Are longer working lives a response to changing financial incentives? Exploiting micro panel data from the UK

James Banks[†], Carl Emmerson[‡] and David Sturrock[1]

Abstract

In the UK, as in many advanced economies, male and female employment rates at older ages have been rising consistently since the mid-1990s. To what extent do changes to pensions systems explain this trend? In this paper, we use the English Longitudinal Study of Ageing to construct individual-level measures of the financial incentives to remain in paid work that pensions systems provide. We use UK state pension rules, alongside the rich information on individual private pensions that is elicited in ELSA, to calculate the overall implicit tax or subsidy of earnings that each individual faces due to the change in their pension entitlements that results from an additional year of work. Linear probability and probit models find a positive association of wealth with labour market exit, and a mild but positive association with implicit taxes. Including individual fixed effects suggests that the association of labour market exit with wealth is not causal, but likely to be driven by unobserved heterogeneity, and that a 10ppt rise in the implicit tax rate causes a statistically significant 0.4ppt rise in the one-year-ahead exit rate (compared to a baseline rate of 8% amongst 55 to 74 year-olds). We find consistent evidence that crossing the State Pension Age threshold causes a 4-5ppts rise in the labour market exit rate, even when controlling for the changing financial incentives to retire that occur at that age.

Acknowledgments

This paper forms part of the National Bureau of Economic Research's International Social Security (ISS) Project. The authors are grateful to the other participants of the ISS project for useful comments and advice. We are grateful to the ESRC-funded Centre for the Microeconomic Analysis of Public Policy at IFS (grant reference ES/M010147/1) and the "Policies for longer working lives: understanding interactions with health and care responsibilities" project (grant ES/P001688/1) that is part of the Joint Programme Initiative "More Years, Better Lives" for providing funding or this project. Data from the Family Expenditure Survey (FES) and the Labour Force Survey (LFS) were made available by the UK Data Archive. ELSA was developed by a team of researchers based at the National Centre for Social Research, University College London and the Institute for Fiscal Studies. The data were collected by the National Centre for Social Research. The data creators, depositors, copyright holders and funders bear no responsibility for the analysis or interpretation of the data presented here. Responsibility for interpretation of the data, as well as for any errors, is the authors' alone.

[†] University of Manchester and Institute for Fiscal Studies.
[‡] Institute for Fiscal Studies.
[1] Institute for Fiscal Studies and University College London.

1. Introduction

Recent decades have seen rapid rises in longevity at older ages across advanced economies. The UK is no exception to this trend, with life expectancy at age 60 increasing by around one year with every five years that have passed since 1980. The consequent ageing of the population is projected to put stress on healthcare and pensions systems in the coming decades.

Alongside this increased longevity, since the early-to-mid 1990s, employment rates at older ages have been rising. The employment rate for UK men aged 55 to 69 has risen from 43% in 1995 to 57% in 2019. Female employment at the same ages has risen even more rapidly, from 28% in 1995 to 48% in 2019. But this strong increasing trend has not always been the predominant pattern. In the 20 years from the mid-1970s to the mid-1990s, the trend in the UK, as in a number of countries, was for a decreasing proportion of older men to be in work. The female employment rate over the same period remained broadly flat.

Why have employment rates at older ages been on a strong rising trend since the early 1990s and do changes to the UK pensions system explain the dramatic turnaround in the trends in employment rates? More generally, what role do incentives provided by the state and private pensions system play in determining whether or not individuals at older ages stay in, or leave, employment? The UK pension system has undergone near-constant reform since 1975 and a number of changes have the potential to drive some of the employment patterns observed. As reforms continue to roll out and governments around the world consider further policy changes in the face of pressures from demographic change, it is crucial to understand the role that financial incentives play in shaping individual choices.

In this piece of research, we exploit a panel dataset containing detailed information on pension entitlements of a representative sample of the English population over the period 2002–03 to 2018–19. Constructing individual-level measures of the implicit tax on earnings provided by the state and private pension systems, we assess whether changing financial incentives explain individual choices to exit the labour force.

The UK pension system is composed of both publicly-funded state pensions and private pension provision. Unlike in many other advanced economies, private occupational pensions are an integral part of the UK system, providing the main replacement of earnings for a substantial proportion of individuals. The Basic State Pension (BSP), the first pillar of the UK system, is paid at the same rate to all individuals with the same number of years of credited activities (including most paid work and, in some circumstances, caring for a child or being in receipt of unemployment or disability benefits). The second pillar of the UK pension system is made up of state earnings-related pensions and private occupational or personal pensions. In the period where state earnings-related schemes were open to accrual, individuals could either build up entitlement to additional state pension income or accrue a private pension and, in return, pay a reduced rate of payroll tax. Occupational pension schemes for the birth cohorts we analyse are split between defined benefit schemes that typically pay a pension based on job tenure and final salary, and defined contribution schemes. Occupational DB schemes typically have a normal retirement age (NRA) and penalise the drawing of pension income before this age and, for those who have reached a maximum number of years of accrual in the pension, penalise delaying drawing the pension beyond that point.

We construct a panel dataset which includes measures of the financial incentives to stay on in paid work that are provided by these pension systems, for a representative sample of the English household population. The English Longitudinal Study of Ageing (ELSA) provides us with rich information on the private pension coverage of a sample of around 8,000 to 10,000 individuals on a biennial basis. A range of questions are designed to elicit the parameters of each pension scheme of which an individual is a member. Using these, we are able to construct a measure of the future pension income that a worker would receive both in the case where they left the labour market immediately and began to draw their pension, and in the case where they worked for an additional year before retiring and claiming. This then allows us to calculate a measure of the implicit tax or subsidy of earnings that is created by these pensions. To calculate the equivalent incentives provided by the state pension system, we need to know an individual's full history of employment (and other credited activities) and earnings, as these determine their entitlements. For the majority of individuals in our sample, we are able to use employment histories elicited as part of the 'ELSA Life' survey that was a follow-up to wave 3 of the main ELSA survey. Unfortunately, we do not have individuals' full earnings histories but we impute these on the basis of their reported earnings in the ELSA survey.

Having constructed a measure of individuals' pension incomes in the case where they leave work, and in the case where they stay on in work, we obtain two measures of the overall financial incentive to work that an individual faces, as a result of the pension system. The first is the individual's already-accrued pension wealth, measured as the net present value of the future income to which they would be entitled if they retired immediately. The second is the implicit tax (or subsidy) on earnings that they face due to the reduction (or increase) in pension wealth that they would accrue by working for an additional year (once social security taxes are netted off). Financial incentives vary across individuals and across time for a host of reasons. For example, whether an individual stands to increase their state pension entitlements by working for an additional year depends on whether they have already achieved the 'full' contribution history required to receive a full BSP (and the number of years required has varied over time). The level of this accrual in turn depends on the system of uprating, earnings accrual factor, and state pension age (SPA) that applied to an individual of that sex and cohort under the policy regime in place at that particular point in time. Private pension incentives depend on the individual's career history in terms of job tenure and final earnings, and the way in which this interacts with the schemes offered by their current and past employers.

We use the variation in our measures of pension wealth and implicit tax to examine the relationship between incentives and labour force exit. Using the sample of individuals in ELSA aged between 50 and 74 and in paid work, we estimate both linear probability and probit models where our outcome variable is a binary indicator for whether an individual left work in the year following their interview. We examine the effect of wealth (pension wealth plus net financial wealth) and implicit tax on this probability, while controlling for potentially important covariates such as individual health, sex and job type. There is a potential concern that macroeconomic trends, changing labour force attachment across cohorts, or norms around retirement at particular ages, could confound our estimates. Fortunately, as we have rich variation in incentives even within time-, cohort- and age-groups, we can assess the effect of incentives while controlling for these in various ways. Exploiting the panel nature of the data, we run further specifications that control for unobserved heterogeneity across individuals.

In both our linear probability and probit models, we find a positive relationship between wealth and labour force exit and a negative relationship between implicit tax and labour force exit. These findings are robust to the inclusion of various controls for birth cohort, age and time, although implicit tax is only marginally statistically significant in some specifications. In the linear probability model, the inclusion of individual fixed effects leads us to find a larger effect of implicit tax that is highly statistically significant and a smaller, and no longer significant, effect of wealth. One interpretation of such a finding is that those who have a higher propensity to remain in work (perhaps because they have lower disutility from working, all else equal) tend to have lower wealth and lower implicit tax rates and that this confounds the estimates that do not control for fixed effects.

We also examine the impact on labour force exit of an individual reaching their SPA – the age where they can first draw their state pension. We are able to identify this effect separately from age effects, even while controlling for age non-parametrically with a full set of age dummies, due to the fact that SPA for women increased during much of our data period. We find a quantitatively large effect of crossing the SPA threshold, leading to an increased propensity to retire, consistent with Cribb et al., (2016). We demonstrate for the first time that this threshold has a strong effect on behaviour, even when controlling for individual-specific measures of the financial incentives to retire at the SPA, consistent with it being a 'signal' to retire.

We seek to assess the impact of pension reforms by using our estimated regression model to predict the labour force exit rates that would have been seen in a counterfactual situation where reforms to pension systems did not take place. We find that despite the significant incentive effects that we identify, the change in incentives over our data period that is attributable to pension reforms is modest. Consequently, we find that pension reforms are not a primary explanation for the changing employment rates that have been seen in the UK.

This paper contributes to the literature that seeks to understand the drivers of employment at older ages (Coile et al., 2019*,* and Börsch-Supan and Coile, 2021) and, in particular, the effects of financial incentives provided by pension and social security systems (Wise, 2016; Gruber and Wise, 1999; 2005). These previous cross-country analyses have found that countries with pension systems that provided stronger incentives for work at older ages had higher employment rates of older workers and that, within countries, individuals with stronger incentives to stay on in work were more likely to do so. In the UK context, Blundell and Johnson (1999) calculated the financial incentives that state pensions and occupational pensions provide for the decision to retire and compared these to patterns in retirement behaviour. Blundell et al., (2002) uses the 1988–89 and 1994 waves of the UK Retirement Survey to construct individual incentive variables and finds a strong role for financial incentives, and in particular the incentives provided by occupational pension schemes, in the retirement decision. The present paper takes a similar approach in constructing individual-level incentives and relating them to labour market behaviour though does so in a multiple wave panel setting. This is in contrast to the approach in Banks and Emmerson (2021), which calculated a time series of incentives for some typical 'example' UK individuals and looked at the correlation between incentives and aggregate employment rates.

This study is able build upon earlier work and to address some of the limitations resulting from the data available at the time. The ELSA panel contains more detailed occupational pensions than the earlier Retirement Survey and, for the first time, allows us to exploit comprehensive information on the scheme rules of occupational pensions at the individual level when analysing older individuals' employment choices. The longer ELSA panel allows us to capture greater variation in incentives across time and across cohorts than in earlier, individual-level studies, and to exploit a number of reforms that took place over the data period. Furthermore, with individuals observed multiple times, we can control for individual-level unobserved heterogeneity. While the present research is able to construct incentives for a representative panel of actual individuals, as opposed to the example individuals examined in Banks and Emmerson (2021), this does come at the cost of restricting our study to the period since 2002. Hence we are focused solely on the period during which employment rates were rising and cannot directly analyse the both the periods of falling and rising employment rates.

This paper is organised as follows. Section 2 describes trends in employment, pension coverage and policy, giving the context for our study. Section 3 describes our method for constructing individual-level measures of the incentives to stay on in work. Section 4 looks at the trends in our incentive measures and our outcome of interest: labour force exits. In section 5, we set out our approach to estimating the effects of incentives on the probability of leaving work. In section 6, we attempt to give a more direct quantification of the effect of pension reforms on labour market behaviour by using our estimated regression model to simulate exit rates in counterfactual scenarios where pension reforms had not taken place. Section 7 concludes.

2. Context: labour market trends, pension membership, and policy

In this section, we set out the aggregate trends in employment that we are examining, and the pensions environment and policy reforms that are a potential explanation for observed trends. The period for the data that we use in our regression analysis is 2002–03 to 2018–19, but in this section we also set both policy and labour market trends in a longer-term context.

2.1 Labour Market Trends

Figure 2.1 and Figure 2.2 use UK Labour Force Survey data to plot aggregate employment at older ages since 1975 for men and women, respectively. Male employment at age 55 to 69 has followed a well-documented "U-shape" pattern over the past 40 years. Employment fell from 59% in 1975 to 42% in 1994 before rising again to 57% in 2019. Splitting out these trends by 5-year age bands reveals that while employment at ages 65 to 69 has risen from a low of 11% in 1988 to reach 27% in 2019, substantially exceeding its 1975 level of 21%, employment at ages 55 to 59 stands at 79% in 2019, up from a low of 66% in 1995 but still below the 87% reached in 1977.

Female employment rates show a different pattern, having stayed approximately flat from 1975 until the early 1990s, they have risen consistently across old age-groups since then. Over ages 55 to 69, employment was lowest in 1988 at 23% and has risen to 48% in 2019. We can see the particularly sharp rise in employment rates for the 60 to 64 age-group from 2011 onwards that coincided with the increase in the female SPA (outlined below). Employment for this group stood at 33% in 2010 before rising to 50% by 2019.

Figure 2.1 Employment rates of men, by age band, 1975–2019

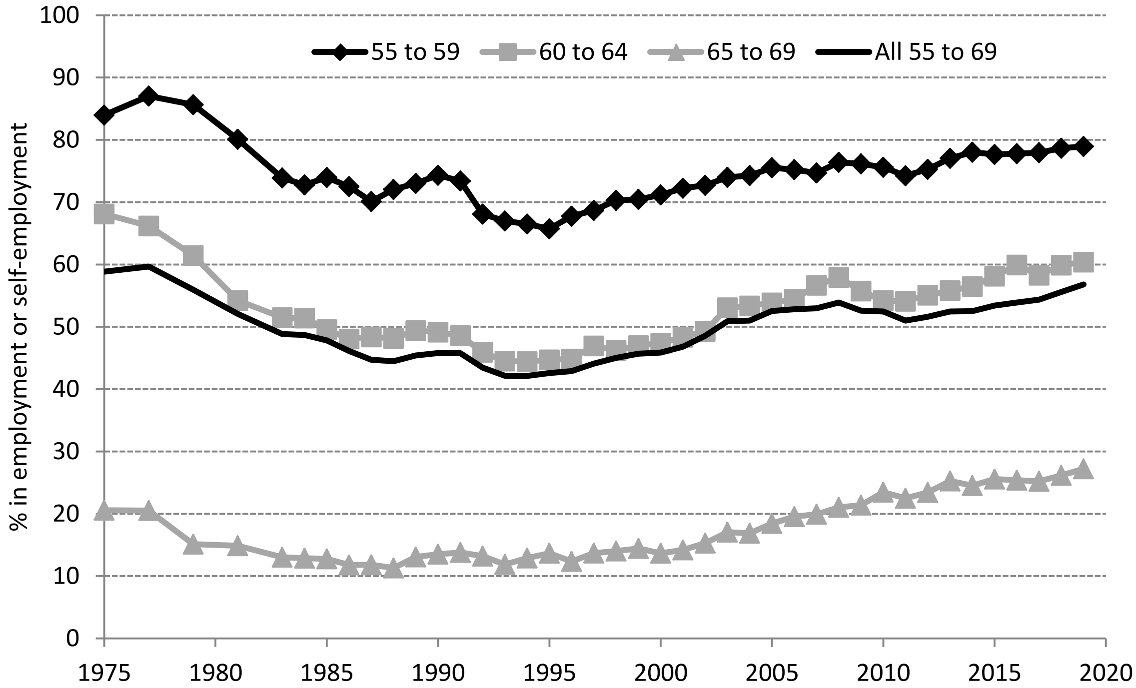

Source: Labour Force Survey.

Figure 2.2 Employment rates of women, by age band, 1975–2019

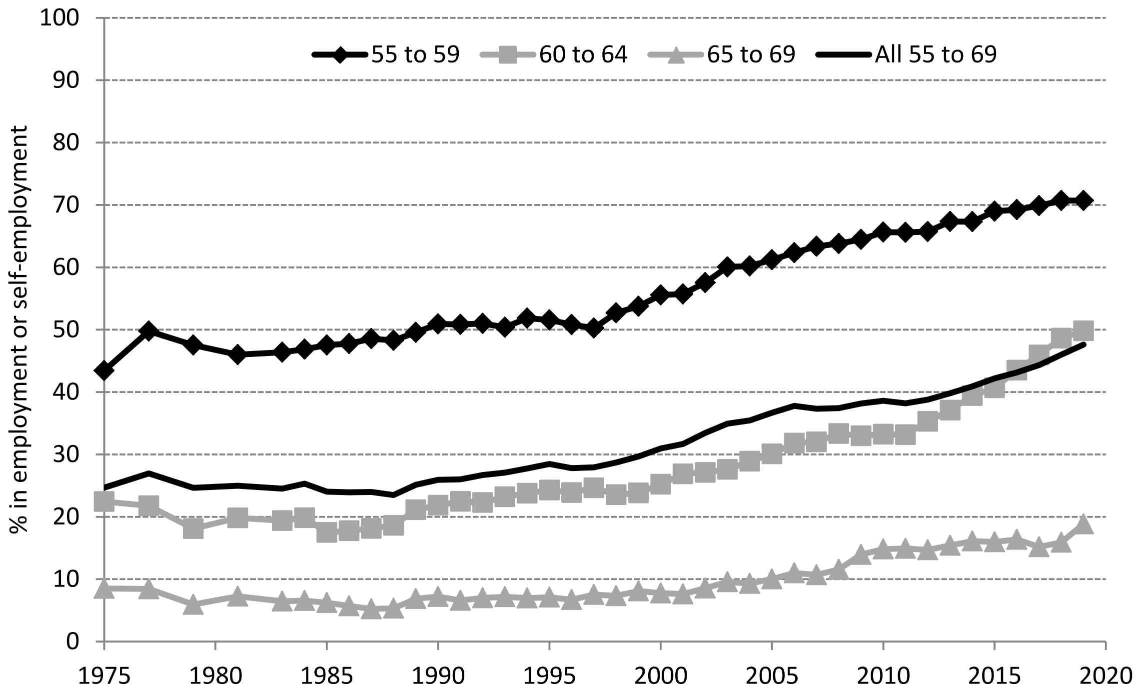

Source: Labour Force Survey.

We next document the relationship between work and pension receipt at different ages. Figure 2.3 shows the distribution of individuals into four possible states in terms whether or not they are working and whether or not they are in receipt of any pension income (state or private). Behaviour in relation to state pension claiming is fairly straightforward and uniform. Individuals cannot begin to draw their state pension before their SPA and almost all individuals begin to draw at this age. This explains the sharp drops in the proportion receiving no pension at ages 60 and 65 in 2002–03 (and shown more clearly in Appendix Figure A.1 which splits the figure by sex). While it is possible to defer receipt of the state pension in return for an actuarial adjustment, less than 1% of individuals in our sample choose to do this. The reasons for the near-universal take-up of the state pension at individuals' SPA are not well understood; individuals must actively claim their state pension, meaning that inertia is not a good explanation, but there may be a lack of awareness of the opportunity and benefits of deferral, or it could be that receiving the state pension at the SPA is engrained as a social norm.

Figure 2.3 Labour force participation and pension receipt by age, in 2002–03 and 2018–19

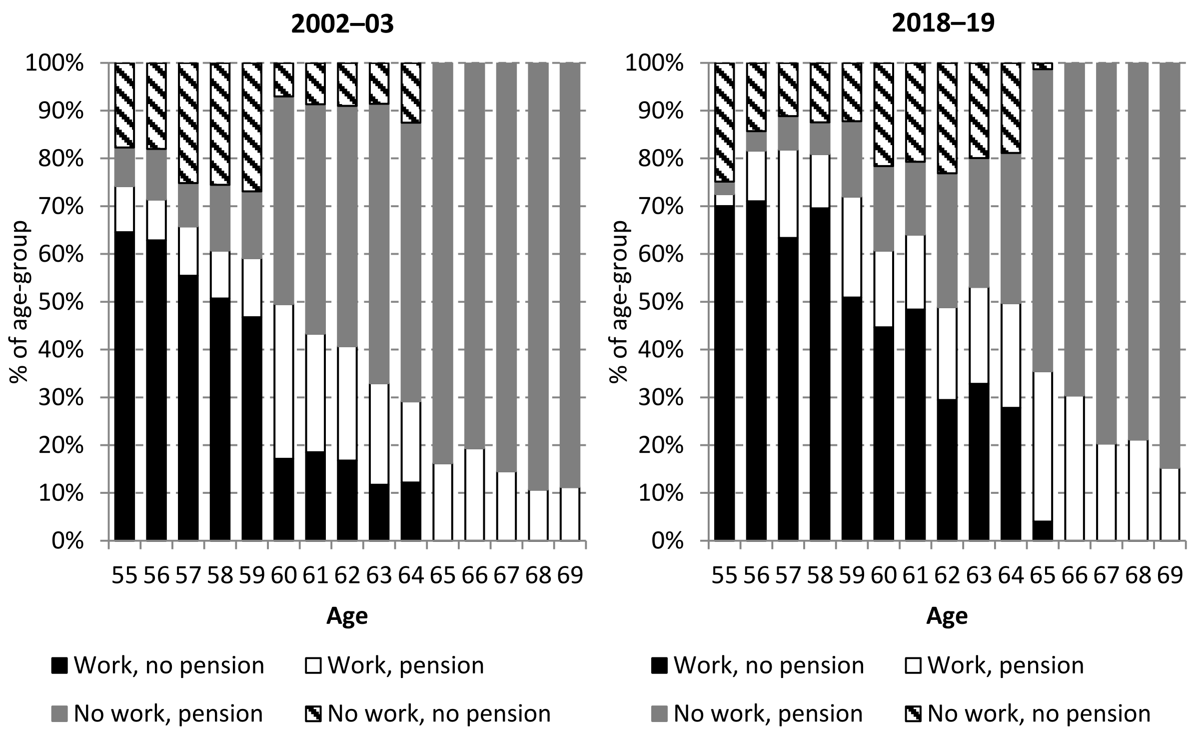

Source: ELSA, waves 1 and 9.

Note: "Work" is defined as being in work or temporarily away from work, in the last month. "Pension" is defined as being in receipt of private or state pension income.

Focusing on private pensions, Figure 2.4 shows the breakdown of working individuals aged 55-74 into those who have no private pension coverage, those who are in receipt of private pension income, and those who have a private pension but are not yet receiving any income from it (and hence they are either contributing to the pension or the pension is being 'retained'). For the group not yet receiving income, we can split individuals into those whose private pension(s) are DB, those with DC and those with both. Across our data period, around 80% to 90% of working men and around 70% of working women have some private pension entitlement. Around 30% of working men and 25% of working women are in receipt of private pension income. Amongst those with pensions that are retained or in contribution, DB pensions are more common for working women than working men.

Figure 2.4 Private pension membership over time, for men and women aged 55–74 in work

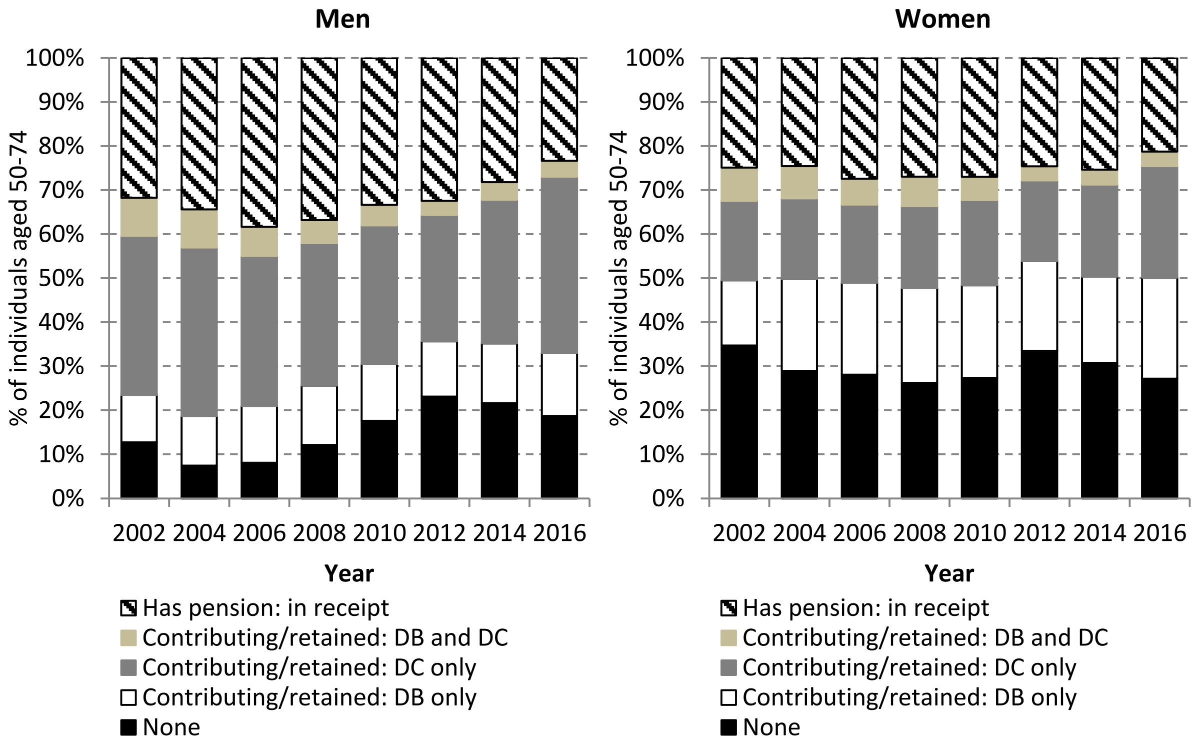

Source: ELSA, waves 1-8.

2.2 Institutional Changes and Pension Reforms

In this section we give a brief overview of the main reforms that affected financial incentives to remain in paid work over the period 2002 to 2019. Incentives to leave work at a particular age changed over time during this period, both due to reforms enacted during the data period but also because earlier reforms can mean that different cohorts faced different rules over their lifetimes. For a fuller discussion of the UK pension system over the period from 1948 to 2018 see Banks and Emmerson (2021). The main reforms that took place within the data period are summarised in Table 2.1.

The broad direction of reforms and their implications can be summarised as follows. The Basic State Pension (BSP), first legislated for in 1948, was payable to men (from age 65) and to women (from age 60) who had worked and paid National Insurance Contributions (NICs) for a sufficient number of years. An individual's rate of BSP was dependent upon their number of years of contributions but not the level of contribution. The `full' BSP was set at a level designed to prevent poverty in retirement, under the expectation that private saving would provide any additional income. Prompted by concerns about the half of workers not offered a workplace pension, from 1978 individuals could accrue entitlement to an additional earnings-related part of the state pension through the State Earnings-Related Pension Scheme (SERPS). As the cost implications of this scheme became clearer, it was made much less generous through reforms legislated in 1986 and 1995.

In the 2000s, the focus of policy reform became ensuring that those with lower earnings or periods out of employment were provided with adequate retirement income. In 2000, the government legislated to replace SERPS from 2002 with the State Second Pension (S2P), which was more generous towards lower earners and to those with caring responsibilities or a long-term illness or disability. In addition to the reforms described in the table below, one of the main sources of variation in retirement incentives attributable to policy during our data period comes from the changes to earnings-related pension schemes and their differential effects on different cohorts.

As set out in Table 2.1, further changes were made to the state pension system during this period which had the effect of broadening out entitlement to the basic state pension and moving the state pension system back towards a 'flat-rate' system. A major policy change, in terms of the generosity of the state pension system was the introduction of the 'triple lock' method of indexation for the BSP in 2011. The 'triple lock' method increases the value of the BSP each April by the highest of nominal earnings growth, consumer price inflation, and 2½%.[2]

Alongside changes to the rates at which parts of the state pension were accrued and paid, the SPA – the earliest age from which individuals can, and in practice do, begin to claim their state pension – has changed. In 1995, it was legislated that the female SPA would gradually rise from 60 to 65 between 2010 and 2020, such that it would become equalised with the male SPA. While equalisation was required by ECHR law, concerns about the sustainability of state pension led to legislation for gradual rises in the (equalised) SPA to 68, and, subsequently, the acceleration of the timetable of increases. In 2011, the scheduled equalisation of the male and female SPA was accelerated such that it was completed by 2018 and the rise in the equalised SPA to 66 was brought forward to take place between 2018 and 2020.

The broad context for UK private pensions is the shift from DB to DC schemes and increasing individual flexibility and autonomy in pension saving. Two particular reforms in this direction during the period 2002 to 2019 were the ending of the requirement that an individual leave their current employer in order to begin drawing the defined benefit pension associated with it and the ending (in two steps) of the requirement that individuals with DC pensions use these to purchase an annuity.

[2] The UK's Office for Budget Responsibility projected that this will have the effect of increasing the value of the BSP by a rate that averages 0.23 percentage points higher than real earnings growth in future years.

Table 2.1: Notable reforms to state and private pension affecting retirement and pension drawing incentives in the period 2002 to 2019

Year	State pensions	Private pensions
2000	• State Second Pension reformed to give greater credit to caring and become more generous to lower earners from 2002	
2002		
2003		
2004		
2005		
2006		• From 2007, no longer compulsory to leave job in order to claim DB pension
2007	• SPA to rise from 65 to 68 between 2024 and 2046 • State second pension will become flat-rate in long-run i.e. phasing out of earnings-related pension • Basic state pension earnings indexed from April 2012 • Number of years required for full BSP reduced from 44 (men) and 39 (women) to 30 • BSP credit given to carers	
2008		
2009		
2010		• Annuitisation of DC fund no longer 'compulsory' if >£20,000 guaranteed income
2011	• Female SPA rise to 65 brought forward to 2018 • Male and female SPA to rise from 65 to 66 over 2018 to 2020 • "Triple lock" indexation of BSP	• End of 'contracting out' of S2P for DC pension schemes
2012		
2013		
2014	• "Single Tier" state pension: No new accrual of earnings-related state pension from 2016 • 35 years required for full state pension	• End of 'contracting out' for all pension schemes from 2016 • End of 'compulsory' annuitisation from 2015
2015		
2016		
2017		
2018		
2019		

3. Construction of individual-level incentives

We conduct our analysis using ELSA, a survey conducted with a representative sample of the English non-institutionalised household population aged 55 and over. There are currently nine waves of this biennial panel available, covering the years 2002/03 to 2018/19, with around 8,000-10,000 observations in each wave. The main survey waves include information on individual earnings and employment as well as detailed information about private pension income and wealth and state pension income, when in receipt. In addition to the nine waves of the main survey, we also exploit an additional 'life history' questionnaire, which asks individuals further questions about their earlier life, including which years they were in employment.[3]

We use the ELSA data to construct individual-level measures of the financial incentive to stay on in work over the coming year, for each of the individuals in our sample. This section describes our method for constructing these incentive variables and validates our method. The calculation of some elements of our incentive variables requires information about individuals' full histories of employment, their earnings, and time spent out of work and caring for a child. We first outline our approach to constructing these elements before going on to describe the way in which incentives are defined and calculated.

3.1 Constructing employment, earnings and carer status histories

As described in section 2.2, individuals can accrue entitlement to the basic state pension and single tier pension by spending time carrying out certain activities that include working and paying National Insurance contributions, but also being in receipt of certain welfare benefits or being the primary carer for a child of a certain age. Individuals could also accrue entitlement to different incarnations of the state earnings-related pension if they had not chosen to 'contract out' of this system. The state pension income that an individual receives upon retirement depends on their full history of employment, earnings, carer status and contracted out status, in a way that is complex and varies across cohorts. In order to determine the additional amount of state pension entitlement that an individual in our ELSA sample may accrue by working for an extra year, we therefore need information on their history along these various dimensions for each year of their adult life.

We construct an employment history using the information in the ELSA life survey for those individuals for whom it is available. The ELSA life survey elicits the dates at which individuals began and finished each spell of employment in their career, enabling us to construct a history of employment status for each financial year. For a full discussion of the method of constructing employment histories from the ELSA life data, see Banks et al. (2020) and Brugiavini et al. (2013). For the individuals who did not respond to the ELSA life survey, we impute an employment history. For men, and for women with no children, we assume that those classified as 'low'-, 'mid'- and 'high'-educated enter employment at ages 16, 18, and 22, respectively, and are continuously employed until the age of their ELSA observation. For women with children, we assume that their employment rate depends on the age of their youngest child. In years when none of her children are aged under

[3] Individuals who took part in the wave 3 ELSA survey were asked to participate in the life history survey. In addition, individuals who refused to take part in the main wave 3 survey, but were in the same household as someone who responded to wave 3, were eligible to take part in the life history. 9,771 individuals took part in wave 3 and 7,855 took part in the life history. 51% of the individuals in our final estimation sample responded to the life history, making up 64% of our observations (due to the fact that those who completed the life history were more likely not subsequently not to attrit from the survey).

18, we assume that a woman is in paid work (so long as she is older than the education-specific entry ages, assumed the same as for men and women who do not have children). In years where one or more of her children are aged under 18, a woman's employment status is randomly imputed with a probability of employment that depends on her educational attainment and the ages of their children. These probabilities are estimated using the sample of women born between 1930 and 1959 and aged 16-45, in the Family Expenditure Survey (subsequently the Expenditure and Food Survey and Living Costs and Food Survey), which measured age, sex, labour market status and number and ages of children contemporaneously over the period 1968 to 2017, and are set out in Table 3.1.

Table 3.1: Assumed employment probabilities for women with young children

	Probability of working	
Education level of mother	Has child under age 5	Has child under age 18 and no child under age 5
Low-education	0.31	0.65
Mid-education	0.41	0.75
High-education	0.53	0.80

Source: Family Expenditure Survey, Expenditure and Food Survey and Living Costs and Food Survey, years 1968-2017. 209,462 observations.

Clearly, our assumed earnings histories for men, and for women without children, who do not have a life history response, are very stylised and represent an upper bound on their number of years spent in employment. For these individuals, we are likely to overstate their accrued state pension entitlements and also to assume away some heterogeneity that arises due to spells out of work.

To construct an earnings history, conditional on these employment trajectories, we use the observed earnings from ELSA where available in the years from 2002 onwards. For other years, we impute a level of earnings in the following way. We first use the FES/EFS/LCFS data to run a median regression of earnings amongst different groups according to sex, birth cohort (measured in 3-year cohort groups), education level and year of observation, according to the following specification, following Blundell et al. (2002):

$$log(y_{i,t}) = \alpha + \beta_1\{education \times cohort \times year\} + \beta_2\{education \times sex \times year\} + \varepsilon_{i,t}$$

where $y_{i,t}$ is the level of individual gross earnings, and each of "education", "cohort", "year" and "sex" represent dummy variables for each value of these variables such that β_1 and β_2 represent a vector of coefficients on each of the possible combinations of the variables they multiply. For each individual in ELSA, we calculate the mean, over the years in which they have positive earnings in ELSA, of the difference between their reported log earnings and their log earnings as predicted by the estimated equation above. Calling this mean 'residual' $\varepsilon_{i,ELSA}$, we then impute earnings for other years of their life in which we assume that they are employed, but do not observe them in ELSA, as:

$$\log(y_{i,t}) = \widehat{\log(y_{i,t})} + \varepsilon_{i,ELSA}$$

This method has the effect of imputing for each individual, in each of their historic years of work, the same percentage of predicted earnings, for someone of their characteristics, as they were observed to have in the ELSA sample, on average. This is equivalent to assuming that, conditional on cohort, education, sex and year, there are no idiosyncratic shocks to earnings.

Finally, being the primary carer for a child contributes to state pension entitlement over some of the period of our analysis. We assume that women who are not employed are the primary carer for their child, if they have one.

3.2 Constructing incentive variables

The incentive variables that we calculate are based on the increase or decrease in pension entitlement that accrues to an individual if they stay in work for an additional year. We calculate for each pension that an individual has, the income that the individual would receive in each future year from that pension, both in the case where they retired immediately and in the case where they retired in one year's time. The approach to calculating the income received varies across the type of pension.

In the case of occupational private pensions, the pension income is dependent on the parameters of the individual's pension scheme. ELSA elicits detailed information on each individual pension held by each sample member, if they are currently contributing to that pension (a 'current' pension). In the case of defined benefit occupational pensions, individuals will typically accrue an additional fraction of entitlement (often 1/80th or 1/60th) of their final salary, for each year that they work in the job with which the pension is associated, up to some maximum number of years. Individuals may face a percentage reduction in their pension income (typically 4% per year) if they begin drawing the pension before the pension's NRA, which is often either 60 or 65 with no bonus for starting to draw the pension at an age above the NRA. Pension income, when being drawn, will usually increase each year in line with price inflation. All of these parameters are elicited for each pension to which an individual is currently contributing during the ELSA interview, allowing the calculation of the pension income that would be received if the individual began drawing their pension and if they worked for an additional year before drawing.[4]

In the case of current defined contribution pensions, pension income depends on the current fund size and rate of contributions (which are reported in ELSA) as well as the return on investments and the cost of an annuity (for which we use market data).[5]

As described in section 2, individuals can, and in some cases do, begin to draw on a private pension while still working. Individuals who are in work and have already begun to draw from one of their private pensions do not accrue additional pension entitlement (for that particular pension) regardless of whether they choose to work an additional year, and so pensions in payment do not contribute to their work incentives through the channel of accrual, though they will form a component of their pension wealth.

Some individuals hold private pensions for which they are no longer making contributions (often because the individual has left the associated job) but have not yet begun to draw pension income.

[4] In some cases (where the individual does not know), some of these parameters are imputed using a conditional hot-deck imputation method.

[5] We assume that individuals will buy a 'level' (i.e. non price indexed), single-life annuity, as was the norm over the data period. The annuity rate we use is specific to the individual's age (and also their sex prior to 2010; since 2011 it is illegal to price annuities differentially based on sex).

In cases where these pensions were observed as current in previous waves of ELSA, we are able to construct the income that would be received (according to the rules of the scheme, and individual tenure and final salary etc.) in the case where the individual begins to draw the pension income immediately and when they choose to draw it in a year's time. If the pension was never observed as current, then more limited information is obtained in the survey and further assumptions must be made. In constructing the income that will be received from each private pension, we follow the methodology in Banks et al. (2005), which can be consulted for further details of our approach. But given that these pensions cannot affect rates of accrual, and as they have been found typically to be a relatively small component of pension wealth, our estimates should be robust to the assumptions made in valuing these pensions.

In the case of state pensions, an individual's entitlement depends on the system of rules that was in place in each year of their life and their history of earnings and employment, caring status, and contracted out status.[6] We construct a pension calculator which calculates, for each individual, their accrued-to-date entitlement to the basic state pension and earnings related state pensions (SERPS and the State Second Pension), under the assumption that the individual will begin to claim their pension at their SPA. We also construct their entitlement if they continue to work for an additional year. This calculation is made under the state pension rules that were in place at the start of the year, on the basis that incentives depend on individuals' *expected* returns from working, and that individuals do not anticipate any reforms or future changes that were not already announced prior to the beginning of the year.

For those individuals who have reached their SPA, we are able to compare the state pension income which they report receiving in ELSA to the income that our calculator predicts they will receive in that year. Figure 3.1 shows the distribution of reported and predicted state pension income for individuals over their SPA. There is a greater concentration of the predicted values at the level of the full basic state pension than is seen in the reported state pension income. There is a mass of individuals reporting zero state pension income (15% of the sample over their SPA) that is not replicated in our predictions. This is because these individuals must not have been in paid work, nor taking part in any other creditable activity (which includes caring for a child and receiving welfare benefits) for much of the relevant period. We are likely to understate the proportion of individuals in this situation due to the fact, if we do not have life history data telling us that an individual is not employed then we assume that they are either employed or caring for a child. There does not seem to be any strong association between reporting zero state pension income characteristics that might be expected to be associated with extending periods without creditable activity in the UK, such as sex, education level, or being born outside the UK. Rather, this group is made up of a range of different types of people.

[6] In reality, an individual's entitlement also depends on their history of receipt of certain benefits. As we assume that all individuals are either in work or caring for a child, we do not need to calculate entitlements stemming from benefit receipt.

Figure 3.1 Histogram of predicted and reported state pension incomes for individuals over state pension age

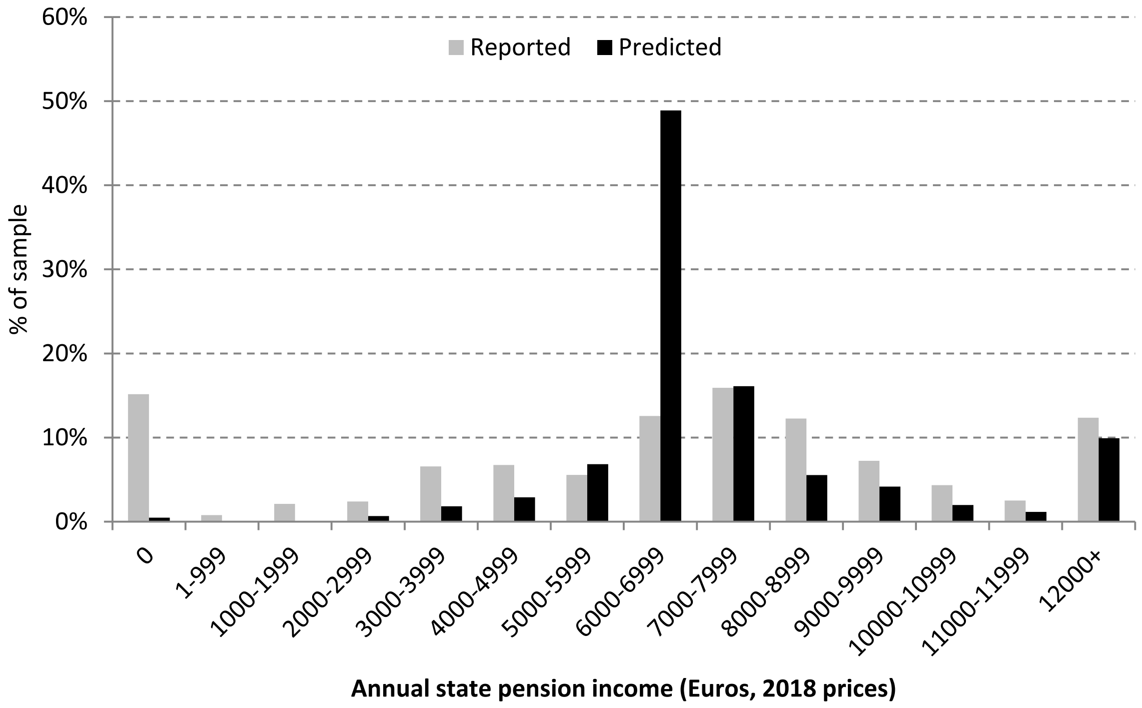

Source: Authors' calculations. ELSA, waves 1-9. Individuals with an imputed level of pension income are excluded from the sample.

Figure 3.2 illustrates the distributions of differences between pension income as predicted by our state pension calculator and as reported in ELSA. We see again the mass of individuals who report zero state pension income in ELSA. The mean percentage difference between reported and predicted state pension income is 8% and median difference is 5% (i.e. our pension calculator overestimates pension income, as compared to the ELSA report, on average).

Figure 3.2 Histogram of percentage difference between reported and predicted state pension incomes for individuals over state pension age

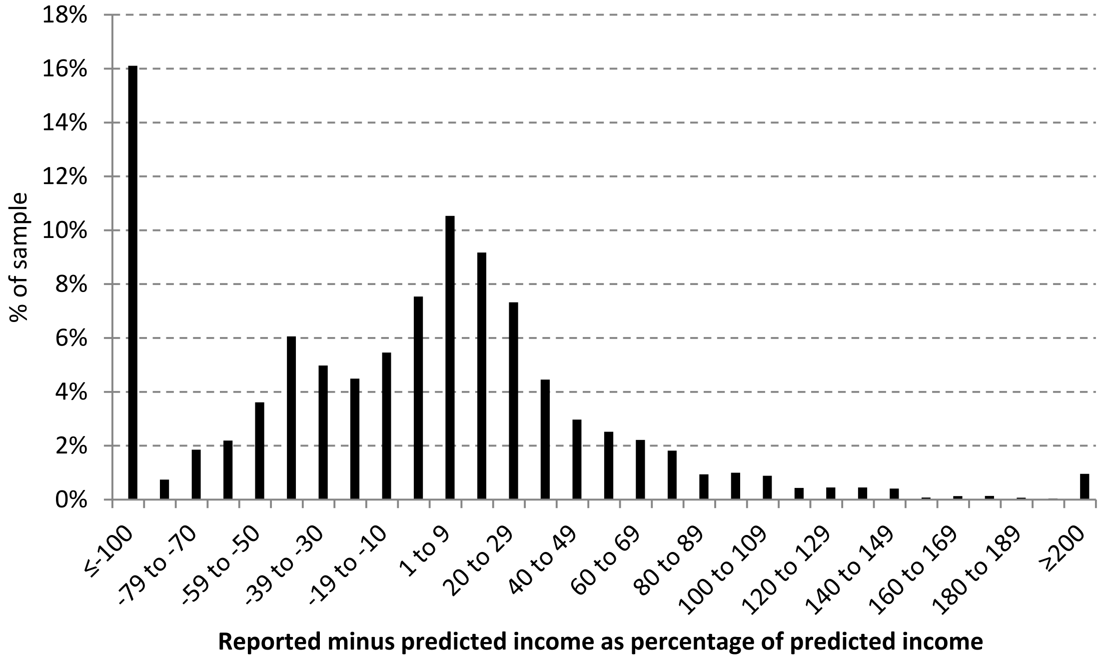

Source: Authors' calculations. ELSA, waves 1-9. Individuals with an imputed level of pension income are excluded from the sample.

By summing together an individual's income from all of their pensions, we obtain measures of total pension income that would be received in each future year both in the case where an individual retires immediately and where they work for one additional year before retiring.

We now set out how we use the constructed measures of pension entitlement to create our incentive variables. We first calculate a measure of pension wealth, defined as the expected discounted sum of future pension income, for both the case of immediate retirement and retirement in one year's time. In calculating this, we discount future income at a rate of 3% per year and take expectations using sex-and-cohort-specific survival curves.[7] We define *gross accrual* as the change in pension wealth that results from an extra year's work. *Net accrual* is defined by subtracting employer and employee National Insurance contributions (NICs) from gross accrual. National Insurance contributions are paid at different rates on different bands of earnings. An individual paid a reduced rate of NICs if they 'contracted out' of the SERPS/S2P (meaning that they did not build up any entitlements under these schemes) in the years where this was permitted.[8] We do not have a direct measure of whether individuals were contracted out, but we assume that individuals enrolled in private pension schemes were contracted out, as was the norm. The *implicit tax* rate (ITAX) is then defined as the negative of net accrual divided by net individual earnings (where earnings are made net of direct taxes i.e. NICs and Income Tax).[9]

[7] We use the UK Office for National Statistics 2014-based cohort survival curves for England and Wales.
[8] Contracting out was possible during the period 1978 until 2016.
[9] We windsorise ITAX for values less than -100% and over 100%.

It is possible for an individual to experience negative gross accrual on some private pensions. This will happen for an individual who has passed their NRA and has already built up a full tenure in their final salary pension scheme but does not begin to draw their pension. By delaying receipt of their pension by a year, such an individual forgoes one year of income but does not receive any higher pension income when they begin to draw. Before 2006, an individual had to leave their job in order to begin to draw the defined benefit pension associated with it. For these years, we deem the work and pension receipt decision to indeed be joint (we implicitly assume that individuals cannot move employer in order to begin to draw their pension while continuing to work) and so negative accrual as a result of an extra year of work is possible. From 2006 onwards, there was no such requirement. We therefore assume that if an individual will enter into negative accrual as a result of staying on in work and not drawing their pension, that they would begin to draw this pension even if they continued to work, and hence have zero accrual for that pension.

4. Variation in individual incentives and employment outcomes

This section sets out the key descriptive trends in our retirement incentives and in labour market exit rates. As in previous studies of work incentives using 'representative' individuals (see Banks and Emmerson (2021) for the UK), we exploit the fact that the reforms to the state and private pension systems have led to different cohorts facing different incentives to stay on in work and that the different rules faced by men and women result in variation in incentives across sexes. This paper goes further by constructing our incentive measures for a panel of microdata and therefore taking advantage of the individual-level variation in incentives that results from heterogeneous characteristics and idiosyncratic histories of earnings and private pension provision. This yields a wealth of additional variation in incentives within demographic groups, which will be exploited to identify the effect of incentives on exit from work.

Figure 4.1 and Figure 4.2 show the mean levels of ITAX at each age from 55 to 69 for each year of our data from 2002–03 to 2016–17 (which are the years of data for which we observe whether individuals subsequently retire), for men and women, respectively. Before 2010, individuals below their SPA (60 for women and 65 for men) generally face lower implicit tax rates because they have the opportunity to accrue entitlement to the state pension (though this is counterbalanced by the fact that employee NICs are only paid by individuals under their SPA). In 2010, we see the mean level of ITAX fall for women aged 59 and 60 because the female SPA began to rise in 2010 and so some of these individuals would accrue additional state pension entitlement if they remained in work.

Another prominent trend is the fall in mean ITAX for men aged 65 and over, from 2004 to 2006. This is largely explained by a small proportion of men who have a DB pension, have accrued a full tenure and are over the NRA for their scheme. Under the system before 2006, these men had to forgo pension income if working but, from 2007, were able to draw their pension while continuing to work. This is seen to a lesser degree for women, too.

Figure 4.1 Mean ITAX level by age and year for men

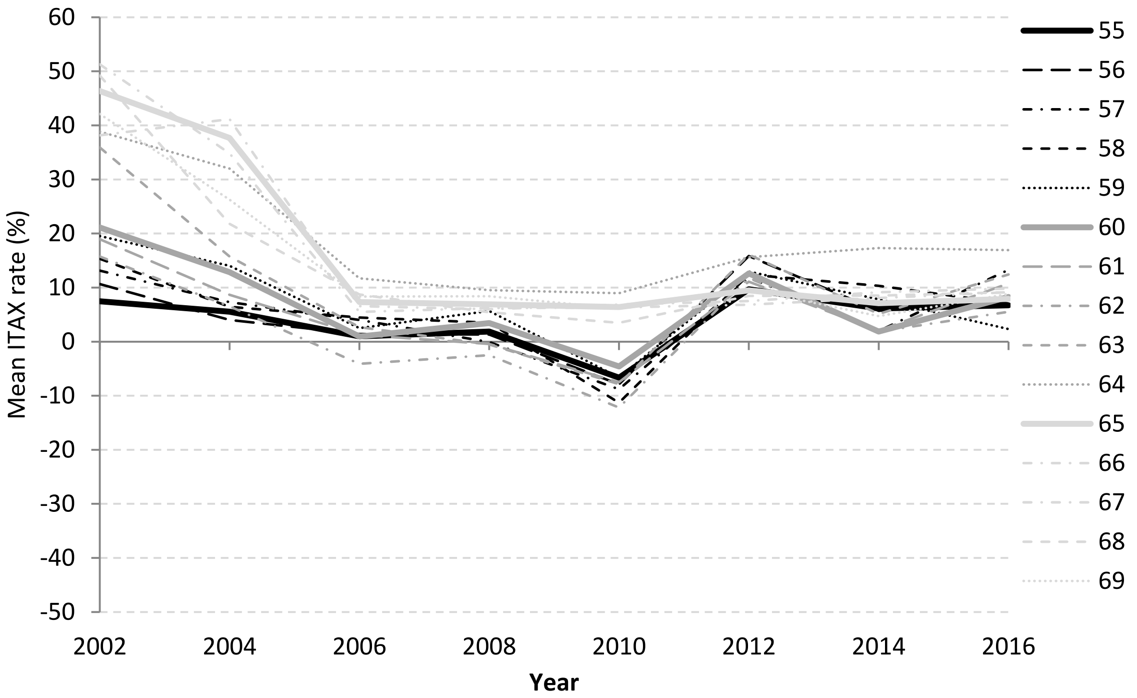

Source: Authors' calculations. ELSA waves 1-8.

Note: Sample includes individuals in work.

Figure 4.2 Mean ITAX level by age and year for women

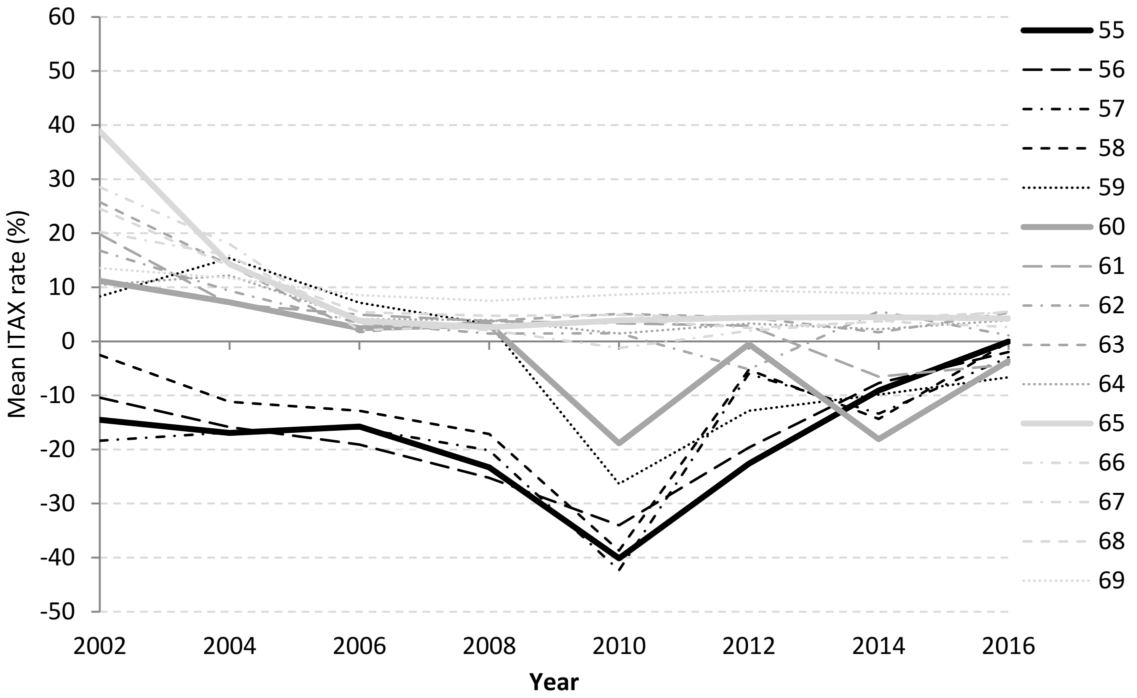

Source: Authors' calculations. ELSA waves 1-8.

Note: Sample includes individuals in work.

Figure 4.3 illustrates the degree of cross-sectional and inter-temporal variation in ITAX that exists even within sex and age-groups by showing selected percentiles of the ITAX distribution. There is more variation at younger ages, as there is greater potential for these individuals to be accruing additional pension entitlements. In particular, we see that a significant proportion of women aged 55 receive a substantial implicit subsidy to their earnings as they are more likely to have spent a number of years out of the labour market. To the extent that some of this variation reflects common shifts in incentives over time, this is not variation that will be exploited to identify incentive effects, once time trends are controlled for.

Figure 4.3 Selected percentiles of ITAX for selected ages, for men and women

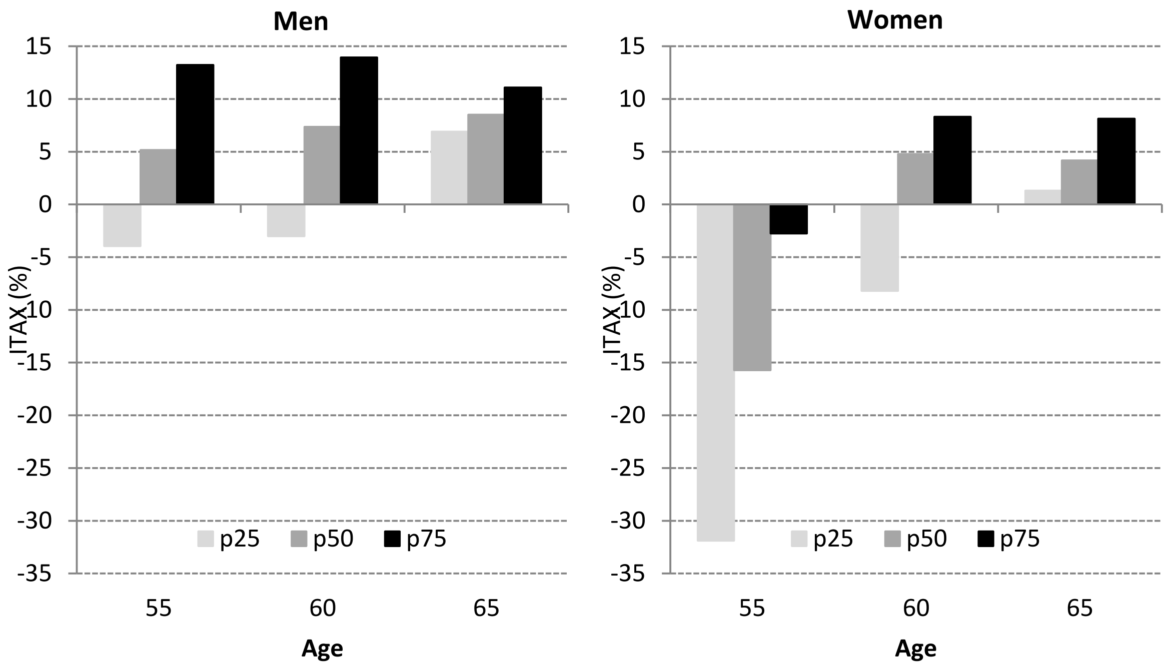

Source: Authors' calculations. ELSA waves 1-8.

Note: Sample includes individuals in work.

We now turn to examine our outcome variable, one-year-ahead labour force exit. We define labour force exit in the following way. Individuals who are in work in the subsequent wave of ELSA are defined as not exiting. Individuals who are not in work in the subsequent wave of ELSA, and whose reported date of leaving their last job was within 12 months of their prior ELSA interview date are defined as having exited. Figure 4.4 shows one-year-ahead labour force exit rates by age and sex for the first and final year of our panel. Unsurprisingly, exit rates rise with age. We can see that in 2002–03, female exit rates rose sharply at age 59. In 2016–17, a similar pattern is seen, but at age 63. This appears to be consistent with the increase in the female SPA from 60 to 63 between these two data periods.

Figure 4.4 One-year-ahead labour force exit rate by age, in 2002–03 and 2016–17

Source: ELSA, waves 1-9.

Note: An individual is defined as exiting the labour force if they are no longer in work at the time of their next ELSA interview and report having finished their last job within one calendar year of their previous interview.

In general, the labour force exit rates appear to be more often lower than higher in the final data period, as compared to the first. Quantifying the overall difference between exit rates in 2002–03 and 2016–17, and their implications for employment is not easily achieved only by examining Figure 4.4. Clearly changes in aggregate exit rates at younger ages are more important for overall participation levels as more people are in work at these ages. In order to summarise the information contained in the labour exit rates in each year of our data in a useful way, we calculate the expected retirement age for a man and for a woman, aged 59 who experienced the cross-sectional mean labour force exit rates for that year at each subsequent age. We assume that any individual in work at age 75 does not work past age 75. Figure 4.5 plots these calculated 'expected retirement ages' for each wave of our data, for both men and women. The male expected retirement age rises from 65.1 to 67.7, an increase of 2.5 years. The female expected retirement age rises from 64.1 to 65.7, an increase of 1.6 years. This measure captures the trend of longer working at these ages. Women in particular see a large increase from 2010–11 to 2016–17, most likely due to the SPA rises that took place over those years. This measure shows the extent of change over time but of course does not capture changes in the proportion of individuals working at age 59 (and the consequences for employment at older ages), which is also an important part of the picture when examining aggregate employment trends.

Figure 4.5 Implied mean retirement age by year, for men and women who were in paid work at age 59

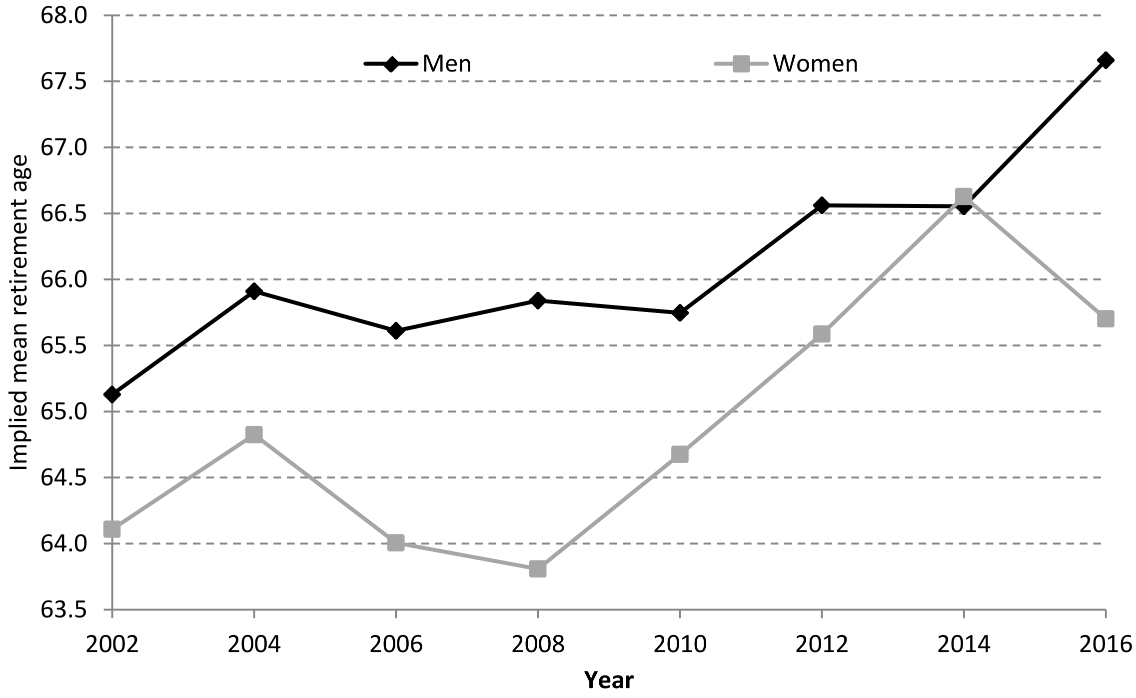

Source: ELSA waves 1-9.

Note: Implied mean retirement ages for a given sex and year are calculated by assuming that an individual faces the one-year ahead retirement hazard rate at each age as measured in ELSA in that year for someone of the given sex. We assume that anyone not retired before age 75 retires at that age with probability 1.

5. Regression Analysis

In this section, we test whether retirement incentives can explain individual decisions to leave or remain in work. Using our constructed incentive measures and our measure of labour force exit, for the individuals in work in our ELSA sample, we conduct regression analysis of the association between incentives and labour force exit. We are interested in the effect on labour force exit of the implicit tax rate but also of individual wealth (composed of both the accrued pension wealth that we calculate and of other net financial wealth). Standard theory leads us to expect both implicit tax and wealth effects.

5.1 Methodology

The basic empirical specification that we test is set out in Eq. (1). Our dependent variable is a binary variable for one-year-ahead labour force exit. Our independent variables of interest are the implicit tax rate ("ITAX") and individual wealth in hundreds of thousands of euros in 2018 prices ("W"). Individual wealth is the sum of accrued pension wealth (both state and private) and the individual's level of net financial wealth (for those in a couple, net financial wealth is the joint financial wealth of the couple). The vector X contains controls for sex, for self-reported health (using dummies for 3 possible categories of report), for occupation type (dummies for 3 categories), for lifetime earnings, and includes a dummy for whether the individual responded to the ELSA Life survey.[10]

$$Exit_{i,t} = \alpha_1 ITAX_{i,t} + \alpha_2 W_{i,t} + X_{i,t}'\beta + \varepsilon_{i,t} \quad (1)$$

We first assume a linear probability model and estimate the above equation using OLS. Secondly, we estimate a probit model, giving the above equation the latent variable interpretation and assuming normality of the individual error term, $\varepsilon_{i,t}$. The key identifying assumption in both models is that the individual error term is uncorrelated with the covariates. Here, this means that any unobserved factors that might affect an individual's propensity to leave work – such as their taste for work or leisure – are uncorrelated with our incentive variables, conditional on the level of the other covariates. We interpret the controls for self-reported health as picking up any work-limiting health conditions, or the effect of health on the disutility from working. Similarly, the controls for occupation type will pick up some aspects of the working environment which may be relevant for the disutility from work and hence decision of whether to leave work. The control for sex is intended to capture differences in labour market attachment or norms around working between sexes which could confound our incentive variables. Conditional on these covariates, there will be variation in our incentive measures induced by policy cut-offs and reforms and variation in the rules of occupational pensions across employers, as set out in section 3 and section 4. This exogenous variation in incentives allows us to estimate their effect.

In additional specifications we introduce, in various combinations, controls for age (either a quadratic function of age or single-year dummies) interacted with sex, and for cohort (linear term in year of birth) also interacted with sex. We also add a set of single year dummies. While age *per se* may not be expected to affect the decision to leave work, these controls might pick up residual variation in the disutility of work driven by other unobserved factors related to age, such as social norms around retirement timing or age-specific eligibility for other programmes or benefits. Time controls pick up effects of aggregate macroeconomic and other shocks that are relevant to individual choices beyond their impacts on earnings and wealth. The sex-specific cohort controls are intended to capture changes in labour market attachment, or any other changes that occur smoothly across cohorts that might confound the incentive variables.

In the further specifications that include age controls, we also include a dummy variable which takes the value 1 if an individual will reach their SPA in the year following their interview. This picks up the additional 'signalling' effect of the SPA threshold beyond its impact on financial incentives. We include this variable only when controlling for age so that it is not confounded by any changes in exit rates that are driven by unobserved factors correlated with age. In the models where we control for the single year of age the effect of the SPA threshold is separately identified from age because the

[10] Our occupational classification is based on the UK National Statistics Socio-economic classification of the individual's current job. We control for 3 dummy variables corresponding to whether an individual is in a managerial/professional job (NS-SEC codes 1-6), an intermediate job i.e. clerical/administrative/employer in small organisation (NS-SEC codes 7-9) or a routine or semi-routine job (NS-SEC codes 10-13). Lifetime earnings is the sum of (imputed) real terms earnings in each year from age 16 to the year of observation.

female SPA was changing over our data period (while where we control for age with a quadratic it is also identified by the parametric restriction of the age effects).[11]

It is possible that, conditional on the level of covariates, there is unobserved variation across individuals in the propensity to stay in, or leave, work. If this unobserved heterogeneity is uncorrelated with our covariates then our estimates are unbiased, but they are inefficient.[12] We therefore show results for both the linear probability and probit models allowing for random effects, which is the efficient estimation under the assumption of strictly exogenous individual effects.

Identification would not be achieved by the models set out above if there were unobserved heterogeneity in the propensity to stay in, or leave, work that were correlated with our incentive variables, conditional on other covariates. This could happen if, for example, conditional on health and occupation type, some individuals liked to work more than others and so were more likely both to have accrued a full state pension entitlement (and so have high wealth and high ITAX) and to stay on in work at older ages. To address this concern, we estimate our linear probability model including individual-specific fixed effects (using the within estimator). Here, we exploit the within-individual variation in incentives that results from variation in individual earnings over time and its impact on pension accrual, from individuals reaching certain thresholds or discontinuities in the state and private pension systems, and from policy reforms. We cannot estimate our probit model with fixed effects due to the incidental parameters bias that results from estimating a non-linear model with fixed effects on a relatively short panel.

In all specifications, our sample is composed of all observations of individuals aged between 50 and 74 who are in work and for whom we observe of their earnings and wealth and for whom we observe at least one subsequent ELSA interview (necessary to determine whether they subsequently exited work). There are 29,863 person-year observations in ELSA who are in the relevant age-group and in work. Of these, 24,318 observations have a subsequent ELSA interview and, of these, 20,642 have wealth and earnings information. This estimation sample of 20,642 observations is composed of 7,354 unique individuals who are observed on average 3 times.

5.2 Results

We present our results for the linear probability model specifications in Table 5.1 and the probit specifications in Table 5.2. Both tables are split into sections corresponding to specifications without controls for individual unobserved heterogeneity, specifications including random effects, and (for the linear probability model only) specifications including individual fixed effects. Standard errors (heteroscedasticity-robust and clustered at the individual level) are shown in grey below the estimated coefficients. An additional set of results with an alternative specification of control variables is presented in the Appendix.

[11] For further discussion of the identification of the SPA effect using the change to the female SPA see Cribb et al. (2016).

[12] While these estimates are inefficient, our standard errors are heteroscedasticity and cluster robust and hence correct.

Table 5.1. Linear probability regression results: effects on one-year-ahead exit

	Implicit Tax	Wealth (€,000k)	Reach SPA next year	Age controls	Time controls	Cohort controls
A: No controls for unobserved heterogeneity						
(1)	0.094***	0.001**	-	None	None	None
	0.008	<0.001	-			
(2)	0.009	0.001*	0.068***	Quadratic	None	Linear
	0.008	<0.001	0.012			
(3)	0.011	0.001*	0.068***	Quadratic	Dummies	Linear
	0.008	<0.001	0.012			
(4)	0.012*	0.001*	0.075***	Dummies	Dummies	None
	0.008	<0.001	0.019			
B: Random effects						
	Implicit Tax	Wealth (€,000k)	Reach SPA next year	Age controls	Time controls	Cohort controls
(5)	0.090***	0.001**	-	None	None	None
	0.008	<0.001	-			
(6)	0.015**	0.001*	0.063***	Quadratic	None	Linear
	0.008	<0.001	0.010			
(7)	0.024***	0.001*	0.074***	Quadratic	Dummies	Linear
	0.008	<0.001	0.018			
C: Individual fixed effects						
	Implicit Tax	Wealth (€,000k)	Reach SPA next year	Age controls	Time controls	Cohort controls
(8)	0.054***	0.001	-	None	None	N/A
	0.010	0.001	-			
(9)	0.027***	<0.001	0.052***	Quadratic	None	N/A
	0.010	<0.001	0.011			
(10)	0.045***	0.001	0.050***	Quadratic	Dummies	N/A
	0.010	<0.001	0.011			

Note: *** denotes that the coefficient is significantly different from zero at the 1% level, ** at the 5% level and * at the 10% level. All specifications include controls for self-reported health, occupation type and lifetime earnings. Specifications (1)-(7) also include controls for sex and whether the respondent was part of the "ELSA LIFE" sample. Age and cohort controls are sex-specific. Standard errors are clustered at the individual level and shown in grey. Sample size for all specifications is 22,629 observations.

The coefficient on implicit tax is the effect of a 100 percentage point increase in the implicit tax rate on the probability of exit in the following year. Specification 1 in Table 5.1, which corresponds to Eq. (1) above, therefore implies that a 10 percentage point increase in the implicit tax rate increases the probability of exit by 9.4 percentage points. For comparison, the mean exit rate across all individuals in our sample is 7.5%, so a 10 percentage point increase in implicit tax would increase the exit rate from 7.5% to 8.4%, an increase of 12 per cent. The interpretation of the wealth term is that a 100,000 euros increase in wealth results in a 0.1 percentage points increase in the probability of exit. These effects are of the expected sign (as they are in all specifications).

Specifications 2, 3 and 4 show the results when adding in selected controls for age, time and cohort. Changes across cohorts related to the changing labour force attachment of women, or changing levels of health or expected longevity would be expected to occur in a gradual and smooth manner, hence we choose to control for cohort with a simple linear term. In specification 2, we also add a quadratic term in age. Adding these two controls reduces the estimated effect of implicit tax substantially and it is no longer statistically significant (at even the 10% level). The effect of wealth is unchanged and significant at the 10% level. Specification 3 adds in time controls in single year dummies. In specification 4, we show our most 'flexible' set of controls for age and time effects by controlling for age in single year dummies. In this specification, we omit the cohort control as the collinearity of age, time and cohort mean that there may not be sufficient independent variation to estimate credibly all three coefficients at this level of granularity. Specifications 2, 3 and 4 all show similar results in terms of the size and precision of the estimated effects of implicit tax and wealth, showing robustness to the more flexible controls for age and time. Across specifications 2 to 4, we find a consistent and large effect of crossing the SPA threshold on labour force exit. Individuals who reached their SPA in the year following their interview had around a 7 percentage point higher exit rate in that year, even when controlling for the associated financial incentives.

Turning to section B of the table, and the random effects estimates, we find effects that are broadly similar to those in section A. When controlling flexibly for time effects in specification 7, we find a substantially larger effect of implicit tax (compared to the equivalent specification without random effects) of 0.024 percentage points that is significant at the 1% level.

Turning to the fixed effects estimation in section C, our incentive variables have the same, expected signs, as in the other specifications, but the magnitude of the estimated implicit tax effects are quite different and wealth is no longer found to have a significant effect. In specification 9, which includes a quadratic term in age, the estimated effect of a 1% rise in implicit tax is a 0.027 percentage point rise in the exit rate. This is larger than the effect found in the equivalent specifications without fixed effects. When controls for time are added, the estimated effect of implicit tax is larger still, at 0.045 percentage points and is significant at the 1% level. The estimated SPA effects are only slightly smaller than in the specifications without fixed effects and remain large (at around 5 percentage points) and statistically significant. One interpretation of these results would be that a substantial amount of the association between wealth and labour force exit is in fact driven by those who have low wealth having a larger (unobserved) propensity to stay in work. A positive association between a taste for work and earnings (which we would expect if taste for work were correlated with skill or ability) could drive a positive or a negative correlation between taste for work and wealth. On the one hand, greater lifetime earnings would lead us to expect higher wealth, holding fixed the length of an individual's career. On the other hand, those with a greater taste for work may retire later and hence accumulate their wealth more slowly, leading to a negative correlation between wealth and taste for work. Our fixed effect results are consistent with the latter effect dominating.

The probit model allows us to assess the robustness of our findings in a specification that explicitly allows for the binary nature of our outcome variable (although requires the assumption of normally distributed errors). Table 5.2 sets out the results from seven probit specifications that are equivalent, in terms of variables included, to the first seven linear probability models in Table 5.1. We report estimated average marginal effects (integrated over the distribution of covariates in the sample). The results are quantitatively very similar to the equivalent linear probability results, in terms of the size of the implicit tax and wealth effects. Across all probit specifications, the estimated SPA effect is 4–5 percentage points. This is smaller than in the linear models but still economically and statistically significant.

Table 5.2. Probit regression results: average marginal effects on one-year-ahead exit

A: No controls for unobserved heterogeneity						
	Implicit Tax	Wealth (€,000k)	Reach SPA next year	Age controls	Time controls	Cohort controls
(1)	0.097***	0.001**	-	None	None	None
	0.008	<0.001	-			
(2)	0.014	0.001***	0.043***	Quadratic	None	Linear
	0.009	<0.001	0.006			
(3)	0.015**	0.001***	0.043***	Quadratic	Dummies	Linear
	0.009	<0.001	0.006			
(4)	0.014	0.001***	0.050***	Dummies	Dummies	None
	0.009	<0.001	0.012			
B: Random effects						
	Implicit Tax	Wealth (€,000k)	Reach SPA next year	Age controls	Time controls	Cohort controls
(5)	0.097	0.001**	-	None	None	None
	0.008	<0.001	-			
(6)	0.015*	0.001***	0.046***	Quadratic	None	Linear
	0.009	<0.001	0.006			
(7)	0.018**	0.001***	0.045***	Quadratic	Dummies	Linear
	0.009	<0.001	0.006			

Note: *** denotes that the coefficient is significantly different from zero at the 1% level, ** at the 5% level and * at the 10% level. All specifications include controls for self-reported health, occupation type, lifetime earnings, sex and whether the respondent was part of the "ELSA LIFE" sample. Age and cohort controls are sex-specific. Standard errors are clustered at the individual level and shown in grey. Sample size for all specifications is 22,629 observations.

We can illustrate the quantitative importance of the SPA effect by comparing it to the estimated effects of age. Figure 5.1 plots the estimated coefficients on the dummy variables for each age (for each sex) from probit specification 4, telling us the average 'effect' of being at that age, as compared to age 50, on the probability of labour force exit. We show alongside this the magnitude of the SPA coefficient. The size of the SPA coefficient is comparable to the increase in the probability of labour force exit associated with an increase in age of six years (the figure shows that the likelihood of leaving the labour market increases by 20ppt over the 25 years from age 50 to age 74, this implies a 5ppt increase over 6 years which is equal to the estimated SPA effect). Our estimated SPA effect is very similar, or if anything slightly lower than the estimated SPA effect found by Cribb et al. (2016), who exploit the same reform to the UK female SPA to find that being below the SPA increases female employment by 6.3 percentage points. The fact that we find this magnitude of effect even when controlling for the financial incentives provided by the state pension system and by individuals' private pensions is further evidence that the SPA acts as a signal to retire.

Figure 5.1 Comparison of estimated effects of age and of reaching the State Pension Age on labour force exit

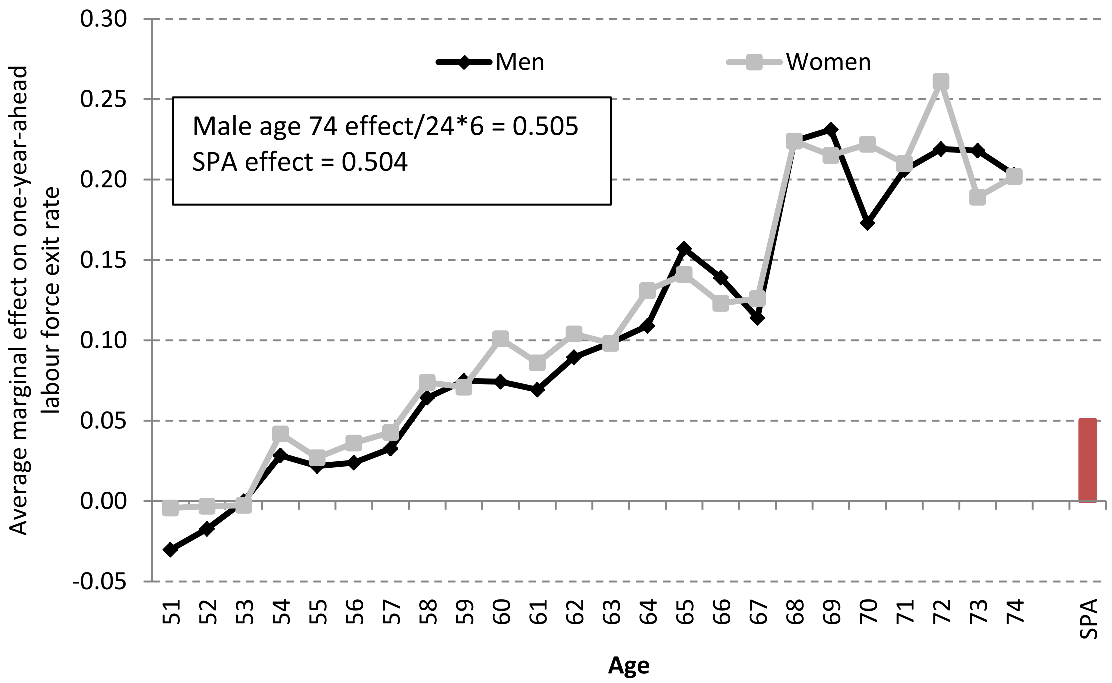

Note: Age effects are the average marginal effects of the dummy variables for each age, interacted with the dummy for sex, from probit regression model 4, where age 50 is omitted. "SPA" is the coefficient on the dummy variable that indicates that the individual will reach their SPA for the first time in the coming year.

Figure 5.2 shows the fit of our model (using coefficients from probit specification 4). We calculate the predicted exit probabilities for each individual and then compare the mean probabilities by age to the actual exit rates for our estimation sample, for both men and women, in the first and final waves of our data. The model captures the main features of the data in 2002–03: the spike in male exits at age 64 and the spike in female exit rates at age 59 in 2002–03 (before the SPA reforms). It also picks up the general increase in exit rates with age, including the sharp increase in the late 60s.

Figure 5.2 Actual and predicted one-year-ahead labour force exit rate by age, for 2002–03 and 2016–17, for men and women

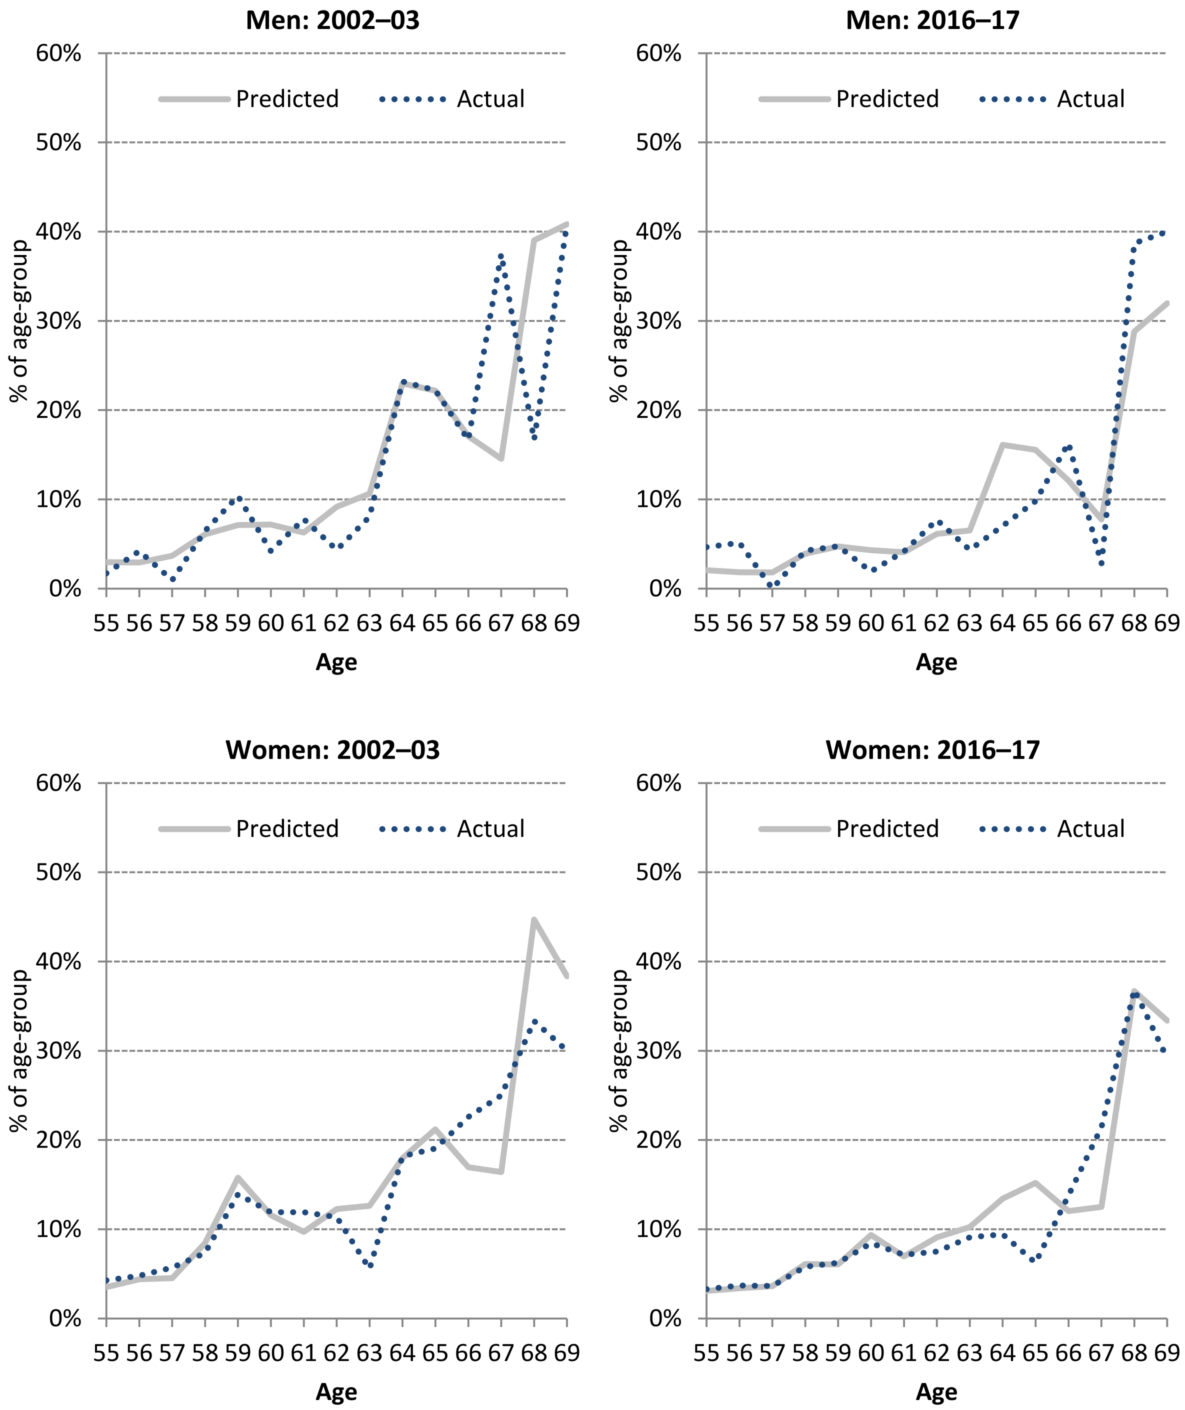

Source: ELSA, waves 1-9.

Note: Predicted exit rates are created using probit regression model 4.

6. Counterfactual analysis

To what extent can reforms to pension systems explain changes in labour force participation over time? The estimated models of the previous section enable us to give an answer to this question. In this section, we conduct a 'counterfactual' analysis where we calculate, for each individual, the incentives that they would have faced under earlier pension systems and the resulting labour force exit rates predicted by our regression model. We can then aggregate these to compare labour force exit rates as predicted by our model under the actual pension system to the labour force exit rates we predict would have obtained in the absence of reforms. We summarise these predicted labour force exit rates as an expected retirement age, in order to interpret the effect of reforms in terms of the additional years of work they induce.

We calculate individuals' pension entitlements and incentives under three pension systems: the actual pension system, the 2002 system – selected as it is the start of our data period – and the 1979 system – which was the 'peak' of the generosity of the SERPS system. The 'actual' system, is the system of incentives that we have already calculated i.e. the incentives that individuals actually faced. Calculating incentives under the 2002 system means that we calculate incentives as they would have been for each individual if there had been no reforms to the pension system after 2002. By this we mean that we calculate each individual's state pension incentives assuming that they accrued entitlements according to the 'actual' state pension rules until 2002 but that for years after 2002, they accrued entitlements according the rules that prevailed in 2002. Individual entitlements and incentives can change in the years after 2002, but this happens only due to changes in, for example, an individual's earnings or an indexation of the state pension that was in place by 2002.

On the private pensions side, rules after 2002 are also kept unchanged as they were under the 2002 system. The effect of this is to 'turn off' the 2006 reform which allowed individuals to begin drawing their DB pension while still working in the related job (the effect of this on incentives is that individuals may face a negative accrual on a DB pension after 2006 if they continue to work). Under the 1979 system, we follow the same principles and calculate state pension entitlements as they would have been if there were no reforms after 1979. Calculating private pensions under a 1979 system is complicated by the dramatic changes in the private pensions landscape over this time period, including the fact that private DC pensions did not exist in 1979. We therefore use the same private pension incentives as in our 2002 counterfactual. The difference between the 1979 and 2002 counterfactuals is therefore solely attributable to state pension reforms.

With calculated incentive measures for the actual pension system and our two counterfactual systems, we can use our regression model's estimated coefficients to calculate each individual's predicted probability of labour force exit under the three different systems. Comparing the mean predicted exit rates under these three systems gives us a measure of the impact of policy reforms on exit rates. Figure 6.1 plots the predicted exit rates by age for the final year of our data under the actual system of incentives as well as the 1979 and 2002 counterfactual systems, using the estimated coefficients from probit model 4. We can see the estimated impact of the reform to the female SPA. This reform was legislated in 1995 and so is included in the actual and 2002 systems but not the 1979 system. Under the 1979 system, predicted one-year-ahead exit rates 'spike' at age 59 due to the bunching of retirement at the female SPA of 60. Under the later systems, the female SPA for those who reached their SPA in the year after their interview in 2016–17 was between 63 and 64. For the 2002 system, we see a 'spike' in predicted one-year-ahead exit rates at ages 62 and 63. The current system features a predicted exit rate at these ages that is higher than the 1979 system but lower than the 2002 system. The 2011 Pension Act introduced a faster rate of increase of the female SPA from 63 to 65 (so the female SPA reached 65 in 2018-19 rather than 2020-21) meaning that fewer women aged 62 and 63 reached their SPA in 2017-18 than would have under the 2002 system.

Figure 6.1 Predicted one-year-ahead labour force exit rate by age, for 2016–17, under 1979, 2002, and actual pension systems

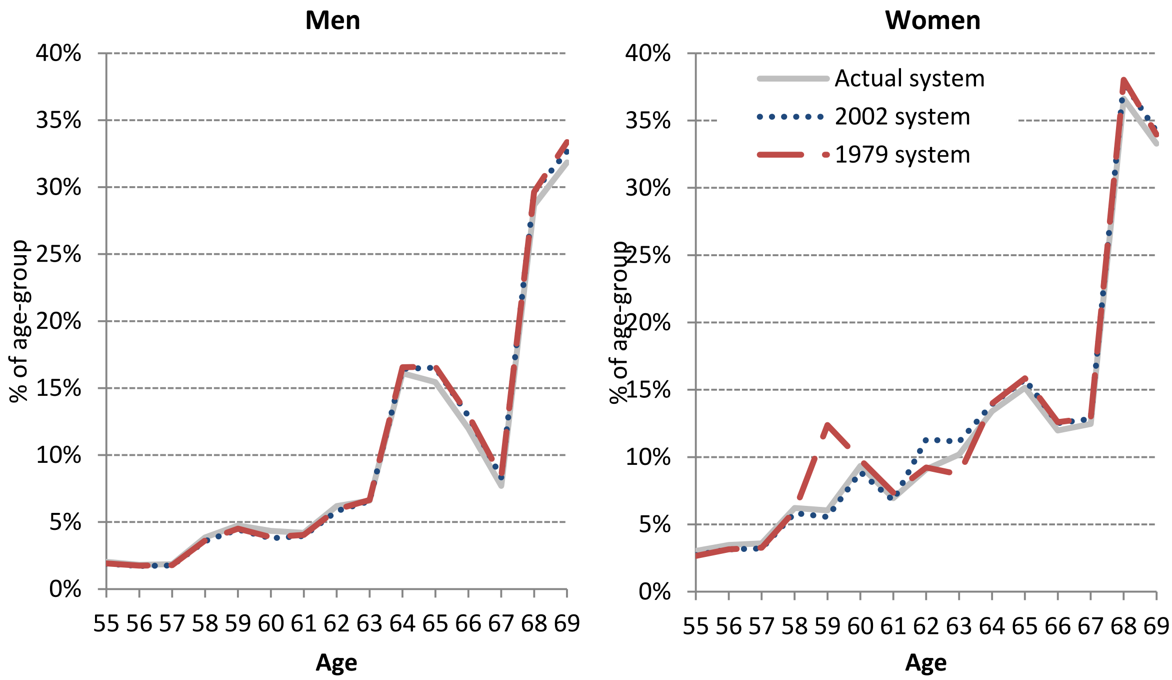

Source: ELSA, waves 1-9.

Note: Predicted exit rates are created using probit regression model 4.

Figure 6.2 Predicted one-year-ahead labour force exit rate by year, for men and women aged 55 to 69, under 1979, 2002, and actual pension systems

Source: ELSA, waves 1-9.

Note: Predicted exit rates are created using probit regression model 4.

To assess the impact of reforms on the evolution of exit rates over time, we can examine mean predicted exit rates in each year of our data, under the actual and counterfactual pension systems. Figure 6.2 plots these predicted exit rates, for both men and women, taking the average exit rate across the individuals in our sample aged 55 to 69. Both figures show lower exit rates for later pension systems in each year, with the gap between the 1979 and actual systems getting larger until 2010. There is then a mild increase in exit rates under the actual system, as compared to the 2002 system, in 2012, which strengthens in 2014. One explanation is this is due to increases in pension wealth that resulted from the introduction of the 'triple lock' indexation in 2011.

Of course, looking at mean exit rates by themselves is not necessarily very informative of the impact of retirement on longer working. What matters is the age at which exit rates are elevated, as well as their mean level. One way of interpreting results that takes account of this fact is to calculate, analogously to the construction of Figure 4.5, the predicted mean retirement age for an individual aged 55 and in work who experienced the exit rates we predict for each system. Such a calculation, analogous to the 'period' version of a survival curve, is a parsimonious way to summarise the implications of our counterfactual simulations that does not require us to predict exit rates outside of our data period.

Figure 6.2 shows the predicted mean exit rates for each year, under each system, for both men and women. We see that the increase in male and female mean retirement ages is a feature of our predicted model under the actual system as well as the counterfactual systems and the differences between systems are modest. While we have found effects of incentives, there have not been large changes in incentives as a result of reforms meaning that longer working is not overwhelmingly driven by changes to pension systems. While the increase in the female SPA that was legislated for in 1995 has a substantial impact on the exit rates of women at particular ages, its explains only a minority of the increase in employment of older women over this period.

Figure 6.3 Implied mean retirement age, for men and women, under 1979, 2002, and actual pension systems

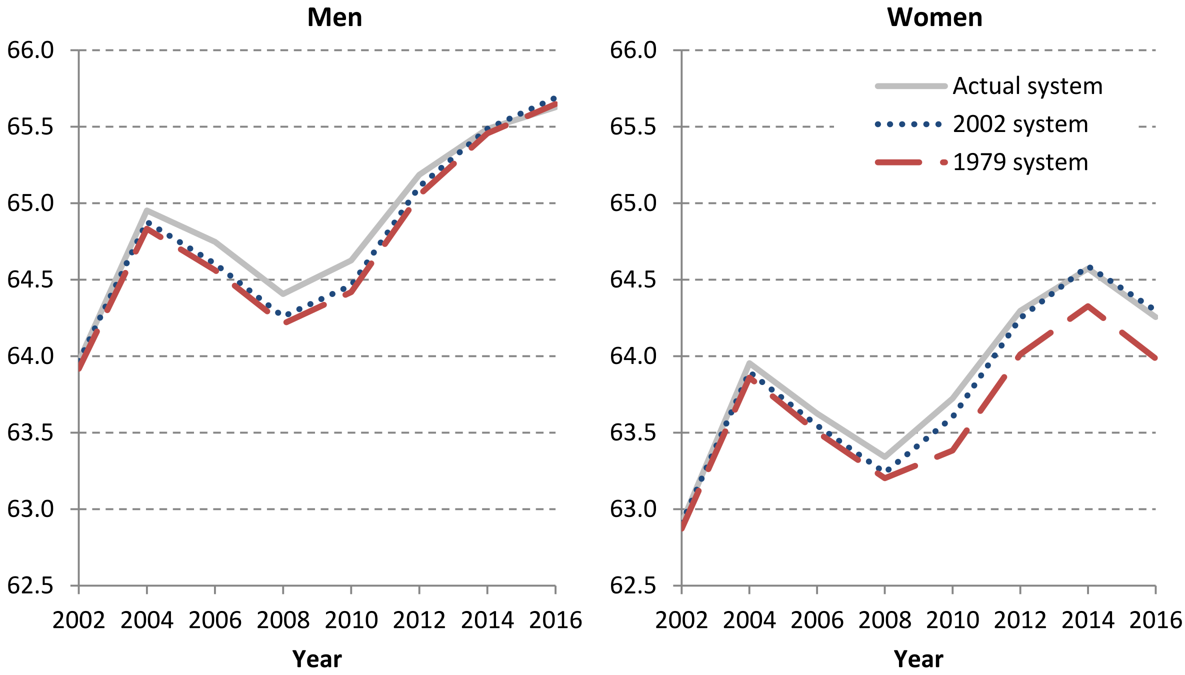

Source: ELSA, waves 1-9.

Note: For a given pension system, implied mean retirement ages for a given sex and year are calculated using the mean predicted exit rate at each age for each sex, amongst our regression sample, according to probit model 4, when using the calculated incentives for that pension system. We then assume that an individual of each sex faces the mean predicted one-year ahead retirement hazard rate at each age and that anyone not retired before age 75 retires at that age with probability 1. We compute the predicted exit rates, and resulting implied mean retirement age, for each of the 1979, 2002, and actual systems.

7. Conclusions

What role do the incentives provided by pensions systems play in driving the decision of whether to leave work at older ages and to what extent are the rising employment rates of recent decades a consequence of policy reforms? In this paper we have drawn on the ELSA panel dataset of 2002–03 to 2018–19 to make an empirical assessment of the relationship between financial incentives and labour force exit. Using linear probability and probit models, we have found a positive association of wealth with labour market exit, and a mild but positive association with implicit taxes. However, the inclusion of fixed effects suggests that the relationship between wealth and exit from work is driven by those who have a tendency to stay in work – perhaps due to smaller disutility from working – tending to have lower wealth. When fixed effects are included, we finder a larger effect of implicit taxes: a 10ppt rise in the implicit tax rate is found to cause a 0.4ppt rise in the one-year-ahead exit rate. This is significant, although not extremely large, when set against a baseline exit rate of 8% amongst 55- to 74-year-olds. Furthermore, for the first time, we have made an estimate of the effect of reaching the SPA on exit from work, while controlling for the individual financial incentives associated with crossing this threshold. Our findings lend further support to the idea that focal points in pensions systems can have real effects on behaviour even if their financial implications are limited.

We then used our estimated regression model to predict the labour force exit rates that would have been seen in a counterfactual situation where reforms to pension systems did not take place. Over the period from 2002–03 to 2016–17, the expected age of retirement for a 59-year-old man rose by 2.5 years, from 65.1 to 67.7. The female expected retirement age rose from 64.1 to 65.7, an increase of 1.6 years. However, our counterfactual simulations find that these changes are not primarily attributable to changes in incentives as a result of pension reforms and that differences between the 1979, 2002 and present-day systems are modest in their implications for aggregate working patterns.

Overall, our findings suggest a significant role for incentives in determining the timing of retirement. While cross-sectional variation in lifetime earnings and a number of relatively small reforms enable us to identify the effect of incentives, including the signalling effect of the SPA, the UK pension reforms in the period we examine have not changed marginal incentives to work much, and certainly not as much as in other countries. For this reason, much of the pattern of longer working at older ages does not seem to be attributable to reforms. Even in a setting where the impact of reforms on longer working has been modest, we have seen substantial increases in employment at older ages, indicating that other factors, alongside incentives, are also important in driving longer working.

References

Banks, J., Brugiavini, A., and G. Pasini, (2020), 'The Powerful Combination of Cross-country Comparisons and Life-History Data', *Journal of the Economics of Ageing*, Volume 16, https://doi.org/10.1016/j.jeoa.2019.100206.

Banks, J and C. Emmerson, (2021), 'A Lifetime of Changes: State Pensions and Work Incentives at Older Ages in the UK, 1948–2018' in *Social Security Programs and Retirement around the World: Reforms and Retirement Incentives*, Chicago: University of Chicago Press

Banks, J., Emmerson, C. and G. Tetlow, (2005), 'Estimating pension wealth of ELSA respondents', Working Paper no. 05/09. London: Institute for Fiscal Studies

Börsch-Supan, A. and C. Coile (eds), (2021), *Social Security Programs and Retirement around the World: Reforms and Retirement Incentives*, Chicago: University of Chicago Press

Blundell, R. and P. Johnson, (1999), 'Pensions and Retirement in the UK', in Gruber, J., and D. Wise (eds), 1999, *Social Security Programs and Retirement around the World*, Chicago: University of Chicago Press

Blundell, R., Meghir, C. and S. Smith, (2002), 'Pension Incentives and the Pattern of Early Retirement', *The Economic Journal*, vol. 112, pp. 153-170

Brugiavini, A., Cavapozzi, D., Pasini, G., and E. Trevisan, (2013). 'Working life histories from SHARELIFE: a retrospective panel', SHARE WP Series (pp. 11-13).

Coile, C., Milligan, K., and D. Wise (eds), (2019), *Social Security Programs and Retirement around the World: Working Longer*, Chicago: University of Chicago Press

Cribb, J., Emmerson, C., and G. Tetlow, (2016), 'Signals matter? Large retirement responses to limited financial incentives', *Labour Economics*, vol. 42, pp. 203–12 (https://doi.org/10.1016/j.labeco.2016.09.005).

Gruber, J., and D. Wise (eds), 1999, *Social Security Programs and Retirement around the World*, Chicago: University of Chicago Press

Gruber, J., and D. Wise (eds), 2004. *Social Security Programs and Retirement around the World: Micro Estimation*, Chicago: University of Chicago Press

Office for Budget Responsibility (2022), *Fiscal risks and sustainability*, July 2022 (https://obr.uk/frs/fiscal-risks-and-sustainability-july-2022/).

Wise, D. (ed), 2016, *Social Security Programs and Retirement around the World: The capacity to work at older ages*, Chicago: University of Chicago Press

Appendix

Figure A.1 Labour force participation and pension receipt by age and sex, in 2002–03 and 2018–19

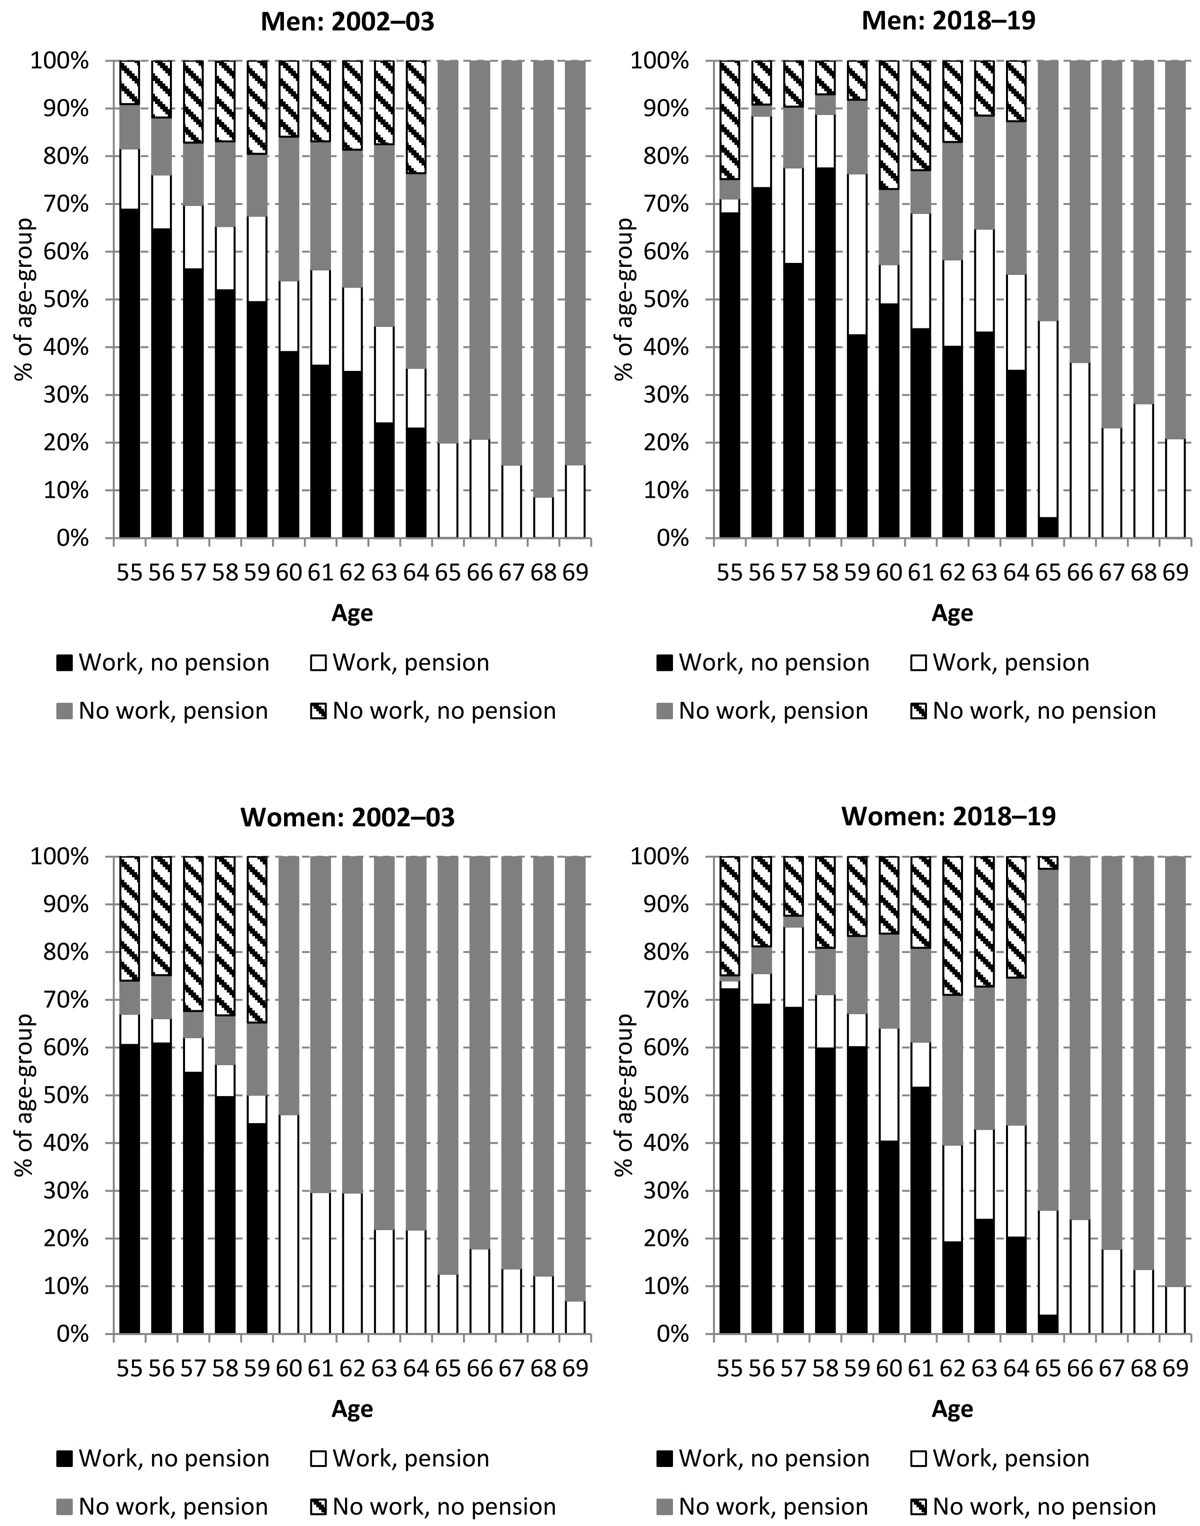

Source: ELSA, waves 1 and 9.

Note: "Work" is defined as being in work or temporarily away from work, in the last month. "Pension" is defined as being in receipt of private or state pension income.

Alternative results: Tables A.1 and A.2 show an alternative specification in which we control (and show the effects of) log current earnings and log wealth (and exclude total earnings from the controls). The implicit tax effects are similar in magnitude to the specifications in the main text. Log wealth is significant in more specifications than in the main text, indicating that the effect of wealth may be diminishing as wealth rises.

Table A.1. Alternative linear probability regression results: effects on one-year-ahead exit

A: No controls for unobserved heterogeneity							
	Implicit Tax	Log earnings	Log wealth	Reach SPA next year	Age controls	Time controls	Cohort controls
(1)	0.089***	-0.037***	0.015***	-	None	None	None
	0.008	0.003	0.003	-			
(2)	0.009	-0.018***	0.014***	0.068***	Quadratic	None	Linear
	0.007	0.003	0.003	0.012			
(3)	0.013	-0.019***	0.015***	0.068***	Quadratic	Dummies	Linear
	0.008	0.003	0.0027	0.012			
(4)	0.014*	-0.019***	0.014***	0.074***	Dummies	Dummies	None
	0.008	0.003	0.003	0.018			
B: Random effects							
	Implicit Tax	Log earnings	Log wealth	Reach SPA next year	Age controls	Time controls	Cohort controls
(5)	0.077***	-0.033***	0.021***	-	None	None	None
	0.008	0.003	0.003	-			
(6)	0.012	-0.015***	0.012***	0.064***	Quadratic	None	Linear
	0.008	0.003	0.003	0.011			
(7)	0.021***	-0.017***	0.013***	0.075***	Quadratic	Dummies	Linear
	0.008	0.003	0.003	0.018			
C: Individual fixed effects							
	Implicit Tax	Log earnings	Log wealth	Reach SPA next year	Age controls	Time controls	Cohort controls
(8)	0.040***	-0.019***	0.023***	-	None	None	N/A
	0.011	0.004	0.004	-			
(9)	0.023**	-0.007*	0.004	0.053***	Quadratic	None	N/A
	0.010	0.004	0.004	0.012			
(10)	0.042***	-0.006*	0.005	0.050***	Quadratic	Dummies	N/A
	0.011	0.004	0.004	0.012			

Note: All specifications include the set of control variables set out in Eq. (1) and in the note to Table 5.1, with the exception that we here include log current earnings and log wealth, and exclude total lifetime earnings. Standard errors are clustered at the individual level and shown in grey. Sample size for all specifications remains 22,629 observations.

Table A.2. Alternative probit regression results: average marginal effects on one-year-ahead exit

A: No controls for unobserved heterogeneity							
	Implicit Tax	Log earnings	Log wealth	Reach SPA next year	Age controls	Time controls	Cohort controls
(1)	0.090***	-0.030***	0.013***	-	None	None	None
	0.008	0.002	0.003	-			
(2)	0.011	-0.014***	0.014***	0.045***	Quadratic	None	Linear
	0.008	0.002	0.003	0.006			
(3)	0.013	-0.015***	0.015***	0.044***	Quadratic	Dummies	Linear
	0.008	0.002	0.003	0.006			
(4)	0.012	-0.015***	0.014***	0.050***	Dummies	Dummies	None
	0.008	0.002	0.003	0.012			
B: Random effects							
	Implicit Tax	Log earnings	Log wealth	Reach SPA next year	Age controls	Time controls	Cohort controls
(5)	0.090***	-0.030***	0.013	-	None	None	None
	0.008	0.002	0.003	-			
(6)	0.012	-0.014***	0.015***	0.047***	Quadratic	None	Linear
	0.009	0.002	0.003	0.006			
(7)	0.014*	-0.015***	0.016***	0.046***	Quadratic	Dummies	Linear
	0.009	0.002	0.003	0.006			

Note: All specifications include the set of control variables set out in Eq. (1) and in the note to Table 5.1, with the exception that we here include log current earnings and log wealth, and exclude total lifetime earnings. Standard errors are clustered at the individual level and shown in grey. Sample size for all specifications remains 22,629 observations.