probability of a women having left work in the last three months by 1.1 percentage points, consistent with the positive employment effects we found in Section VI.A. All of this reduction comes from women being less likely to retire when under the ERA by 1.6 percentage points; in fact, leaving work due to dismissals or redundancies if anything is slightly more common among women who are under the ERA as opposed to above it. In the right-hand column, we repeat this analysis for the restricted sample of those women who were in paid work three months prior, and again find only a negative effect of being under the

ERA on the probability of having left work due to retirement. This table therefore shows that all of the employment response to the ERA comes from women retiring at this point, rather than employers dismissing workers once they reach the ERA.

Table 5 Percentage point effect of increasing the early retirement age for women from 60 to 66 on probability of leaving work within the last 3 months, for different reasons

Outcome	*Effect of being under ERA, conditional on being in paid work at 59*	*Effect of being under ERA, conditional on being in paid work 3 months earlier*
Left work in last 3 months	–0.011***	–0.031***
	[0.002]	[0.003]
Due to dismissal/redundancy	0.002**	0.001
	[0.001]	[0.002]
Due to retirement	–0.016***	–0.035***
	[0.001]	[0.003]
Due to health reasons	0.002***	0.002***
	[0.000]	[0.001]
Due to other reasons	0.001	0.001
	[0.001]	[0.001]
Observations	344,246	124,635

Note: This table shows how the increase in the ERA affected the probability of having left work in the last 3 months for different reasons. The table contains regression results from estimating equation (1) on the sample of women who were in paid work at age 59, but where the outcome variables are whether the woman had left work in the last three months (for a specific reason). The results in the right-hand column are equivalent but are only estimated on the sample of women in paid work three months prior to interview. ***, ** and * denote that the effect is significantly different from zero at the 1%, 5% and 10% level respectively. Robust standard errors clustered at the year-and-month-of-birth level in parentheses.

E. The contribution of behavioural factors to the employment response to a higher ERA

Overall, the results in this paper indicate that credit constraints are an important mechanism for explaining why increases in the ERA increase employment rates during women's early 60s. In contrast, there is little evidence of wealth effects or employer responses playing an important role, and the marginal financial incentive to leave work does not increase at the

ERA in the UK. Despite this, women in wealthy families still exhibit significant employment responses to increases in the ERA, with an 11 percentage point boost in employment for women who were in the top quartile of the liquid wealth distribution in their 50s. This group is not credit constrained, nor forced out of work by employers. Wealth effects are not important for them and they do not face a particular financial incentive to retire at the ERA. This leaves a residual explanation for this group. This is suggestive, but not conclusive proof, that behavioural factors (such as signalling, social norms, or peer effects) are likely playing an important role in explaining the employment responses, at least for some people.

VII. Conclusion

There is increasing evidence of people continuing in employment and delaying retirement in response to governments increasing ERAs in public pension systems. Substantial increases in employment rates driven by these reforms have been found in ex-post analyses of reforms in Australia, Austria, Denmark, Estonia, France, Germany, Norway, and the UK. But to date there is much less concrete evidence on the mechanisms that drive these effects.

We exploit ten years of reforms in the UK which increased the ERA for women by six years (from 60 to 66) in just a decade (between 2010 and 2020), alongside large-scale household survey data. Together, this reform and data means we have power to identify heterogeneities precisely in responses to the reforms that shed light on the mechanisms driving them.

We do this in a context – the United Kingdom – where there is no earnings test on state pension benefits, and these pension benefits are not means-tested. This means there is little financial incentive to retire at the ERA (Banks et al., 2023), particularly in comparison to many European countries where are high implicit taxes on remaining in paid work in people’s early 60s (Borsch-Supan et al., 2021; 2023).

We find that women who are more likely to be credit constrained are more likely to delay their retirement as a result of the reform. Among those in paid work at age 59, those in the bottom quartile of the liquid financial wealth distribution see higher increases in employment rates as a result of the reforms (of 24 percentage points) compared to those in the top quartile of the distribution, who see an increase of 11 percentage points.

To identify the role of wealth effects, we compare birth cohorts that see substantially different wealth losses as a result of the reforms (such as those with an increase in their ERA of less than a year with those whose ERA rises by more than 5 years) and find no cross-cohort differences in how the higher ERA generates higher employment rates between 60 and their new ERA. This lack of relationship between wealth losses and employment responses holds even looking at women with particularly low levels of financial wealth before reaching their 60s, for whom the loss in state pension wealth represents a more significant share of their lifetime wealth. There is therefore little evidence of the wealth effects caused by reduced state pension wealth driving the increased employment rates resulting from ERA increases.

The UK is a country in which mandatory retirement at particular ages is all-but abolished, and employment decisions which discriminate on the basis of age are illegal. This makes it unlikely that employers dismiss workers at pension eligibility ages (as happens in the Netherlands). We confirm this by showing that there is no increase in redundancies or dismissals at the ERA, and that instead essentially all of the increased exits from work at the ERA are driven by people retiring, implying that these kinds of employer responses are not important in this context.

Given the lack of importance of financial incentives, wealth effects, and employer responses, the fact that there are still sizeable employment responses for wealthier groups implies there is a residual effect on employment that is not directly explained in this analysis.

This is suggestive that there is also likely an important role for behavioural effects, such as social norms around an appropriate age to retire, or framing, that are driving some people to delay retirement as a result of a higher ERA.

Determining which mechanisms may be driving the employment response to ERA increases has implications for our understanding of retirement behaviour and for improved accuracy when predicting the effect of future ERA increases. Our finding of similar employment responses across cohorts who experienced a ten-fold difference in their loss of public pension wealth may be considered in the context of the existing literature on wealth effects at older ages. Our results are somewhat at odds with Brown et al. (2010), who find that unexpected inheritances increase the probability of retirement, and Becker et al. (2023), who find large falls in employment for women with higher pension wealth in Germany, but are in line with the more mixed literature on retirement effects of unanticipated stock market gains (see Coile, 2015, for a review). Speculatively, one possible explanation for the disparate findings is that older workers may respond differently to the loss of pension wealth than the loss of financial wealth, perhaps for mental accounting reasons. Alternatively, wealth losses from state pension wealth in the UK (where the state pension sits alongside a large system of private pensions) may not be a particularly salient form of wealth loss.

Understanding whether credit constraints play a role in the employment response to ERA increases is important for understanding the distributional effects of pension reforms. Our finding that credit constrained individuals have a much stronger employment response indicates that these reforms can have a highly unequal effect on retirement behaviour. Given the recent implementation of automatic enrolment into employer-facilitated pensions for most employees in the UK (Cribb and Emmerson, 2020), credit constraints could be less important in the response to potential future ERA increases, as workers will likely accumulate larger pots of defined contribution wealth that can be flexibly accessed from a few years before the

ERA. This finding suggests that it may be important to look at the availability of other financial resources in a given country in order to predict the effect of ERA increases there.

Finally, our finding that ERA increases affect employment behaviour even among high-wealth individuals who are not credit constrained indicates a potential role for behavioural factors. This is a reminder that choice architecture within the pension system is a tool at policy makers' disposal that will influence retirement timing. Future research will be needed to understand better whether the root cause of this residual effect has to do with (for example) signalling, framing, or social norms around retiring at the ERA, and whether employer expectations play any role in this regard.

References

Amin-Smith, Neil and Rowena Crawford. 2018. 'State pension age increases and the circumstances of older women', in (eds) Banks, Batty, Nazroo, Oskala, and Steptoe *The Dynamics of Ageing: Evidence from the English Longitudinal Study of Ageing 2002–16 (wave 8).* IFS: London.

Artmann, Elisabeth, Nicola Fuchs-Schündeln, and Giulia Giupponi. 2023. 'Forward-looking labor supply responses to changes in pension wealth: Evidence from Germany', *CESifo Working Paper* No. 10427.

Atalay, Kadir and Garry F. Barrett. 2015. 'The impact of age pension eligibility age on retirement and program dependence: evidence from an Australian experiment', *The Review of Economics and Statistics*, 97(1): 71-87

Baker, Andrew C., David F. Larcker, and Charles C. Y. Wang. 2022. 'How much should we trust staggered difference-in-differences estimates?', *Journal of Financial Economics,* 144(2): 370-395.

Banks, James, Carl Emmerson, and David Sturrock, 2023. ‘Are Longer Working Lives a Response to Changing Financial Incentives? Exploiting Micro Panel Data from the UK’ in (eds) Börsch-Supan and Coile “*Social Security Programs and Retirement Around the World: The Effects of Reforms on Retirement Behavior*” University of Chicago Press. Chicago: 2023.

Becker, Sebastian, Hermann Buslei, Johannes Geyer, and Peter Haan. 2021. ‘Employment Responses to Income Effect: Evidence from Pension Reform’. *Discussion Papers of DIW Berlin* No. 1941.

Börsch-Supan, Axel and Courtney Coile. 2021. Introduction to “Social Security Programs and Retirement around the World: Reforms and Retirement Incentives” in (eds) Börsch-Supan and Coile “*Social Security Programs and Retirement around the World: Reforms and Retirement Incentives*” University of Chicago Press. Chicago, IL.

Börsch-Supan, Axel and Courtney Coile. 2023. Introduction and Summary in (eds) Börsch-Supan and Coile “*Social Security Programs and Retirement Around the World: The Effects of Reforms on Retirement Behavior*” University of Chicago Press. Chicago, IL.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess. 2024. ‘Revisiting Event Study Designs: Robust and Efficient Estimation’, *Review of Economic Studies*, 91(6): 3252-3285.

Brown, Jeffrey R., Courtney C. Coile, and Scott J. Weisbenner. 2010. ‘The Effect of Inheritance Receipt on Retirement’, *Review of Economics and Statistics* 92(2): 425-434.

Coile, Courtney C. 2015. ‘Economic Determinants of Workers’ Retirement Decisions’, *Journal of Economic Surveys* 29(4): 830-853

Cribb, Jonathan. 2023. ‘Understanding retirement in the UK’, *IFS Report 284*

Cribb, Jonathan and Carl Emmerson. 2019. ‘Can’t wait to get my pension: the effect of raising the female early retirement age on income, poverty and deprivation’, *Journal of Pension Economics and Finance*, 18(3): 450-472.

Cribb, Jonathan and Carl Emmerson. 2020. ‘What happens to workplace pension saving when employers are obliged to enrol employees automatically?’, *International Tax and Public Finance,* 27: 664-693.

Cribb, Jonathan, Carl Emmerson, and Gemma Tetlow. 2016. ‘Signals Matter? Large Retirement Responses to Limited Financial Incentives’, *Labour Economics,* 42: 203–12.

Cribb, Jonathan and Paul Johnson. 2022. ‘Policies to preserve worker-firm links during the pandemic: lessons from the UK’s Coronavirus Job Retention Scheme’ in (eds) Strain and Veuger *Preserving Links in the Pandemic: Policies to Maintain Worker-Firm Attachement in the OECD.* AEI Press: Washington DC.

Deshpande, Manasi, Itzik Fadlon, and Colin Gray. 2024. ‘How Sticky Is Retirement Behavior in the United States?’, *The Review of Economics and Statistics* 106(2): 370–383.

Geyer, Johannes, Peter Haan, Anna Hammerschmid, and Michael Peters. 2020. ‘Labor Market and Distributional Effects of an Increase in the Retirement Age’, *Labour Economics* 65: 101817.

Geyer, Johannes, and Clara Welteke. 2021. ‘Closing Routes to Retirement for Women: How Do They Respond?’, *Journal of Human Resources* 56(1): 311–41.

Gruber, Jonathan, Ohto Kanninen, and Terhi Revaska. 2022. ‘Relabeling, retirement and regret’, *Journal of Public Economics,* 211:104677

Johnsen, Julian Vedeler, Kjell Vaage, and Alexander Willén. 2022. ‘Interactions in Public Policies: Spousal Responses and Program Spillovers of Welfare Reforms’, *The Economic Journal* 132(642): 834–64.

Kanabar, Ricky. 2015. ‘Post-Retirement Labour Supply in England’, *The Journal of the Economics of Ageing* 6: 123–32.

Manoli, Dayanand S. and Andrea Weber. 2018. ‘The Effects of the Early Retirement Age on Retirement Decisions’. *NBER Working Paper* No. 22561.

Mastrobuoni, Giovanni. 2009. ‘Labor Supply Effects of the Recent Social Security Benefit Cuts: Empirical Estimates Using Cohort Discontinuities’, *Journal of Public Economics* 93(11): 1224–33.

Morris, Todd. 2022. ‘Re-Examining Female Labor Supply Responses to the 1994 Australian Pension Reform’, *Review of Economics of the Household* 20(2): 419–45.

Pensions Commission. 2004. “Pensions: Challenges and Choices: The First Report of the Pensions Commission” The Stationery Office.

Rabaté, Simon, Egbert Jongen, and Tilbe Atav. 2024. ‘Increasing the effective retirement age: key factors and interaction effects’, *American Economic Journal: Economic Policy*, 16 (1): 259–91.

Rabaté, Simon, and Julie Rochut. 2020. ‘Employment and Substitution Effects of Raising the Statutory Retirement Age in France’, *Journal of Pension Economics and Finance* 19 (3): 293–308.

Roth, Jonathan. 2024. ‘Interpreting Event-Studies from Recent Difference-in-Differences Methods’. arXiv:2401.12309.

Sæverud, Johan. 2024. ‘The Impact of Social Security Eligibility and Pension Wealth on Retirement’, *CEBI Working Paper* 05/24.

Seibold, Arthur. 2021. ‘Reference Points for Retirement Behavior: Evidence from German Pension Discontinuities’, *American Economic Review* 111(4): 1126–65.

Soosaar, Orsolya, Allan Puur, and Lauri Leppik. 2021. ‘Does Raising the Pension Age Prolong Working Life? Evidence from Pension Age Reform in Estonia’, *Journal of Pension Economics and Finance* 20(2): 317–35.

Staubli, Stefan, and Josef Zweimüller. 2013. ‘Does Raising the Early Retirement Age Increase Employment of Older Workers?’. *Journal of Public Economics* 108: 17–32.

Steptoe, Andrew, Elizabeth Breeze, James Banks, and James Nazroo. 2013. ‘Cohort profile: the English longitudinal study of ageing’, *International Journal of Epidemiology*, 42(6): 1640-48.

Appendix A - Supplementary Tables and Figures

Appendix Table 1 Illustrative change in the incentive to work at the early retirement age in the UK for a single individual (£ per week)

	In paid work	*Not in paid work with no private income*	*Income change from working*	*Change in incentive above ERA*
Pre April 2016				
Below ERA	Salary	JSA (£73)	Salary - £73	
Above ERA	Salary + BSP (£116)	BSP (£116) + PC (£35)	Salary - £35	+£38
Post April 2016				
Below ERA	Salary	JSA (£75)	Salary - £75	
Above ERA	Salary + NSP (£175)	NSP (£175)	Salary	+£75

Note: This table shows incomes in and out of work for a single person with no income other than state pension (either basic state pension, BSP, or new state pension, NSP), pension credit (PC), jobseeker's allowance (JSA), or income from employment. Throughout we assume the individual has accumulated the maximum number of qualifying years for the state pension. Example individual is assumed to have no housing costs, not to be disabled, to have financial assets below £16,000, and to have no other adults or children living in their household. This does not consider the impact of taxes, where the only relevant change is that employee National Insurance Contributions are no longer paid on earnings once an individual reaches the ERA, which also strengthens rather than weakens the financial incentive to work at this point.

Appendix Table 2 Effect of increasing the early retirement age for women from 60 to 66 on different types of employment – comparing TWFE and BJS estimates

Outcome	*TWFE estimates*	*BJS estimates*
In paid work	0.147***	0.155***
	[0.007]	[0.006]
Job tenure < 1 year	0.004**	0.001
	[0.002]	[0.003]
Full-time work	0.088***	0.107***
	[0.007]	[0.008]
Part-time work	0.059***	0.049***
	[0.007]	[0.008]
Employed	0.141***	0.156***
	[0.007]	[0.007]
Self-employed	0.006	-0.001
	[0.005]	[0.005]
Public sector	0.071***	0.084***
	[0.007]	[0.007]
Private sector	0.076***	0.071***
	[0.007]	[0.007]

Note: This table shows the effect of being under the ERA on different types of employment, estimated using TWFE and the BJS estimator. ***, ** and * denote that the effect is significantly different from zero at the 1%, 5% and 10% level respectively. Robust standard errors clustered at the year-and-month-of-birth level in parentheses. TWFE regressions contain 214,177 observations while BJS regressions contain 168,885 observations. The sample for the BJS regression includes 59- to 65-year-old women who were in paid work at age 59, while the sample for the TWFE regression includes 59- to 67-year-old women who were in paid work at age 59. When estimating the BJS regressions, we specify the treatment as being over, rather than under the ERA, but reverse the sign of the results in this table to aid comparability with the TWFE estimates. The BJS regressions also only include controls for age (in quarters), time (in months), ethnicity, educational qualification level, housing tenure, region, dataset, relationship status, and local area deprivation as set out in Appendix Section B.

Appendix Table 3 Percentage point effect of increasing the early retirement age for women from 60 to 66 on employment rates by quartile of financial wealth during mid-50s – comparing TWFE and BJS estimates

Group	*TWFE estimates*	*TWFE observations*	*BJS estimates*	*BJS observations*
All	0.157***	5,698	0.123***	4,075
	[0.021]		[0.025]	
Lowest wealth quartile	0.237***	1,429	0.166***	1,046
	[0.040]		[0.038]	
Second wealth quartile	0.158***	1,428	0.125**	989
	[0.045]		[0.051]	
Third wealth quartile	0.143**	1,421	0.077	946
	[0.048]		[0.052]	
Highest wealth quartile	0.106**	1,415	0.078	1,026
	[0.050]		[0.055]	

Note: This table shows the effect of being under the ERA on employment and how this varies by initial wealth quartile, estimated using TWFE and the BJS estimator. We first restrict our sample to women aged 59-67 who were in paid work at age 59, and split women into quartiles based on their level of financial wealth in the wave when they are closest to age 55. The sample for the BJS regression then includes 59- to 65-year-old women who were in paid work at age 59, while the sample for the TWFE regression includes 59- to 67-year-old women who were in paid work at age 59. We estimate separate regressions for each quartile of initial net liquid financial wealth. When estimating the BJS regressions, we specify the treatment as being over, rather than under the ERA, but reverse the sign of the results in this table to aid comparability with the TWFE estimates. The BJS regressions also only include controls for age (in quarters), time (in months), ethnicity, educational qualification level, housing tenure, region and marital status as set out in Appendix Section B. ***, ** and * denote that the effect is significantly different from zero at the 1%, 5% and 10% level respectively. Robust standard errors clustered at the year-and-month-of-birth level in parentheses.

Appendix Table 4 Effect of increasing early retirement age for women from 60 to 66 on employment for different birth cohorts, separately for women in the bottom two quartiles of the liquid net financial wealth distribution

Birth cohort	*ERA*	*Effect of being under ERA for bottom wealth quartile*	*Effect of being under ERA for second wealth quartile*
Before 6 Apr 1952	60-62	0.224***	0.151***
		[0.050]	[0.052]
After 6 Apr 1952	62-66	0.257***	0.146***
		[0.058]	[0.074]

Note: This table shows the effect of being under the ERA on employment for two different birth cohorts of women from the lowest two quartiles of initial liquid financial wealth. We first restrict our sample to women aged 59-67 who were in paid work at age 59, and split women into quartiles based on their level of financial wealth in the wave when they are closest to age 55. Data is from the English Longitudinal Study of Ageing. For women in the bottom two quartiles of initial financial wealth, we then separately estimate Equation (2) and show the estimated coefficients and standard errors on the interaction between birth cohort and $underERA$. There are 1,429 observations of 511 women in the bottom wealth quartile, and 1,428 observations of 517 women in the second wealth quartile. ***, ** and * denote that the effect is significantly different from zero at the 1%, 5% and 10% level respectively. Robust standard errors clustered at the year-and-month-of-birth level in parentheses.

Appendix Figure 1 Calculated employment shares in the month reaching age 59, for women born in 1952, by current age

Notes: This figure shows the estimated share of women born in 1952 that were in different labour market states in the month they reached their 59th birthday in 2011, by current age. The blue bar shows the share of women born in 1952 who are in paid work when we observe them at their current age. Then, among those women who are not in paid work at that age, the red bar shows the proportion who report last having worked on or since the month they reached their 59th birthday, the green bar shows the proportion who report last having worked before their 59th birthday, and the purple bar shows the proportion who report having never been in paid work. Women in the blue and red bar groups form our sample of women who were in paid work at age 59.

Appendix Figure 2 Employment rates over time of women who were in paid work at age 59, by single year of age

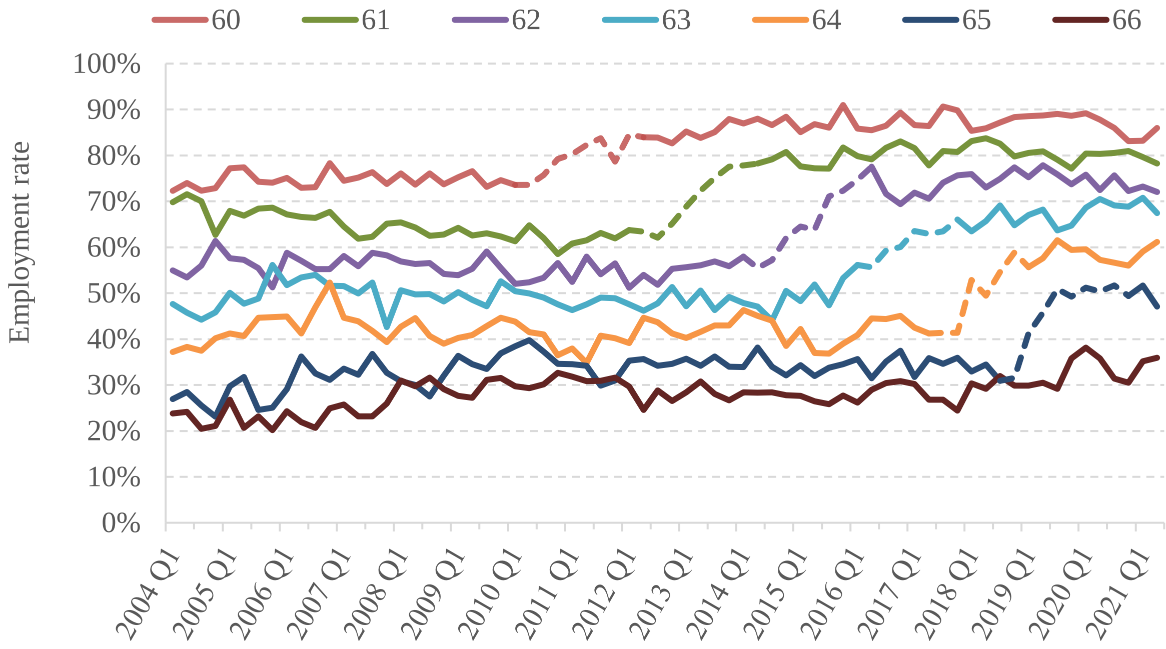

Notes: This figure shows the employment rates of 60- to 66-year-old women in the UK who were in paid work at age 59, by quarter and single year of age between 2004 Q1 and 2021 Q2. The lines are dashed from the last quarter in which all women of that age were over the ERA until the first quarter in which all women of that age were under the ERA.

Appendix Figure 3 Estimated effect of distance to early retirement age on the probability of being in paid work – BJS estimates

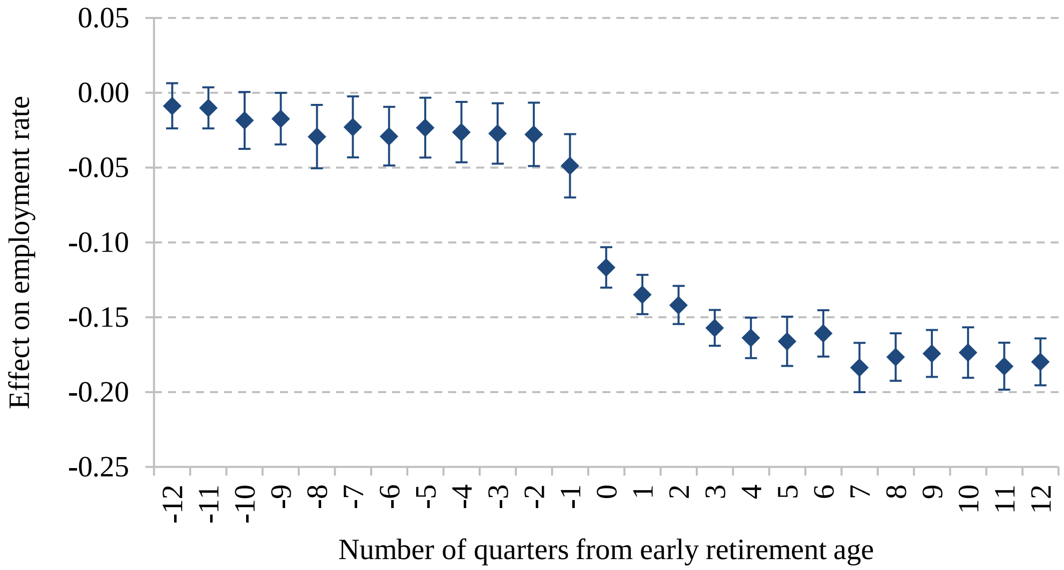

Note: This figure shows the effect of distance to the ERA on the probability of being in paid work, estimated using the BJS method. The sample is women aged 59 to 65 who were in paid work at age 59 for the sample period 2009Q2 to 2021Q2. We only plot the coefficients and 95% confidence intervals for 12 quarters each side of the early retirement age for clarity. Error bars show 95% confidence intervals, where robust standard errors are clustered at the year-and-month-of-birth level.

Appendix Figure 4 Estimated effect of distance to early retirement age on the probability of being in paid work for different birth cohorts

Panel A: TWFE estimates

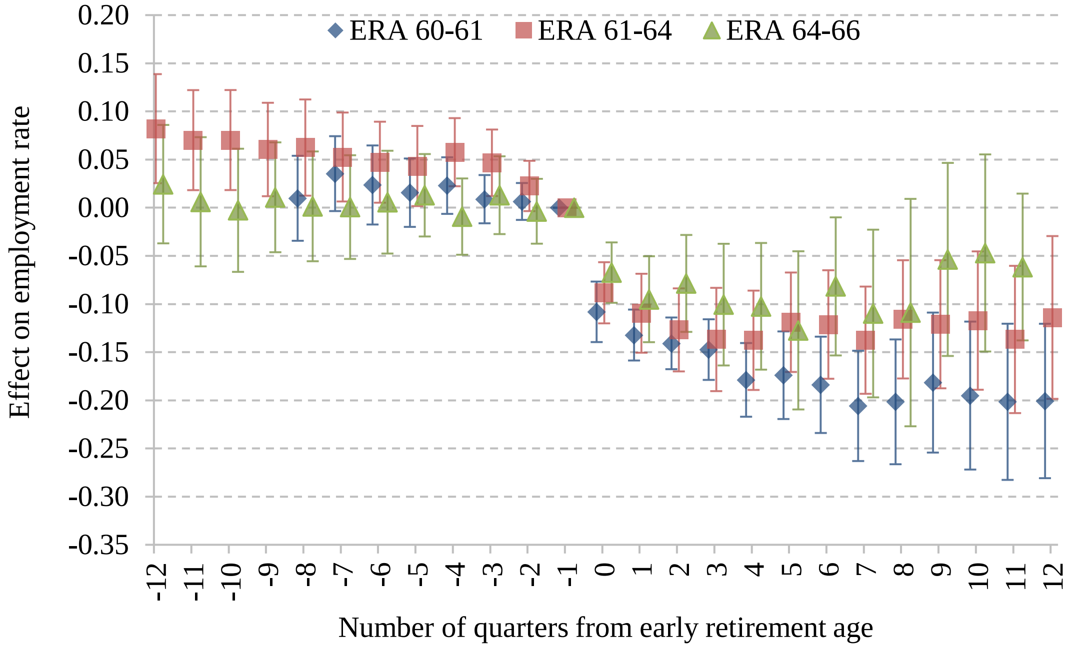

Panel B: BJS estimates

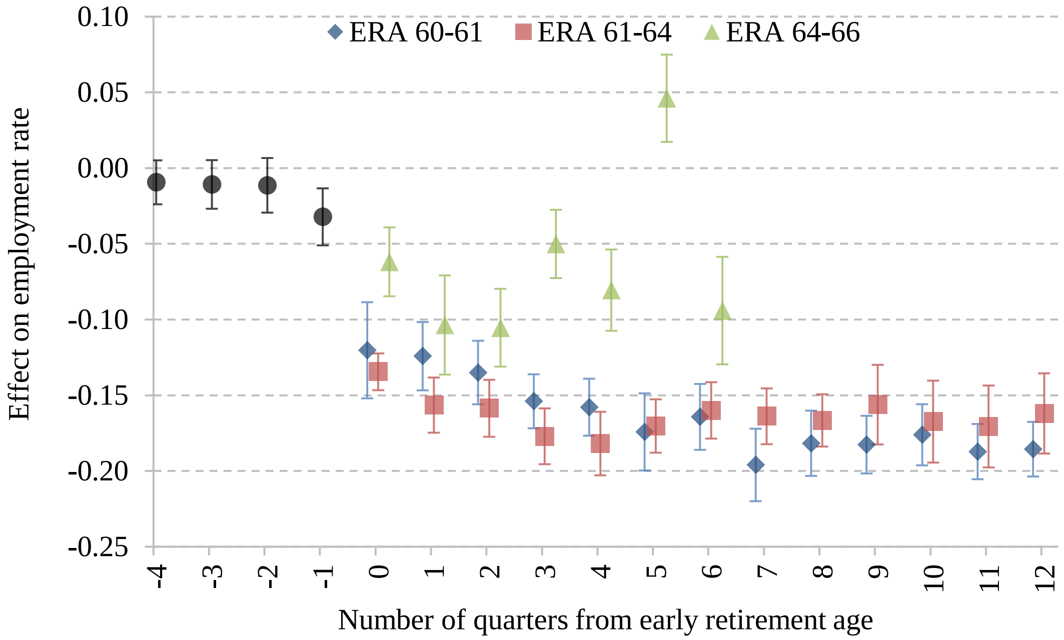

Note: This figure shows the effect of distance to the ERA on the probability of being in paid work for different cohorts of women, based on their early retirement age. Panel A shows the results estimated using Equation (4), while Panel B shows the results estimated using the BJS estimator. Error bars show 95% confidence intervals, where robust standard errors are clustered at the year-and-month-of-birth level. The sample includes women who were in paid work at age 59 for the period 2009Q2 to 2021Q2, restricted to those aged 59 to 67 in Panel A and 59 to 65 in Panel B. The BJS regressions also only include controls for age (in quarters), time (in months), ethnicity, educational qualification level, housing tenure, region, dataset, relationship status, and local area deprivation as set out in Appendix Section B.

Appendix Figure 5 Comparing wealth across birth cohorts

Notes: This figure shows the share of women in each quartile of the mid-50s wealth distribution, by ERA, calculated using the English Longitudinal Study of Ageing. We first restrict our sample to women aged 59-67 who were in paid work at age 59, and split women into quartiles based on their level of financial wealth in the wave when they are closest to age 55. We then calculated the share of women in each of these quartiles for four groups defined by their ERA, determined by their date of birth. Women with an ERA of 60 were born before 6 April 1950; between 60 and 62 were born between 6 April 1950 and 5 April 1952; between 62 and 65 were born between 6 April 1952 and 5 December 1953; of at least 65 were born between from 6 December 1953 onwards.

Appendix B: Using the Borusyak, Jaravel, and Spiess (2024) estimator

In Appendix Tables 2 and 3, we re-estimate our key empirical results using the estimator proposed by Borusyak et al. (2024; henceforth "BJS"). In this section of the appendix, we first describe the intuition behind this estimator, and then provide more details on the method we use to perform this estimation.

The BJS estimator is used to estimate the causal effects of a binary treatment in settings of staggered treatment rollout and potentially heterogeneous treatment effects. We define treatment as being over the ERA in our context.[16] Denoting the period-t (stochastic) potential outcome of individual i if she is never treated (i.e. always under ERA) by $y_{iact}(0)$, the causal effect of treatment for this woman is then $\tau_{iact} = \mathbb{E}[y_{iact} - y_{iact}(0)]$. The overall treatment effect is then the average of all these observation-level treatment effects for women over ERA.

We can represent the BJS estimator as an imputation procedure. The first step involves specifying a model for $y_{iact}(0)$, and estimating the parameters of this model. We specify this model as

(5) $\mathbb{E}[y_{iact}(0)] = \alpha + \gamma_t + \theta_a + X_{iact}\delta,$

where γ_t and θ_a are age and time-period fixed effects, measured in years and month for time and year and quarter for age. X_{iact} contains ethnicity, educational qualification level, housing tenure, region, dataset, relationship status, and local area deprivation. Note that we therefore use a more limited set of controls when using the BJS estimator compared to when estimating using TWFE – this is for computational reasons. Using this narrower set of controls rather than our full set of controls for TWFE estimation has little effect on our

[16] Note that when estimating using TWFE, we define treatment as being under the ERA. While this only changes the sign (but not the magnitude) of coefficients estimated using TWFE, this definition can affect both the magnitude and sign of coefficients estimated using the BJS estimator. Results estimated using BJS but where treatment is defined as being under the ERA are, however, similar to the results presented in Appendix Tables 2 and 3 (but of opposite sign) and are available on request.

results; these results are available upon request. The parameters α, γ_t, θ_a and δ are estimated in the first step using OLS on the sample of all untreated observations. Note that we cannot estimate age-specific fixed-effects for women older than 65, as they are always over the ERA and are therefore always treated. For this reason, we restrict our sample for BJS estimation to women aged 59 to 65.

In the second step of the procedure, $y_{iact}(0)$ is estimated as $\hat{y}_{iact}(0) = \hat{\alpha} + \hat{\gamma}_t + \hat{\theta}_a + X_{iact}\hat{\delta}$ for all treated observations (i.e. all observations of women over the ERA). Then, individual treatment effects are estimated as $\hat{\tau}_{iact} = y_{iact} - \hat{y}_{iact}(0)$, and the overall average treatment effect is obtained as the average of these simulated individual treatment effects over all treated observations. We can also obtain the average treatment effect at event time $j \geq 0$ by averaging the estimated treatment effects only for the set of treated observations such that $distERA_{iact} = j$, and we can obtain average treatment effects for different subgroups by averaging the treatment effects only among treated observations in each subgroup. We follow the methodology outlined in Borusyak et al. (2024) to calculate standard errors clustered at the month-of-birth level.

There are two key assumptions for identification. The first is that treated and untreated observations would have had the same expected counterfactual had treatment not occurred, i.e. Equation (5) holds for all i and t. This is essentially the (generalized) parallel trends assumption common to all difference-in-differences estimation. The second assumption is that there are no anticipation effects, that is $y_{iact} = y_{iact}(0)$ for all untreated observations.