

Joint Retirement of Couples: Evidence from Discontinuities in Denmark

Esteban García-Miralles, Jonathan M. Leganza

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

<https://www.cesifo.org/en/wp>

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: <https://www.cesifo.org/en/wp>

Joint Retirement of Couples: Evidence from Discontinuities in Denmark

Abstract

We study how social security influences the retirement behavior of couples. First, we exploit over two decades of full-population data and a discontinuity design to document sizable retirement spillovers to spouses when individuals reach pension eligibility age. Next, we explore underlying mechanisms. We find age differences within couples to be a fundamental determinant of joint retirement, which is driven by older spouses working longer. Accounting for these age differences reveals a strong gender gap, which prevails after controlling for relative earnings. Finally, in a complementary analysis we show that a reform increasing eligibility ages induces similar spillovers to spouses.

JEL-Codes: J140, J260, D100, H550.

Keywords: joint retirement, pension eligibility age, couples labor supply.

Esteban García-Miralles
University of Copenhagen / Denmark
egm@econ.ku.dk

Jonathan M. Leganza
Clemson University / SC / USA
jleganz@clemson.edu

July 11, 2021

We thank our advisors, Gordon Dahl, Itzik Fadlon, Miriam Gensowski and Mette Gørtz for support and guidance. We also thank Martin Browning, Julie Cullen, Meltem Daysal, Mette Ejrnæs, Maria Fitzpatrick, Alex Gelber, Thomas Høgholm Jørgensen, Gaurav Khanna, Claus Kreiner, Rafael Lalive, Søren Leth-Petersen, Torben Heien Nielsen, Krislert Samphantharak, Benjamin Ly Serena, Adam Sheridan, Jakob Egholt Søgaaard, and participants at the CESifo Area Conference on Public Economics, at the Institute for Fiscal Studies, at the CEBI and Graduate Seminars at the University of Copenhagen, at the DGPE conference, and at the 45th Spanish Economic Association Meeting for helpful comments. García-Miralles gratefully acknowledges funding from the Novo Nordisk Foundation (grant no. NNF17OC0026542) and from the Danish National Research Foundation through its grant (DNRF-134) to the Center for Economic Behavior and Inequality (CEBI).

1 Introduction

In recent decades, aging populations have led to widespread pension reforms. These reforms, and pension systems in general, are often designed at the individual level, but if couples coordinate their retirement, it can have sizable implications for projections of labor supply, budgetary estimations, and welfare analyses. It is therefore important to understand how social security influences the retirement behavior of couples.

A large body of literature takes the first step and documents the existence of coordinated retirement behavior within the household (e.g. [Gustman and Steinmeier, 2000](#); [Coile, 2004](#); [Stancanelli and Soest, 2012](#); [Stancanelli, 2017](#); [Lalive and Parrotta, 2017](#); [Honoré et al., 2020](#); [Willén et al., forthcoming](#)). However, convincing reduced-form evidence is still relatively scarce, and much less is known about the mechanisms that underlie the retirement behavior of couples. Understanding these mechanisms can help to successfully design and evaluate social security policies and can inform models of household labor supply and retirement.

Providing causal evidence on joint retirement and its underlying mechanisms is difficult. First, there is an identification challenge, as there are many reasons why one may observe associations between spouses retirement, such as correlated age, shocks to health or income, shared assets, and common preferences for leisure. Second, even with a successful identification strategy, taking the next step to shed light on mechanisms requires a rich dataset that allows for complementary analyses, such as investigating heterogeneity, and that is large enough to obtain precise estimates.

In this paper, we overcome these identification and data challenges to provide causal evidence on how social security impacts the retirement of couples and on the mechanisms that give rise to joint retirement behavior. Our context is the public pension system in Denmark, which is largely comparable in structure to most other developed countries. We focus on the early pension eligibility age, the age at which individuals first become eligible for early retirement benefits. Eligibility ages in general have been shown to greatly impact own retirement across different settings (e.g. [Staubli and Zweimüller, 2013](#); [Manoli and Weber, 2016](#); [Deshpande et al., 2020](#)), and the early pension eligibility age that we study has traditionally been the main pathway to retirement in Denmark.

Our analysis proceeds in three steps. In the first step, we document the existence of joint retirement behavior at the early pension eligibility age of 60. Using over two decades of administrative data on the entire population and an age discontinuity design, we show a sharp increase in own retirement when individuals reach their own pension eligibility age:

they are 20 percentage points more likely to retire right when they become eligible for early retirement benefits. We then estimate a sizable spillover effect to their spouses: we document a sharp 1.5 percentage point increase in spousal retirement right as their partners reach the pension eligibility age. This amounts to a scaled spillover effect on spouses of 7.5%; for every 100 individuals who retire upon reaching the early pension eligibility age, around 8 of their spouses adjust their behavior to retire at the same time. We further document spousal spillovers to claiming of pension benefits and to labor market earnings.

The identifying assumption in our age discontinuity design is that, in the absence of an individual becoming eligible for pension benefits at age 60, own outcomes and spousal outcomes would have evolved smoothly through the age of the individual. A potential concern involves age correlations between spouses. If spouses tend to be of a similar age, then spousal outcomes might exhibit a discontinuity at age 60 due not to the pension eligibility of their partner, but due to becoming eligible for benefits themselves. To address this concern, we control for spousal age using a flexible specification, and we show in a placebo exercise that our design successfully removes any effect induced by age correlations between spouses.

In the second step, we use the same age discontinuity design to explore mechanisms that underlie joint retirement behavior through a detailed sub-sample analysis and the use of reweighting methods. We document several explanatory channels that give rise to joint retirement behavior. Most striking are the results on age differences. We find that age differences between couples are a crucial determinant of joint retirement behavior, as joint retirement at the early pension eligibility age is primarily driven by older spouses who work past their own pension eligibility age while waiting for their younger spouse to become eligible for pension benefits as well. A key implication of this result is that this type joint retirement behavior increases aggregate labor supply.

We also find a strong gender gap: female spouses are more likely to adjust their own labor supply to time their retirement with their male partners. Importantly, this result is only revealed after we control for the age difference within the couple, since older partners are disproportionately males, which confounds the results from a simple comparison of males and females. The gender gap prevails even after controlling for relative earnings within the couple, suggesting that the gap is not explained by traditional economic incentives and that gender norms may be playing a role.

We then study the role of relative earnings and find two relevant results. First, we find patterns consistent with the collective model of household behavior; couples in which the primary earner values joint leisure more are more likely to retire jointly. Structural work

from Gustman and Steinmeier (2000) and Browning et al. (2021) show that males value joint leisure more. Consistent with this, we find that couples where male spouses are the primary earner are more likely to retire jointly. Second, we find patterns suggesting that couples consider the opportunity cost of retirement, seen as foregone earnings. After controlling for gender, we find that younger spouses who are secondary earners are more likely to retire jointly by retiring earlier, while older spouses who are primary earners are more likely to retire jointly by retiring later.

Our third and last step, after having documented joint retirement behavior and its leading mechanisms, is to explore joint retirement in the context of social security reform. One common reform involves raising pension eligibility ages in an effort to increase labor supply. Our initial estimates inform policymakers interested in potential spillovers to spouses from these types of reforms; yet, our setting allows us to further extend our analysis by directly studying this policy variation of interest. Specifically, we study a major 2011 policy change that discontinuously increased the early pension eligibility age by six months, contingent on exact birth date. We use monthly administrative records on employment and public benefits for the entire population and a local difference-in-differences design to study the causal effect of the reform. We find that individuals affected by the reform delay their retirement, and their spouses respond by delaying their retirement as well. We estimate a scaled spillover effect from this design of 9%, which is similar to our previous estimate using the age discontinuity design.

Our paper relates generally to the rich literature that studies the retirement behavior of couples.¹ Early evidence provided by Hurd (1990) documented associations in retirement between spouses. Ample structural work analyzes household retirement behaviors, often finding evidence of coordinated retirement between spouses due to complementarities in leisure (e.g. Zweimüller et al., 1996; Blau, 1998; Gustman and Steinmeier, 2000, 2004; An et al., 2004; Jia, 2005; Bingley and Lanot, 2007; Van der Klaauw and Wolpin, 2008; Casanova, 2010; Michaud and Vermeulen, 2011; Honoré et al., 2020).

Our paper relates most closely to a growing number of recent studies that use a reduced-form approach to document how pension eligibility ages impact the retirement behavior of couples (Banks et al., 2010; Hospido and Zamarro, 2014; Selin, 2017; Lalive and Parrotta,

¹We also naturally relate to papers that study the impact of own pension eligibility on own retirement (e.g. Mastrobuoni, 2009; Behagel and Blau, 2012; Staubli and Zweimüller, 2013; Cribb et al., 2016; Manoli and Weber, 2016; Geyer and Welteke, 2019; Haller, 2019; Nakazawa, 2019; Deshpande et al., 2020).

2017; Atalay et al., 2019; Bloemen et al., 2019; Willén et al., forthcoming).² Two of these papers stand out as the closest to our empirical strategy. Lalive and Parrotta (2017) use census data from Switzerland and an age discontinuity design similar to our main empirical design to document spillover effects from reaching pension eligibility age. They find significant spillovers to female spouses but not to male spouses. Willén et al. (forthcoming) use administrative data from Norway, and exploit a pension reform that lowered the retirement age for workers in some firms but not in others. They explore spillover effects across public programs and across spouses, restricting their analysis of spousal spillovers to younger spouses only, and finding effects on female spouses but not on male spouses. They test whether this gender difference is explained by differences in relative earnings and find that it is.

Our main contribution related to these papers and the literature more generally is to provide novel and comprehensive evidence on the mechanisms that explain joint retirement behavior. In particular, our paper is the first to explore the role of age differences between spouses, which we find to be a fundamental determinant of joint retirement behavior. We also highlight that accounting for age differences is important when studying other dimensions. For instance, we show that age differences can confound estimates from simple gender splits, an area where the existing literature has found particularly mixed results.³ In addition, by accounting for both age differences and gender, we are able to provide new and detailed evidence on the effect of relative earnings that supports collective models of household behavior.

We also contribute by providing causal evidence on joint retirement for both a representative population and a representative pension system, whereas previous work has been restricted to specific groups of the population or to pension systems where eligibility ages differ by gender. Our sample is representative as it includes male and female spouses as well as spouses who are either younger or older than their partner reaching pension eligibility age. Our context is representative of most modern pension systems, as the pension eligibility age of males and females is the same, and taxation and pension benefits are independent between

²Another strand of the literature studies other features of the pension system that indirectly affect retirement. Baker (2002) investigates a spouse allowance that is means-tested jointly against the partner's income. Coile (2004) explores the financial incentives to retire of each spouse and its interrelation. Stancanelli (2017) studies an increase in the contribution period needed to claim full pension benefits. Kruse (2020) studies the removal of the earnings test on early pension benefits.

³Some studies find spillovers effects to female spouses only (e.g. Hospido and Zamarro, 2014; Lalive and Parrotta, 2017; Bloemen et al., 2019; Willén et al., forthcoming) while others find spillovers only to male spouses (e.g. Banks et al., 2010; Atalay et al., 2019). Also, Selin (2017) does not find any significant effects when exploring only spillovers to male spouses.

spouses.⁴

Our findings have two main implications for policy. First, our main results show that there exist sizable spillover effects that should be incorporated into policy design and evaluation to avoid systematic deviations. Our analysis of the reform further shows that changes to the pension eligibility age generate sizable spillover effects. Second, our results on mechanisms indicate that policy makers should consider the age composition of couples when designing and predicting effects of social security reform, especially for countries that are still increasing eligibility ages of females to converge to those of men, as females tend to be the younger member of couples, who we show generate larger spillovers.

The rest of the paper is structured as follows. Section 2 describes the institutional background. Section 3 presents the data and the samples of analysis. Section 4 lays out our empirical strategy for estimating the effect of reaching pension eligibility age and reports the results. Section 5 explores the mechanisms that underlie joint retirement behaviors. Section 6 analyzes the reform that increased pension eligibility ages. Section 7 concludes.

2 Institutional Background

The Danish retirement system is broadly typical of other developed countries (OECD, 2019). The two primary sources of retirement income are benefit payments from public pensions and savings in private retirement accounts, with the latter coming from personal or employer contributions during working life.

There are two sources of public pension benefits. The Old Age Pension (OAP) provides universal retirement income security at old ages, and the Voluntary Early Retirement Pension (VERP) provides early retirement benefits for those who choose to participate in the program. The majority of workers participate in the VERP program, about 80% of the birth cohorts we study. As VERP plays a major role in determining labor supply and retirement patterns of the Danish population, we focus our analysis on the VERP eligibility age.

Voluntary Early Retirement Pension. The VERP program, introduced in 1979, provides access to early retirement benefits, traditionally from age 60. Participating in VERP requires making modest contributions to qualified unemployment insurance funds during working life. Benefits are flat-rate and result in a fixed amount paid to all workers equal to roughly \$27,000 annually (in 2010 USD).

⁴In the past, many pension systems had different pension eligibility ages for males and females, but currently most developed countries have the same pension eligibility age for both genders or are in a process of convergence (OECD, 2015).

The decision to claim VERP benefits is tightly linked to retirement, although they are technically separate decisions. The reason for this tight link is that the design of VERP produces strong incentives to retire at the same time as claiming. First, individuals must be “available to the labor market” in order to transition to VERP; that is they must be employed or actively searching for jobs or on a special transition pension (*delpension*). Hence, if individuals choose to leave the labor market before reaching VERP eligibility age, they will potentially forgo 5 years of benefits. Second, there are no actuarial adjustments for deferring claiming, so delaying claiming by one year amounts to a foregone year of benefits. Third, benefits are also subject to substantial means testing against labor market earnings at essentially 100%, which creates strong disincentives to keep working after VERP benefits are claimed, and against private retirement accounts.

The VERP program has remained fairly stable over time. Importantly, during the period 1991–2013, which we use in our first analysis, the VERP eligibility age remained constant at age 60. Two changes occurred during this period that are worth mentioning. First, the number of years that an individual has to contribute to an unemployment fund to qualify for VERP increased over time.⁵ Second, a pension reform in 1999 introduced incentives for individuals to delay claiming of VERP benefits by two years, to age 62. By postponing claiming to age 62 the flat-rate benefits are slightly increased (from approximately \$27,000 to \$29,600) and they are no longer means-tested against private pension accounts. The effect of the reform was a mild decrease in the number of people claiming at age 60, and a new discontinuous increase at age 62. Across our different analyses we show that this reform does not meaningfully affect our results.⁶

In 2011 the Danish government announced a pension reform increasing pension eligibility ages in 6 month steps contingent on birthdate. Both the VERP and OAP ages increased, as well as the incentivized VERP age. In Section 6 we describe this reform in detail, and we exploit the first discontinuity created by the reform to study the effect on joint retirement. We focus on the first cohort affected, those born after the cutoff date of January 1, 1954, whose VERP eligibility age was raised from 60 to $60\frac{1}{2}$, and who are first impacted in 2014 when they turn 60.

Two features of the VERP program make it ideal to study joint retirement behavior.

⁵From 1985, individuals had to contribute for 15 years out of the last 20 years. In 1990 the number of years increased to 20 out of the last 25, and in 1995 it increased to 25 out of the last 30.

⁶While not a reform of VERP, between 1992 and 1996 a transitional benefits program allowed long-term unemployed above age 55 (and above age 50 from 1994) to retire with similar conditions as the VERP program.

First, the pension benefits are independent between spouses. The decision to claim or retire does not have any direct effect on the pension benefits of the spouse. Therefore, we can rule out direct effects on the pension benefits of spouses as a mechanism for joint retirement in our analyses.⁷ Second, the pension eligibility age is the same for men and women over the entire period considered, which has two advantages. First, our setting is representative of modern systems in most OECD countries that have eliminated the gender gap in statutory pension eligibility ages over the last decades (OECD, 2015, 2017). Second, we can study heterogeneous effects by gender, age composition and income shares within the couple that are not affected by differential pension eligibility ages.⁸

Old Age Pension. The OAP provides universal old-age benefits. The eligibility age was traditionally 67, and it was lowered to 65 by the 1999 reform. Less than 5% of the spouses in our samples of analysis are therefore old enough to be eligible for OAP. Benefits are roughly \$15,000 for married or cohabiting individuals and \$20,000 for single individuals. Individuals are eligible for full OAP benefits if they have resided in Denmark for at least 40 years, and benefits are reduced proportionally if individuals have resided for a shorter period. Claiming benefits is an active choice, and the decision to claim is separate from the decision to cease working. From 2004, individuals can defer claiming OAP benefits and receive (approximately) actuarially-fair increases in benefits. Also, the means testing of OAP is less strict than that of VERP.

3 Data and Samples of Analysis

3.1 Data

We use administrative data covering the entire population of Denmark over the period 1986–2014. Using personal identifiers for each individual, we combine different registers with information on labor market outcomes, pension benefits, socio-demographics and family linkages. Variables are third-party reported on an annual basis and contain a large degree of disaggregation. Individuals cannot select themselves out of the registers, and they only exit the registers if they migrate out of the country or die.

⁷This is in contrast to Baker (2002) who studies exactly these direct links between spouses' pension benefits, and also to the second empirical design of Atalay et al. (2019) which is based on the characteristics of Vietnam veterans' pension system.

⁸Note that this is in contrast to Lalive and Parrotta (2017), who study a period where men and women face different statutory retirement ages, and to Atalay et al. (2019), who study a reform that raises women's pension eligibility ages to converge to that of men's.

In addition, we also use monthly-frequency register data on earnings for all employees in Denmark and on pension benefits for all recipients, which are available from 2008. We combine this data with the annual-frequency registers using the same individual identifiers. This allows us to define retirement ages with more precision, which is crucial for the analysis of the 2011 reform that increased the pension eligibility age by 6 months.⁹

3.2 Key Variables.

One advantage of our data is that we can measure different margins of labor supply and retirement behavior. We consider three main outcomes, which are defined either at the end of each calendar year (when using the annual data in the first, age-based empirical design) or as half-year measures (when using the monthly data in the second, reform-based design, since the reform increased the VERP eligibility age by 6 months).

Retirement: We define retirement as ceasing to earn labor market income. For the age-based design we use the annual data to define retirement as the year in which individuals earned income for the last time.¹⁰ Therefore, we define retirement as an absorbing state where the retirement variable takes the value one thereafter. In the robustness section we show that the results are robust to using a flow definition of retirement where we allow individuals to retire multiple times. These definitions are standard in the retirement literature (Coile and Gruber, 2007; Deshpande et al., 2020). For the reform-based design, we use the monthly data to define a dummy that takes the value one if an individual works past the first half of the year (that is, past July 1) in a given year. This accommodates the fact that individuals unaffected by the reform become eligible for benefits at the beginning of the reform year (2014) when they turn 60, whereas individuals affected by the reform become eligible at least 6 months later, when they reach age $60\frac{1}{2}$.

Claiming: We define claiming as receiving pension income, either VERP or OAP. For the age-based design we define an indicator equal to one if an individual receives any pension income in a given year. For the reform-based design we define an indicator that takes the value one if an individual received pension income before July 1 in a given year.

Earnings: In both research designs we use taxable annual labor market earnings from the annual registers. We winsorize this variable at the 1st and 99th percentile to reduce the influence of outliers. We adjust this variable for inflation using 2010 as a baseline and

⁹This new dataset, often referred to as *eIncome*, is described in more detail in Kreiner et al. (2016).

¹⁰We allow for some small positive income, equivalent to 1 month of average earnings, to accommodate the fact that individuals can receive some labor income after they have retired, such as holiday payments or delayed wages.

convert Danish kroner to U.S. Dollars using the exchange rate $1 \text{ USD} = 5.56 \text{ DKK}$.

3.3 Samples of Analysis

We define two samples of analysis, one for each research design. For both of our research designs we start with the full population of Danish couples who reside in Denmark between 1991 and 2014. We define couples as those who are either married, or in a registered partnership, or cohabiting. To avoid endogenous changes in marital status around the time of pension eligibility we identify couples when they are both below age 60 and observe them for as long as they remain together. We restrict the analysis to couples who are up to 8 years apart from each other, which excludes around 5% of the sample on each side of the distribution. We illustrate the distribution of age differences within couples in panel (a) of Appendix Figure A.5, and we show that our results are robust to dropping this restriction in Section 4.4.

We focus the analysis on dual-earner couples. First, we restrict the sample to couples where the reference individual (that is, the focal partner who reaches their own pension eligibility age) has earned labor income at least once between ages 55 and 59. All cohorts in our sample of analysis are observed back to age 55 since we have data from 1986. We also exclude reference individuals who are self-employed or on disability benefits at least once between ages 50 and 59, as they are subject to different rules and regulations of the VERP scheme. Second, we restrict the sample to couples where the spouse has earned labor income at least once between ages 50 and 59. We use this longer period for spouses to ensure that our sample does not exclude younger spouses who retire in their early 50s, as they can potentially retire jointly with their older partners.¹¹

Age-based sample. For our age-based design, we consider the period 1991–2013, where the early pension eligibility age remained stable at age 60. This provides us with more than two decades of observations from individuals who faced the same pension eligibility age. We focus the analysis on couples where the reference individual is 57 to 60 years old, which leads to a sample size of 367,585 couples and 2,206,044 couple-year observations.

Reform-based sample. For our reform-based design, we consider the period 2008–2014, starting in 2008 because the monthly-frequency data is only recorded from that year. To

¹¹Note that there are four cohorts of spouses that we cannot observe before age 60 to impose the restriction, and therefore we keep all those spouses, who represent 0.4% of the sample. Similarly, there are nine cohorts of spouses that we cannot observe during the entire period between ages 50 to 59. In this case, we impose the restriction based on the years that we observe. This affects 12% of the spouses, of which 80% are observed for 5 or more years.

focus on individuals who are more likely to be impacted by the reform, we restrict this sample to reference individuals who have made qualifying contributions to the VERP program at least once between ages 50 and 59. Note that we cannot impose this restriction on the full age-based sample because we do not observe contributions far back in time, but in the robustness section we show that our results from both designs are robust to this decision.¹² In our baseline specification, we focus on individuals born within a 3-month window on either side of the January 1, 1954 cutoff, and we balance the sample, leading to a sample size of 10,321 couples and 73,395 couple-year observations.

Table A.1 presents summary statistics for the two samples of analysis and for their corresponding unrestricted population. The first four columns correspond to the age-based period of analysis (1991–2013) and the last four columns correspond to the reform-based period of analysis (2008–2014). First, we can compare the analysis samples to their corresponding unrestricted population samples. We note that both reference individuals and spouses in the analysis samples have higher earnings, higher education, and are less likely to be retired before age 60. This is mainly a consequence of restricting the analysis to dual-earner couples and to those who did not receive disability benefits in the past. Also note that the age difference between spouses is similar between the analysis sample and the population, but the standard deviation is smaller due to the restriction that drops spouses who are more than 8 years apart. Second, we can compare the two analysis samples to each other. Overall the two samples are similar. The reform-based sample has a smaller share of males (47% against 52%), higher earnings (\$64,156 against \$60,289) and is slightly more likely to be retired before age 60 (16% against 14%), but these differences are not statistically significant. These differences are in line with the effect of restricting the reform-based sample to VERP contributors, as females are more likely to contribute to the program. The age difference between partners in both analysis samples is similar and so are the standard deviations.

4 The Effect of Reaching Pension Eligibility Age

4.1 Age-Based Discontinuity Design

To identify the causal effects of individuals reaching pension eligibility age on their own retirement and on their spouses, we exploit the discontinuity that occurs around the early

¹²Specifically, we show that our age-based results are robust to imposing the restriction for the subsample of observations between 2008 and 2013, for whom we can observe past contributions. We also show that the reform-based results are robust to not imposing the restriction.

pension eligibility age. Specifically, we study the retirement patterns of reference individuals and their spouses around the eligibility age of the reference individuals, that is around age 60. Importantly, when analyzing spouses' retirement patterns we control flexibly for the effect of their age on their retirement behavior.

We lead our analysis with a graphical illustration of the retirement patterns of the reference individuals and their spouses, which then guides our estimation strategy and allows us to evaluate the assumptions of our parametric model.

Note that each member of a couple can potentially appear both as the reference individual and as the spouse in the analysis, as long as they are observed at ages 57–60 during the period considered. This reflects the dual nature of the couples' decision, and our design allows us to study their retirement behavior from both sides, observing them as reference individuals when they reach their pension eligibility age, and as spouses, when their partners reach their own pension eligibility age. In the heterogeneity analysis we will, nevertheless, split the sample by age differences and gender and each member of the couple will appear only as either the reference individual or the spouse.

4.2 The Effect of Reaching Pension Eligibility on Own Retirement

We begin by analyzing the retirement behavior of reference individuals around their own pension eligibility age. Specifically, in Figure 1 we pool individuals for the period 1991–2013 and plot raw means of each outcome variable for the reference individual against their own age. As expected, given the strong incentives to retire exactly at the pension eligibility age, we observe a clear discontinuous jump in all outcomes at age 60. An important feature of the data is that the outcome variables are measured at the end of each calendar year, and so is age, which we round up to months. Hence, individuals who turn 60 early in the year can claim their pension earlier that year than those who turn 60 later in the year. This induces a gradual phasing-in of the exposure to early retirement eligibility as monthly age increases from 60 to 61, a pattern captured by Figure 1.

We are interested in the “full-exposure” effect of being eligible for one entire calendar year. Individuals who are fully exposed are those who turn 60 at the beginning of January, becoming eligible for early retirement at that moment. These individuals are exposed to early pension eligibility for almost 12 months by the time their information is recorded in the administrative data in December. In contrast, individuals who turn 60 later in the year are eligible for a shorter period of time that year, so they are only partially exposed. Our estimation strategy exploits information from both partially and fully exposed individuals

to estimate the full-exposure effect with greater precision.

We quantify the full-exposure effect by estimating the following piecewise linear regression, which is closely guided by the graphical analysis:

$$y_{it} = \alpha + \beta_1 age_{it} + \beta_2 \mathbb{1}\{age_{it} \geq 60\} + \beta_3 \mathbb{1}\{age_{it} \geq 60\} \cdot age_{it} + \sum_{c=1991}^{2013} \kappa_c \cdot D_c + \epsilon_{it}, \quad (1)$$

where y_{it} is the outcome of interest for reference individual i at time t , age_{it} is monthly age of the reference individual at the end of the calendar year, and $\mathbb{1}\{age_{it} \geq 60\}$ is an indicator variable that takes the value one if the monthly age of the reference individual is 60 or above and zero otherwise. The model therefore estimates a discontinuous jump at monthly age 60 and a differential trend thereafter, as suggested by the graphical analysis. D_c are calendar year dummies. We estimate this regression for individuals between monthly ages 57 and just below 61.¹³

The full-exposure effect is then given by $\beta_2 + \frac{11}{12} \cdot \beta_3$. This estimator captures the treatment effect of being eligible for early pension during one full calendar year. It is composed of a sharp change in levels at the eligibility-age cutoff, captured by β_2 , and a change in trends, captured by the slope parameter β_3 that captures the effect of one year of eligibility from age 60. We plot the parametric fit of this model in Figure 1. The full-exposure effect corresponds to the vertical distance between the solid line and the dashed line just below age 61.¹⁴

The first row of Table 1 reports the full-exposure estimates for the different outcomes of the reference individual. The first column reports the full-exposure effect on retirement. The estimate is 0.2034, which means that reaching pension eligibility age increases the share of retired individuals by around 20 percentage points. Note that the share of retired individuals before they reach pension eligibility is also positive, around 16% before age 60, as illustrated in panel (a) of Figure 1. This shows that individuals can also retire before they reach pension

¹³Because the outcome variables are measured in December, individuals who turn 60 in December often do not have time to receive pension income until the next year. This is clearly seen in Figure 1, panel (b), where the dot for December is much lower. To prevent this from biasing our estimates, we exclude these individuals by adding a dummy variable that takes the value one if their monthly age is exactly 60. In Table 2 of the robustness section we show that the results are largely unaffected if these individuals are kept.

¹⁴A similar methodology is used by [Fadlon et al. \(2019\)](#) to study the effect of Social Security’s survivors benefits on labor supply in the U.S. Also, [Nielsen \(2019\)](#) explores the same age discontinuity for Denmark to study the effect of retirement on health.

eligibility.¹⁵ The second outcome of interest, pension claiming, is reported in the second column. The point estimate is 0.35, so around 35% of individuals claim VERP benefits by the end of their first year of eligibility. The effect for claiming is larger than for retirement for two reasons. First, it is not possible to claim VERP benefits before age 60, as illustrated in panel (b) of Figure 1, and second, individuals who claim can still have positive earnings in the same year. Finally, the third column reports the full-exposure effect on annual labor market earnings, which can potentially reflect responses both on the extensive margin and on the intensive margin. We estimate a decrease of \$8,642 in annual earnings after one year of exposure to pension eligibility.

Overall our results show that reaching pension eligibility leads to a strong first stage. Individuals are discontinuously more likely to retire after age 60. We now turn to estimate the causal effects of pension eligibility on spousal retirement behavior.

4.3 The Effect of Reaching Pension Eligibility on Spouses

For the spillover effect on spouses, we follow a similar empirical strategy as for reference individuals. The main difference is that we need to control for the effect of spouse’s own age on their retirement behavior so that we can isolate the causal effect of their partner’s pension eligibility.

We lead the analysis with a nonparametric illustration of spouse retirement patterns around their partners’ age, after accounting for the effect of the spouses’ own age. Specifically, we plot the residuals from the following regression:

$$y_{it}^s = \alpha + \sum_{a=49}^{69} \delta_a \cdot D_a^s + \sum_{a=49}^{69} \gamma_a \cdot D_a^s \cdot D_g + \sum_{c=1991}^{2013} \kappa_c \cdot D_c + \varepsilon_{it}, \quad (2)$$

where y_{it}^s is the outcome variable of interest for spouse s of individual i at time t , D_a^s are dummy variables for spouses’ monthly age, and D_g is a gender dummy. The residuals $\hat{\varepsilon}_{st}$ therefore capture the spouses’ retirement behavior that is not explained by their own age and gender.

The dots in Figure 2 plot spousal residuals $\hat{\varepsilon}_{it}$ binned over the monthly age of reference

¹⁵We have argued in Section 2 that there exists strong incentives to claim right at the early pension eligibility age, but individuals might cease to earn labor income earlier than 60 for a number of reasons: they might become unemployed or claim a partial pension until they turn 60, they might voluntarily stop working even if that implies the inability to claim VERP later on, and lastly, not all individuals in our sample qualify for VERP, as explained in Section 3, around 80% of the individuals in the age-based sample of analysis made contributions to qualify for VERP.

individuals. This illustrates the spouses’ retirement patterns that are driven by their partner’s age. We observe that spousal residuals change discontinuously right when their partner becomes eligible for early pension at age 60, resembling the same pattern we observed for the reference individuals themselves.¹⁶

Guided by this graphical analysis, we estimate a parametric model that quantifies the causal effect of one partner reaching pension eligibility age on the retirement behavior of their spouse. The estimating equation is similar to equation (1) for the reference individual, but with spouses’ outcomes as the dependent variables and additional controls for spouses’ age and gender that do not impose any functional form. The estimating equation is:

$$y_{it}^s = \alpha + \beta_1 age_{it} + \beta_2 \mathbf{1}\{age_{it} \geq 60\} + \beta_3 \mathbf{1}\{age_{it} \geq 60\} \cdot age_{it} + \sum_{a=49}^{69} \delta_a \cdot D_a^s + \sum_{a=49}^{69} \gamma_a \cdot D_a^s \cdot D_g + \sum_{c=1991}^{2013} \kappa_c \cdot D_c + \epsilon_{it}, \quad (3)$$

where y_{it}^s is the outcome of interest for spouse s of individual i , age_{it} is age of the reference individual in months, and $\mathbf{1}\{age \geq 60\}$ is an indicator variable that takes the value one if the reference individual is 60 or older (in terms of monthly age) and zero otherwise. D_a^s are dummy variables for spouses’ monthly age, and D_g is a gender dummy. We estimate this regression for the same sample of reference individuals, between ages 57 to 61, as before.

The full-exposure effect is again given by $\beta_2 + \frac{11}{12} \cdot \beta_3$. For illustrative purposes, Figure 2 superimposes the parametric fit of the model estimated in equation (3) over the residuals from equation (2). The full-exposure effect corresponds to the vertical distance between the solid and dashed lines just below age 61. The second row of Table 1 reports the full-exposure effect on spouses from their partner reaching pension eligibility age. The effects on all three spousal outcomes are statistically significant at the 1% level. These point estimates can be viewed as the reduced-form effects on spouses.

To judge the size of joint retirement behaviors, we report “scaled effects” in the last row of Table 1, defined as the full-exposure effect on the spouse divided by the full-exposure effect on the reference individual. These scaled effects are our preferred measure for reporting and interpreting joint retirement spillovers, as they are comparable across different outcomes, samples of analysis, and empirical strategies, including our reform-based design. We compute standard errors for these scaled estimates by bootstrapping (Andrews and Buchinsky, 2000;

¹⁶An alternative approach to this methodology is to estimate equation (2) adding age dummies for the reference individual and plotting those coefficients. Appendix Figure A.1 shows that the result is similar.

MacKinnon, 2006).¹⁷

The scaled effect on the retirement outcome is 7.5%. This is a sizable spillover, for every 100 individuals who retire right when they reach their early pension eligibility age, about 8 of their spouses are induced to retire as well. This is after controlling for the effect of the spouses' age on their own retirement behavior.

Claiming leads to scaled effect of 3.4%. This effect is smaller than the one for retirement for two reasons. First, the denominator is larger, that is, the full-exposure effect on the reference individual is larger for claiming than for retirement as discussed earlier. Second, the numerator is slightly smaller, the full-exposure effect on the spouses is smaller because of spouses who retire but do not claim. Knowing the spillover effect on claiming is important for policy and fiscal estimations, but for the reasons mentioned above it does not fully capture joint retirement behavior.¹⁸ In Section 5 we explore the interaction between claiming and the age composition of couples and its implications for joint retirement responses.

For earnings, the scaled effect is 9.8%. This outcome potentially captures both extensive margin responses and intensive margin adjustments that can be in the form of hours worked, choice of job, or effort. However, we cannot conclude that there are significant intensive margin responses based on the larger size of the scaled effect for earnings compared to retirement. Note that the size of the scaled effect for earnings depends on the relative earnings within each couple, and even if adjustments occur only through the extensive margin, the scale effect on earnings will be larger than the scaled effect on retirement if the spouses who adjust tend to be the primary earner. This in turn depends on the response heterogeneity and mechanisms underlying joint retirement, which we analyze in Section 5.

4.4 Threats to Identification and Robustness

Our identification strategy relies on the assumption that, once we control flexibly for the spouses' age on their own retirement, the discontinuous behavior that occurs when their partner reaches pension eligibility age is caused by that event, and nothing else. In this section we provide a number of tests to assess the validity of our design.

¹⁷Note that these scaled effects are conceptually similar to the estimates from an instrumental variables approach. We use scaled effects because they allow for a more flexible estimation of the second stage (the spouses' full-exposure effect) by estimating the jump at 60 and the differential trend separately. An instrumental variables approach, instead, imposes that the functional form of the jump and the differential trend estimated in the first stage is maintained in the estimation of the second stage.

¹⁸For an analysis of retirement and claiming in the U.S. see Deshpande et al. (2020). Note that while deferring claiming in the U.S. leads to actuarial adjustments of future pension benefits, in the VERP program there is no such actuarially fair updating, and therefore the decision to claim and retire are more closely related.

Placebo test. To be reassured that we successfully control for the effect of the spouses' age, we carry out a placebo test. We repeat the analysis for the same sample of reference individuals, but we randomly assign them a fake spouse of similar age. Specifically, we assign a spouse of the same age to half of the reference individuals, and we assign spouses who are between 1 and 3 years younger or older to the other half of the reference individuals.¹⁹ In this sample, spouses are likely to retire at the same time because their ages are highly correlated and most of them reach pension eligibility age at the same time. However, we should not observe any joint retirement behavior beyond the one due to this age correlation between spouses, given that fake spouses cannot influence each other. If our empirical strategy successfully controls for the effect of age correlations, then we should not find any evidence of joint retirement in this placebo sample. Reassuringly we do not find any, as reported in Appendix Table A.4 and Appendix Figure A.2.

Alternative specifications. In Table 2 we show that the results are robust to a series of changes in the model specification and in the sample definition. Row A reports the baseline estimates for comparison. In row B we extend the sample of analysis to include reference individuals of ages 55 and 56. In row C we exclude reference individuals aged 59 by adding a dummy variable to the model that takes one if the reference individual is 59 or older. This excludes monthly ages between 59 and 60 from the estimation of the counterfactual behavior. In row D we keep couples with partners that are more than 8 years apart from each other. Row E drops the dummy that identifies reference individuals who turn 60 in December, so that they are included in the estimation of the jump at 60 and the differential trend afterwards. Row F allows for a nonlinear counterfactual before age 60 by adding a second order polynomial of the reference individuals' age to the model. This nonlinear specification reduces our point estimates (e.g. the scaled effect on retirement becomes 5.0%), but note that we are fitting a second order polynomial over a short period of three years (ages 57 to 60). To account for this, in row G we increase the age range to include reference individuals of ages 55 and 56 (as in B) and fit a second order polynomial (as in F), obtaining spillover effects much closer to our baseline estimates. Row H controls for predetermined region and education of the reference individual and their spouse. Row I adds a dummy for individuals born after 1939, who are therefore affected by the 1999 reform that introduced incentives to claim VERP at age 62 and lowered the OAP to age 65. Row J estimates the effect over the period 2008–2013, which is almost the same period considered in the reform-based design

¹⁹Note that we do not use only spouses of the same age to avoid collinearity between the age of both partners.

that we present in Section 6. In row K we present estimates over the same period as in J, and restrict the sample to reference individuals who have made contributions to qualify for VERP at least once between ages 50 and 59. Note that we can only impose this restriction for these later calendar years as we do not observe contributions far back in the past. Finally, in row L we report the scaled effect for retirement defined as a flow variable, which allows individuals to retire multiple times (see Appendix Figure A.3 for the full-exposure effects). Reassuringly, our results are robust to all these changes.

Joint Retirement Over Time. In our analyses we have pooled two decades of observations to obtain precise estimates of the causal effects of reaching pension eligibility age on joint retirement of spouses. We show that the size of the scaled effect estimates has been stable over the entire period and therefore the size of the full-period estimates is also representative of the effect in most recent years, which are of more interest for policy and also the relevant period for comparison with the reform-based estimates derived from the 2011 reform that we present in Section 6.

Specifically, we replicate the analysis over 5-year running windows. We report the evolution of the scaled effects for the three outcomes of interest in Appendix Figure A.4, where each dot at year t corresponds to the scaled effect estimated for the period $t - 4$ to t . For instance, the last dot from 2013 reports the scaled effects estimated for the period 2009–2013.

Attrition. Individuals cannot self-select out of the registers. The only two reasons for an individual to exit the registers are either migrating out of Denmark or dying. If reaching pension eligibility caused either of these two things to happen, we would miss that individual from the sample, but in no case would they be wrongly considered as retired. Note also that [Nielsen \(2019\)](#) finds no evidence of increased mortality or worse health at retirement when studying the same age discontinuity in Denmark.

5 Explaining Joint Retirement: Heterogeneity and Mechanisms

The previous section has documented the existence of joint retirement behavior. We find sizable spillover effects from one spouse to another: for every 100 individuals who retire upon reaching pension eligibility age, 8 of their spouses adjust their labor supply to retire around the same time. But how does this labor supply adjustment occur? And what determines which spouse bears the adjustment?

In this section, we go beyond documenting joint retirement behavior and study the mechanisms that give rise to joint retirement, with the goal of explaining how couples retire to-

gether. To investigate this, we use detailed sub-sample analysis and reweighting methods to explore differences across three relevant characteristics: age differences within couples, gender, and relative earnings.

5.1 Age differences within the couple

We study the effect that relative age within partners has on joint retirement and find that it plays a crucial role. We begin our analysis by splitting our sample based on whether spouses are older or younger than their partners who are reaching age 60. For each of these subsamples we replicate the analysis and report the results in columns (1) and (2) of Table 3. Focusing on the retirement outcome, we observe that the scaled effect is 9.9% for older spouses and only 2.9% for younger spouses, while still highly significant. These results show that in our context, the preferred path to joint retirement is the one where the older spouse works longer, past their own pension eligibility age, while waiting for the younger spouse to reach pension eligibility and then retire together. This is in contrast to the alternative path to joint retirement where the younger spouse retires earlier, before reaching pension eligibility age, when the older spouse first become eligible for early pension benefits.

These results suggest an important role for the ability to claim own pension benefits in the decision to retire jointly. Older spouses who retire right when their younger partners reach pension eligibility age must continue working past their own pension eligibility, and can then claim benefits themselves. In contrast, younger spouses who retire right when their older partners reach pension eligibility age cannot claim their own benefits, since they themselves have not yet become eligible. This is a potentially financially costly decision, especially if it entails giving up the right to claim VERP later on due to the “transition to VERP” rules explained in Section 2.

Another important implication of these results is that, in our context, joint retirement behavior increases aggregate labor supply. Older spouses tend to work longer while waiting for their younger spouse to reach eligibility age, and this happens more often than the alternative pathway to joint retirement where younger spouses work less because they retire as soon as their partner reaches pension eligibility.

To explore the effect of age differences in more detail, we define subsamples based on smaller intervals of age differences and estimate joint retirement spillovers for these subsamples. The results are reported in Figure 3, where we plot the scaled spillover effects as spouses’ age increases relative to their partners. We observe that the largest spillovers are concentrated among spouses who are older, but not too far apart from their partners’ age.

Specifically, focusing on the retirement outcome, reported in panel (a), we find the largest effect (above 10%) for spouses up to 2 years older, followed by spouses who are between 2 and 4 years older. The effect decreases for spouses who are more than 4 years older than their partners. For younger spouses, we do not find evidence of differential spillovers in joint retirement as the difference between partners' age increases. The point estimates remain small and stable around the same size as for the pooled subsample of younger spouses (2.9%), although less precisely estimated due to smaller sample sizes.

Overall, these results point to age differences between partners and the ability to claim as crucial determinants of joint retirement. Policies that aim to account for the joint retirement of couples should consider the economic incentives faced by each age group, and particularly for the ability of each partner to claim own pension benefits. In our setting, younger spouses cannot claim benefits of their own if they retire when their partner reaches early pension eligibility. However, in other settings, such as those centered on later pension eligibility thresholds where younger spouses can also claim their own benefits when they retire at the same time as their older partners, the joint retirement spillover of younger spouses might be larger. In addition, our results indicate that economic analysis of intra-household behavior should account for the effect of the age-composition of couples.

5.2 The effect of gender

Next, we explore heterogeneity by gender, a dimension where previous studies have found particularly mixed results. Some of the difficulties faced by the literature include pension systems where eligibility ages differ by gender, reforms that affected one gender only, lack of statistical power to estimate small effects, and not accounting for confounded effects between age differences and gender. Our analysis overcomes these challenges, as there are no gender differences in the Danish pension system, benefits and taxation are independent between spouses, and we have statistical power to estimate gender differences controlling for other confounding factors such as age differences.

We begin by replicating our analysis over a simple split by gender. Column (3) of Table 3 presents results for the subsample of male spouses and column (4) for female spouses. The scaled effects for both male and female spouses is 7.5%, which could erroneously lead us to conclude that both genders are equally likely to adjust their behavior to retire jointly with their partner.

However, this simple split by gender masks important differences in the composition of relative age between spouses among the two groups. As in most countries, Danish men tend

to be the older member of the couple.²⁰ Specifically, in our analysis sample males are around two years older than females, as we illustrate in panel (b) of Appendix Figure A.5. We have shown that older spouses are much more likely to retire jointly, therefore the estimate found for men confounds the fact that the subsample of male spouses is composed by a larger share of older spouses. Therefore, to explore gender differences in joint retirement, we must control for the confounded age differences. We address this by reweighting the subsample of female spouses to match the distribution of age differences from the subsample of male spouses, and then re-estimating the spillover effect. The result is shown in column (5), where we observe that the scaled effect for females rises from 7.5% to 13%. We can then compare this scaled effect to the scaled effect for male spouses, reassured that the difference is not driven by the age-difference composition of both subsamples. Interestingly, we find that females clearly respond more, contrary to the conclusion that we could have reached from the simple split by gender.²¹ The reweighting strategy assumes that couples where females are the older spouse are comparable to couples where females are the younger spouse. We explore this in Appendix Table A.3, and show that these two types of couples are remarkably similar along observable characteristics such as labor market earnings, educational attainment, retirement probability, or whether they live in the Copenhagen region, all measured before age 57. Female spouses are very similar to each other regardless of whether they are the younger or the older member in the couple, and so are males.

A potential explanation for these gender differences in behavior is that relative earnings within couples confound joint retirement and gender. We study the role of relative earnings in detail in the next subsection, but regarding its impact on the gender gap, we show that the gender gap we find is robust to further reweighting the sample of female spouses to have the same distribution of earnings shares as male spouses. The results are reported in column (6) of Table 3, where the scaled effect estimate for retirement remains high at 13.2%.²² Our results therefore unveil a gender gap that cannot be explained by age or relative earnings

²⁰Hospido and Zamarro (2014) and Coile (2004) consistently find similar age gaps, of around two years on average, for different European countries and for the U.S. respectively.

²¹Furthermore, we find this gender gap both for couples where the female partner is the younger member as well as for couples where the female partner is the older member. We show this in Appendix Table A.2 where, as an alternative to the reweighting strategy, we split the sample in four, by gender and by relative age between partners.

²²The gender gap also remains when we further reweight the subsample of female spouses to ensure that the share who made contributions to qualify for VERP in the past is the same as in the subsample of male spouses. Note that we only observe VERP contributions for the most recent period of time and hence we perform this test for the period 2008–2013 only. The scaled spillover for retirement is 7% for males and 11.1% for females after reweighting.

within couples, suggesting a role for gender norms. This result adds to recent findings of gender differences that cannot be explained by traditional economic incentives such as [Daly and Groes \(2017\)](#), [Kleven et al. \(2019\)](#), [Gørtz et al. \(2020\)](#), and [Lassen \(2020\)](#).

Our results also document a new source of gender differences in earnings and labor supply. Unlike previous studies that focus on childbearing and childcare, we show a gender gap that originates in the dynamics of family formation and its interaction with joint retirement behavior, manifesting itself at the end of working life. Because males tend to be older than their female partners, couples who retire together most often achieve this either by males retiring later or by females retiring earlier, therefore increasing males lifetime earnings relative to females. Note that the “grandchild penalty” found by [Gørtz et al. \(2020\)](#) could explain part of the gap we identify, as grandmothers retire earlier to take care of their grandchildren, but it does not explain it all, as we also find that older female spouses are more likely to retire later, waiting to retire together with their younger partners.

5.3 The effect of relative earnings within couples

We now focus directly on the role of relative earnings within couples for joint retirement. To define the relative earnings of each member of the couple we compute predetermined earnings shares based on the average labor market earnings of each partner between ages 55 and 57, reporting the distribution of these shares in panel (d) of Appendix Figure A.5. We define an indicator for who is the primary earner in the couple based on these shares, excluding couples with very similar earnings shares (those between 47.5% and 52.5%, which represent 14% of the sample), although the results are robust to keeping them.

The interaction between relative earnings and gender. A growing literature studies the decision-making process of households through the lens of a collective model ([Chiappori, 1992](#); [Browning and Chiappori, 1998](#); [Donni and Chiappori, 2011](#)), where members with more negotiation power have more weight in the decision-making process of the household ([Browning et al., 1994](#)). If males and females differ in their preference for joint leisure, we would expect that the member with more power, the primary earner, will have a bigger influence in the joint retirement decision. We explore this in Table 4, where we replicate our analysis to estimate spillover effects over four different subsamples, distinguishing by whether the spouse is the primary or secondary earner and by gender. To avoid composition effects confounding our results, we reweight each primary-earner subsample so that it matches the distribution of the secondary-earner subsample of the same gender in terms of age differences.

For simplicity, we report results only for the retirement outcome.

We find that couples where males are the primary earner are more likely to retire jointly, consistent with the findings of Gustman and Steinmeier (2000) and Browning et al. (2021) that males value joint leisure more than women, and in further support of the collective model as an explanation of couples labor supply.²³ Interestingly, we find that the increased joint retirement among couples where males are the primary earner is driven equally by either males or females in these households adjusting their behavior to retire together. Specifically, we find that male spouses who are primary earners, reported in column (2), are much more likely to adjust their behavior to retire jointly than male spouses who are secondary earners, as reported in column (1). The scaled effect is 9.1% against 4.3%. Correspondingly, female spouses who are secondary earners are much more likely to adjust their behavior to retire together than female spouses who are primary earners, as we see from comparing column (3) to column (4), with scaled effects of 8.2% and 2.3% respectively.

The interaction between relative earnings and age differences. We now explore whether the interaction between relative earnings and age differences within couples affects their preferred path to joint retirement. Specifically, one might expect that older members who are primary earners are more likely to extend their employment while younger members who are secondary earners are more likely to retire earlier, consistent with the opportunity cost of retirement as foregone labor market earnings. We study this by replicating our analysis to estimate spillover effects over four different subsamples, distinguishing by whether the spouse is the primary or secondary earner and whether the spouse is the younger or the older member of the couple. To avoid composition effects confounding our results, we reweight the primary-earner subsamples so that they match the distribution of the secondary-earner subsamples in terms of age differences and gender.

The results are reported in Table 4. Overall, the primary-earner status does not seem to be a major determinant of joint retirement, as the differences between primary and secondary earner spouses are small and not statistically significant. However, interpreting the estimates at face value, we observe patterns consistent with the opportunity cost of retirement. We see that among older spouses, shown in columns (1) and (2), primary earners are 1.1 percentage points more likely to retire jointly. That is, they are more likely to work past their retirement age waiting for their younger spouses to reach their own pension eligibility age. On the

²³Note that the finding that males value joint leisure more than women can also be interpreted as males disliking some forms of independent leisure more than women, such as staying at home while their partners go to work.

contrary, among younger spouses, shown in columns (3) and (4), secondary earners are 2.7 percentage points more likely to retire jointly; that is they are more likely to stop working before they reach their own pension eligibility age to retire when their older partner becomes eligible. These results are consistent with the opportunity cost of retirement seen as foregone earnings. The returns to continued employment are higher for primary earners, who therefore are more likely to work longer, while the foregone earnings from secondary earners are smaller, making it less costly for them to stop working earlier.

6 Impact of Increasing Retirement Ages

We have shown that spouses are more likely to retire right when their partners reach pension eligibility age. What happens to the joint retirement of couples when the pension eligibility age of one partner changes? In this section we study a major reform that discontinuously increased the early pension eligibility age of selected cohorts.

6.1 The 2011 Pension Reform

In May of 2011 the Danish government announced a pension reform that discontinuously increased retirement ages in six-month increments contingent on birthdate. The first increase introduced by the reform provides us with the clearest natural experiment: the early pension eligibility age (that is, the VERP eligibility age) was increased from 60 to $60\frac{1}{2}$ for those born from January 1, 1954, while it remained at 60 for those born right before. The reform also introduced six-month increments in the incentivized early retirement age that was traditionally at age 62 and in the OAP age that was traditionally at age 65, but we maintain our original focus on the prominent early pension eligibility age.²⁴ Other characteristics of the VERP program remained the same, including the pension benefits and its independence between spouses. The duration of VERP remained 5 years in length because the OAP age increased as well.

The design of the VERP program, which we described in Section 2, creates strong incentives to retire right at the VERP eligibility age. Hence, the reform induced strong shifts in claiming and retirement ages of the affected individuals that we can use as a first stage to study spillover effects to their spouses. For more details on this reform and an analysis of how it affects savings of individuals directly affected see [García-Miralles and Leganza \(2021\)](#).

²⁴Cohorts born later than July 1, 1954 experienced additional increases in their pension eligibility ages that we illustrate in Appendix Figure B.1.

6.2 Reform-Based Discontinuity Design

To identify the casual effect of increasing individuals' pension eligibility age, we use a local difference-in-differences framework. The treatment group is composed of individuals born on January 1, 1954 or soon after, whose pension eligibility ages increase by 6 months due to the reform. The control group is composed of individuals born right before January 1, 1954, whose pension eligibility ages remain the same. In our main analysis we consider a bandwidth of three months around January 1, 1954 but we show that our results are robust to different bandwidth choices.

We asses the parallel-trends assumption and the dynamics around the announcement and implementation of the reform by estimating a dynamic difference-in-differences model over the period 2008–2014 of the form:

$$y_{it}^{(s)} = \alpha + \delta \cdot treat_i + \sum_{c \neq 2010} \kappa_c \cdot D_c + \sum_{c \neq 2010} \beta_c \cdot D_c \cdot treat_i + X'_{it} \cdot \psi + \epsilon_{it}, \quad (4)$$

where $y_{it}^{(s)}$ is the outcome variable of interest, either for the reference individuals (y_{it}) or for their spouses (y_{it}^s). D_c are calendar year dummies, and $treat_i$ is an indicator for individuals in the treatment group. The matrix X'_{it} is a set of controls that includes spousal age rounded to quarters interacted with gender when the model is estimated for spousal outcomes.

The results from these dynamic difference-in-differences (Figures 4 and 5) are discussed in detail in the next section. Note that to assess the parallel-trends assumption we must consider the pre-announcement period (2008–2010), for which we find no evidence of differential trends. During the period between announcement and implementation (2011–2013) treated individuals and their spouses could adjust behaviors in anticipation of reaching increased pension eligibility ages. However, we find no evidence of anticipatory responses for the reference individuals, nor for the spouses despite a slight change in the coefficient for 2013, the year just before implementation. Nevertheless, to be on the safe side we quantify the effects of the reform with respect to the pre-announcement period only, and show that our results are robust to including the anticipation period in the pre-period. Specifically, we estimate the following model to quantify the causal effects of the reform:

$$y_{it}^{(s)} = \beta_0 + \beta_1 \cdot treat_i + \beta_2 \cdot ant_t + \beta_3 \cdot post_t + \beta_4 \cdot treat_i \cdot ant_t + \beta_5 \cdot treat_i \cdot post_t + X'_{it} \cdot \psi + \epsilon_{it}, \quad (5)$$

where $y_{it}^{(s)}$ is the outcome variable of interest, either for the reference individual (y_{it}) or for their spouses (y_{it}^s), $treat_i$ is an indicator for individuals in the treatment group, ant_t is an

indicator for years in the anticipation period (2011–2013), $post_t$ is an indicator for implementation year 2014, and X'_{it} is a set of controls that includes spousal age rounded to quarters interacted with spousal gender. When this equation is estimated for the reference individual, the coefficient β_5 identifies the causal effect of the reform on the reference individuals (the first stage). When the equation is estimated for the spousal outcomes, the coefficient β_5 identifies the causal effect on the spouses (the reduced-form).

To obtain scaled spillover effects of the reform, which can be compared to the scale effects estimated in Section 4, we estimate a 2SLS model where the retirement outcomes of the reference individual are instrumented by their treatment status interacted with the calendar year where the reform directly affects them ($treat_i \cdot post_t$). The first stage of the 2SLS model corresponds to equation (5) when it is estimated for the reference individual's outcomes. The second-stage equation is the following:

$$y_{it}^s = \beta_0 + \beta_1 \cdot \hat{y}_{it} + \beta_2 \cdot treat_i + \beta_3 \cdot ant_t + \beta_4 \cdot post_t + \beta_5 \cdot treat_i \cdot ant_t + X'_{it}\psi + u_{it}, \quad (6)$$

where \hat{y}_{it} is the predicted outcome for the reference individual estimated in the first-stage and the coefficient β_1 identifies the scaled spillover effect.

6.3 The Effect of Increasing the Pension Eligibility Age on Own Retirement

The reform induced a strong response from individuals directly affected by the increase in pension eligibility ages. Figure 4 shows the results of the dynamic difference-in-differences model on the retirement outcomes of individuals directly affected by the reform. We confirm that the behavior of the treated and control groups along the three outcomes considered is similar during the period before announcement (2008–2010) as well as before implementation of the reform (2011–2013). The trends of both groups move in parallel and we can rule out any significant anticipatory response.

During the implementation year of 2014, individuals in the treatment group respond to the reform by delaying retirement, consistent with the strong incentives built into the VERP program. Panel (a) of Figure 4 shows that individuals in the treatment group are around 19 percentage points less likely to retire during the first half of the year. Note that the reform increased the pension eligibility age by 6 months and hence we define retirement as stopping to work during the first half of the year, as explained in Section 3. Individuals affected by the reform are also 26 percentage points less likely to claim benefits, and have higher annual labor market earnings, around \$8,140, during the implementation year. In the first row of

Table 5 we report estimates from the pooled difference-in-differences model, which quantify the large and significant effect of the reform on individuals directly affected, providing a strong first stage to analyze spillover effects to spouses.

6.4 The Effect of Increasing the Pension Eligibility Age on Spouses

We now study the effect of the reform on spousal retirement behavior. Figure 5 reports the dynamic effects. In the period preceding the announcement of the reform (2008–2010), spouses of both treatment and control individuals behave similarly, providing evidence in support of the parallel trends assumption. After announcement and before implementation (2011–2013), no coefficient is significantly different from zero, suggesting that spouses do not respond differentially in anticipation of their partners’ increased pension eligibility age, in line with the lack of anticipation of the reference individuals who are themselves affected directly by the reform.²⁵ In the implementation year, 2014, we observe that spouses of individuals who are affected by the reform are induced to delay their retirement, consistent with extending employment in order to retire jointly with their partner. We find evidence of spouses adjusting their behavior along the other two margins as well; spouses claim their pensions later and increase their annual earnings.

The second row of Table 5 reports the difference-in-differences estimates that quantify these spousal effects. The estimates are statistically significant for retirement and claiming, but not for earnings. We report scaled effects from the 2SLS model in the third row of Table 5. The scaled effect on retirement is 9%, indicating that for every 100 individuals who postpone their retirement due to the reform, around 9 spouses will delay their own retirement to make it coincide with that of their partner. The spillover in claiming is 4.2% and the spillover in earnings is 8.5%, although the later is not statistically significant.

Overall, our findings show that the reform induced similar spillover effects as the ones we estimated in a stable context where pension eligibility ages did not change and were known by the couples well in advance. These results are consistent with a lack of significant frictions that could prevent couples from adjusting their behavior to retire jointly. This may be of particular interest to policy makers trying to predict short-run responses of social security reforms based on estimates from stable settings. Conversely, it helps with interpreting other

²⁵Although we do not find evidence of anticipatory responses from spouses, we do observe that in 2013, the year just before implementation, the coefficients tend to move slightly, perhaps suggesting a mild, and not significant, anticipatory response by spouses. This is the reason why in our main model specification to quantify the effect of the reform (equations 5 and 6) we include an indicator variable for the years between announcement and anticipation of the reform.

reform-based estimates in the literature, as it shows that couples' joint retirement behavior can adjust relatively quickly to changes that affect the retirement age of one partner.

We also explore heterogeneity in responses to the reform. Despite the relatively large sample size of our reform-based design (a panel of 10,321 individuals), we are unable to explore mechanisms in as much depth as in our age-based design, where we estimated effects on reweighted samples and from more granular sample splits. However, we are able to replicate our reform-based analysis on a simple sample split by age differences and gender and the results go in the same direction as our age-based analysis, with results reported in Appendix Table B.1. We find that in response to the reform, older spouses respond the most, with a 12% spillover in retirement against 3% for the younger spouses. A simple gender split returns similar spillover estimates for male and female spouses (9% and 9.1% respectively) as was the case in the age-based design. This suggest again that female spouses respond more once we account for the fact that females are most often the younger member of the couple (around 1.8 years younger in this reform-based analysis sample).

6.5 Threats to Identification and Robustness

Identifying assumption. The validity of our empirical approach relies on the assumption that in the absence of the reform, spousal outcomes of the treated and control individuals would move in parallel across time. We already showed that trends are parallel in the years preceding the implementation of the reform. However, in interpreting our outcomes as causal, we also assume that spouse behaviors differ in 2014 only because their partners are differentially affected by the reform. A violation of this assumption occurs if the spouses themselves are directly, and *differentially*, impacted by the same reform.

By construction, treated individuals are 3 months older on average than control individuals, and so are their spouses. Therefore, because the reform affects individuals based on their birth date, older spouses are more likely to be directly impacted by the reform themselves. In this section we show that the differential impact of the reform on the spouses is small and that our results are robust to a series of tests that address this concern.

First, note that only spouses born during the first 6 months of 1954 are affected by the reform that increases their eligibility age from 60 to $60\frac{1}{2}$ and impacts them in 2014. In Appendix Figure B.2 we plot the distribution of spouses' birth dates and show that spouses of treated individuals are only 1.3 percentage points more likely to be born during those 6 months than spouses of control individuals (6.5% against 5.2%). To ensure that our results are not driven by this difference, we do the following two tests. First, we replicate our

analysis reweighting the sample of treated individuals so that they have the same distribution of spousal date of birth as the control group. Second, we replicate the analysis excluding individuals whose spouses are born in the first half of 1954, both from the treatment and control groups. The results are reported in rows B and C of Table B.2 and are very similar to the baseline results.

We also note that spouses born after July 1, 1954 are affected by the reform by experiencing larger increases in their pension eligibility ages (as illustrated by Appendix Figure B.1), but these increases only affect them directly after 2014, and we do not include those years in our analysis. Spouses in the control group are 2.2 percentage points more likely to be born after July 1, 1954 (44.3% against 42.1%). Importantly, this differential impact of the reform on the spouses would only affect our results if the reference individuals or their spouses responded in anticipation to future changes in their pension eligibility age. We address this concern in two ways. First, we note that across our analyses, we do not find evidence of significant anticipatory responses (see Figures 4 and 5). Second, we replicate our analysis for the subsample of individuals whose spouses are more than 3 months older. This subsample ensures that all spouses are born before January 1, 1954 and therefore are totally unaffected by the reform. The results, reported in Appendix Table B.3 show even larger spillover effects. This is to be expected, as we have shown earlier that older spouses are the ones that respond the most. Overall, these tests make us confident that the small share of spouses who are differentially impacted by the reform do not have a substantive impact on our results.

Robustness. We perform a series of robustness tests including changes to the model specification and to the sample definition. Table B.2 shows the results. Row A shows the baseline results for comparison. Rows B and C were introduced earlier. Row D shows that the results are unaffected by estimating the model without the anticipation variable. Rows E and F report the results from decreasing and increasing the bandwidth around the cutoff date of January 1, 1954 by two weeks. Row G shows the results when we do not balance the sample of analysis. Row H adds controls for region and education of the reference individuals and their spouses, defined when they are 57 years old. Finally, row I extends the sample to include reference individuals who did not contribute to the VERP program between ages 50-59. Overall our results are robust to all these changes. We note, however, that although the size of the estimates for claiming remains stable, they turn insignificant in some cases. The estimates for earnings remain insignificant in most cases.

7 Conclusion

We have shown that spouses adjust their behavior to retire together, which implies a significant role for leisure complementarities within couples. Using two different discontinuity designs, we estimate sizable spillover effects from individuals' early pension eligibility age to their spouses of around 8%. This means that for every 100 individuals who retire upon reaching their pension eligibility age, 8 of their spouses adjust their behavior to retire at the same time. We also find spillover effects to claiming of pension benefits and to labor market earnings.

Our analysis of mechanisms shows that age differences between couples are a crucial determinant of joint retirement behavior. We find that older spouses work past their own pension eligibility age while waiting for their younger partner to become eligible for pension benefits as well. Accounting for age differences is important. After accounting for age differences we uncover a strong gender gap, which is not explained by relative earnings, suggesting that gender norms play a role. Accounting for both age differences and gender, we find patterns consistent with a collective model of household behavior and patterns consistent with couples considering the opportunity cost of retirement.

Our results have implications for policy design and models of intra-household decision-making. The existence of sizable joint retirement spillovers informs policy makers, who should account for these spillovers when designing policies and projecting their effects. Moreover, policy makers should account for the demographic characteristics of individuals affected by social security policies, in particular the age composition of couples. This is particularly relevant for countries that are still increasing eligibility ages of females to converge to those of men, as females tend to be the younger member of couples and therefore induce the largest spillovers.


References

- An, M. Y., Christensen, B. J., and Gupta, N. D. (2004). Multivariate mixed proportional hazard modelling of the joint retirement of married couples. *Journal of Applied Econometrics*, 19(6):687–704.
- Andrews, D. W. and Buchinsky, M. (2000). A three-step method for choosing the number of bootstrap repetitions. *Econometrica*, 68(1):23–51.
- Atalay, K., Barrett, G. F., and Siminski, P. (2019). Pension incentives and the joint retirement of couples: Evidence from two natural experiments. *Journal of Population Economics*, 32(3):735–767.
- Baker, M. (2002). The retirement behavior of married couples: Evidence from the spouse’s allowance. *The Journal of Human Resources*, 37(1):1–34.
- Banks, J., Blundell, R., and Rivas, M. C. (2010). The dynamics of retirement behavior in couples: Reduced-form evidence from England and the US . *University College London, mimeo*.
- Behaghel, L. and Blau, D. M. (2012). Framing social security reform: Behavioral responses to changes in the full retirement age. *American Economic Journal: Economic Policy*, 4(4):41–67.
- Bingley, P. and Lanot, G. (2007). Public pension programmes and the retirement of married couples in Denmark. *Journal of Public Economics*, 91(10):1878–1901.
- Blau, D. M. (1998). Labor force dynamics of older married couples. *Journal of Labor Economics*, 16(3):595–629.
- Bloemen, H., Hochguertel, S., and Zweerink, J. (2019). The effect of incentive-induced retirement on spousal retirement rates: Evidence from a natural experiment. *Economic Inquiry*, 57(2):910–930.
- Browning, M., Bourguignon, F., Chiappori, P.-A., and Lechene, V. (1994). Income and outcomes: a structural model of intrahousehold allocation. *Journal of Political Economy*, 102(6):1067–1096.
- Browning, M. and Chiappori, P.-A. (1998). Efficient intra-household allocations: a general characterization and empirical tests. *Econometrica*, pages 1241–1278.
- Browning, M., Donni, O., and Gørtz, M. (2021). Do you have time to take a walk together? private and joint time within the household. *The Economic Journal*, 131(635):1051–1080.
- Casanova, M. (2010). Happy together: A structural model of couples’ joint retirement choices. *Working Paper, Department of Economics, University of California*, (5):1–53.
- Chiappori, P.-A. (1992). Collective labor supply and welfare. *Journal of Political Economy*, 100(3):437–467.
- Coile, C. (2004). Retirement incentives and couples’ retirement decisions. *Topics in Economic Analysis and Policy*, 4(1):441–470.

- Coile, C. and Gruber, J. (2007). Future social security entitlements and the retirement decision. *The Review of Economics and Statistics*, 89(2):234–246.
- Cribb, J., Emmerson, C., and Tetlow, G. (2016). Signals matter? Large retirement responses to limited financial incentives. *Labour Economics*, 42:203–212.
- Daly, M. and Groes, F. (2017). Who takes the child to the doctor? Mom, pretty much all of the time. *Applied Economics Letters*, 24(17):1267–1276.
- Deshpande, M., Fadlon, I., and Gray, C. (2020). How sticky is retirement behavior in the U.S.? Responses to changes in the full retirement age. *NBER Working Paper No. w27190*.
- Donni, O. and Chiappori, P.-A. (2011). Nonunitary models of household behavior: a survey of the literature. *Household Economic Behaviors*, pages 1–40.
- Fadlon, I., Ramnath, S. P., and Tong, P. K. (2019). Market inefficiency and household labor supply: Evidence from social security’s survivors benefits. *NBER Working Paper No. w25586*.
- García-Mirallas, E. and Leganza, J. M. (2021). Public pensions and private savings. *Working Paper 06/21. The Center for Economic Behavior and Inequality (CEBI). University of Copenhagen*.
- Geyer, J. and Welteke, C. (2019). Closing Routes to Retirement for Women: How Do They Respond? *Journal of Human Resources*.
- Gørtz, M., Sander, S., and Sevilla, A. (2020). Does the child penalty strike twice?
- Gustman, A. L. and Steinmeier, T. L. (2000). Retirement in dual-career families: A structural model. *Journal of Labor Economics*, 18(3):503–545.
- Gustman, A. L. and Steinmeier, T. L. (2004). Social security, pensions and retirement behaviour within the family. *Journal of Applied Econometrics*, 19(6):723–737.
- Haller, A. (2019). Welfare effects of pension reforms.
- Honoré, B. E., Jørgensen, T. H., and Paula, Á. (2020). The Informativeness of Estimation Moments. *Journal of Applied Econometrics*.
- Hospido, L. and Zamarro, G. (2014). Retirement patterns of couples in Europe. *IZA Journal of European Labor Studies*, 3(1):12.
- Hurd, M. D. (1990). The joint retirement decision of husbands and wives. In *Issues in the Economics of Aging*, pages 231–258. University of Chicago Press, 1990.
- Jia, Z. (2005). Labor supply of retiring couples and heterogeneity in household decision-making structure. *Review of Economics of the Household*, 3(2):215–233.
- Kleven, H., Landais, C., and Søgaaard, J. E. (2019). Children and gender inequality: Evidence from Denmark. *American Economic Journal: Applied Economics*, 11(4):181–209.
- Kreiner, C. T., Leth-Petersen, S., and Skov, P. E. (2016). Tax reforms and intertemporal shifting of wage income: Evidence from Danish monthly payroll records. *American Economic Journal: Economic Policy*, 8(3):233–57.

- Kruse, H. (2020). Joint retirement in couples: Evidence of complementarity in leisure. *The Scandinavian Journal of Economics*.
- Lalive, R. and Parrotta, P. (2017). How does pension eligibility affect labor supply in couples? *Labour Economics*, 46:177–188.
- Lassen, A. S. (2020). Gender norms and specialization in household production: Evidence from a danish parental leave reform.
- MacKinnon, J. G. (2006). Bootstrap methods in econometrics. *Economic Record*, 82:S2–S18.
- Manoli, D. S. and Weber, A. (2016). The effects of the early retirement age on retirement decisions. *NBER Working Paper No. w22561*.
- Mastrobuoni, G. (2009). Labor supply effects of the recent social security benefit cuts: Empirical estimates using cohort discontinuities. *Journal of Public Economics*, 93(11-12):1224–1233.
- Michaud, P.-C. and Vermeulen, F. (2011). A collective labor supply model with complementarities in leisure: Identification and estimation by means of panel data. *Labour Economics*, 18(2):159–167.
- Nakazawa, N. (2019). The effects of increasing the eligibility age for public pension on individual labor supply: Evidence from Japan. *Available at SSRN 3438100*.
- Nielsen, N. F. (2019). Sick of retirement? *Journal of Health Economics*, 65:133–152.
- OECD (2015). *Pensions at a Glance: OECD and G20 Indicators*. OECD.
- OECD (2017). *Pensions at a Glance: OECD and G20 Indicators*. OECD.
- OECD (2019). *Pensions at a Glance: OECD and G20 Indicators*. OECD.
- Selin, H. (2017). What happens to the husband’s retirement decision when the wife’s retirement incentives change? *International Tax and Public Finance*, 24(3):432–458.
- Stancanelli, E. (2017). Couples’ retirement under individual pension design: A regression discontinuity study for France. *Labour Economics*, 49:14–26.
- Stancanelli, E. and Soest, A. V. (2012). Retirement and Home Production: A Regression Discontinuity Approach. *American Economic Review*, 102(3):600–605.
- Staubli, S. and Zweimüller, J. (2013). Does raising the early retirement age increase employment of older workers? *Journal of Public Economics*, 108:17–32.
- Van der Klaauw, W. and Wolpin, K. I. (2008). Social security and the retirement and savings behavior of low-income households. *Journal of Econometrics*, 145(1-2):21–42.
- Willén, A., Vaage, K., and Johnsen, J. (forthcoming). Interactions in public policies: spousal responses and program spillovers of welfare reforms. *The Economic Journal*.
- Zweimüller, J., Winter-Ebmer, R., and Falkinger, J. (1996). Retirement of spouses and social security reform. *European Economic Review*, 40(2):449–472.


Figure 1: The Effect of Reaching Pension Eligibility Age on Own Retirement




Notes: These figures plot different outcomes for individuals around their own pension eligibility age of 60, pooling individuals over the period 1991–2013. The hollow circles are raw means of the outcome variable measured at the end of each calendar year, grouped in monthly age bins. The solid lines plot the parametric fit estimated with the piecewise linear regression model (1). The dashed line represents the counterfactual behavior in the absence of pension eligibility, based on a linear extrapolation from the observed outcome before age 60. The full-exposure effect of being eligible for early retirement pension during an entire year is represented by the vertical distance between the solid and dashed lines just below age 61.

Figure 2: The Effect of Reaching Pension Eligibility Age on Spouses


(a) Spouse Retirement



(b) Spouse Claiming




(c) Spouse Earnings




Notes: These figures plot different outcomes for spouses around the pension eligibility age of their partner. The dots are the residuals estimated in equation (2) where spousal outcomes are regressed on spousal age and gender. The residuals are grouped in monthly bins of the reference individual's age. The solid lines plot the parametric fit estimated with the piecewise linear regression model (3). The dashed line represents the counterfactual behavior in the absence of pension eligibility, based on a linear extrapolation from the observed outcome before age 60. The full-exposure effect on the spouses of their partners being eligible for early retirement pension during an entire year is represented by the vertical distance between the solid and dashed lines just below age 61.

Figure 3: Joint Retirement Behavior by Age Differences Within Couples




Notes: These figures plot the scaled spillover estimates of joint retirement for different subsamples of couples based on the age difference between spouses. These scaled effects are estimated using the same methodology as for the full sample: first estimating models (1) and (3) to obtain full-exposure effects and then dividing the full-exposure effect on spouses by the full-exposure effect on reference individuals. We report 95% confidence intervals calculated from bootstrapped standard errors.

Figure 4: The Effect of Increasing Pension Eligibility Age on Own Retirement



Notes: These figures plot the β_c coefficients from the dynamic difference-in-differences model (4), estimated on different outcomes for reference individuals. Each coefficient shows the difference between the treated group (whose pension eligibility age increases by 6 months, to age $60\frac{1}{2}$) and the control group (whose pension eligibility age remains at age 60). Individuals turn 60 around the beginning of 2014, therefore the coefficient for 2014 identifies the causal effect of the reform during the implementation year. We report confidence intervals at the 95% level, calculated from robust standard errors clustered at the couple level.

Figure 5: The Effect of Increasing Pension Eligibility Age on Spouses



Notes: These figures plot the β_c coefficients from the dynamic difference-in-differences model (4), estimated on different outcomes for spouses of reference individuals. Each coefficient shows the difference between the treatment group (spouses whose partners' pension eligibility age increases by 6 months, to age $60\frac{1}{2}$) and the control group (spouses whose partners' pension eligibility age remains at 60). The coefficient for 2014 identifies the causal effect of the reform on the spouses on the implementation year. We report confidence intervals at the 95% level, calculated from robust standard errors clustered at the couple level.

Table 1: The Effect of Reaching Pension Eligibility Age

	Retirement	Claiming	Earnings
Reference Individual	0.2034*** (0.001)	0.3496*** (0.001)	-8,642*** (69.431)
Spouse	0.0153*** (0.001)	0.0120*** (0.001)	-848*** (61.165)
Scaled Effect	0.0750*** (0.007)	0.0344*** (0.003)	0.0981*** (0.012)
N. of clusters	367,585	367,585	367,585
Observations	2,206,044	2,206,044	2,206,044

Notes: This table reports the effect of reference individuals reaching pension eligibility age on their own retirement and on their spouses' retirement. Each column reports the results for a different outcome. The first row reports the full-exposure effect to pension eligibility on own retirement estimated in equation (1). The second row reports the full-exposure effect on the spouses from their partners becoming eligible for pension, estimated in equation (3). The third row reports the scaled effect resulting from dividing the spouse full-exposure effect by the reference individual full-exposure effect. Robust standard errors in parentheses, clustered at the couple level. Bootstrapped standard errors for scaled effects. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 2: Robustness to Alternative Specifications
for the Effect of Reaching Pension Eligibility Age

	Retirement	Claiming	Earnings
A. Baseline	0.0750*** (0.0071)	0.0344*** (0.0031)	0.0981*** (0.012)
B. Including Younger Ages	0.0718*** (0.0059)	0.0363*** (0.0028)	0.0928*** (0.010)
C. Excluding Age 59	0.0855*** (0.0097)	0.0389*** (0.0044)	0.117*** (0.017)
D. Unrestricted Age Difference	0.0690*** (0.0066)	0.0311*** (0.0029)	0.0905*** (0.012)
E. No Donut December	0.0736*** (0.0069)	0.0327*** (0.0030)	0.0987*** (0.012)
F. Nonlinear Counterfactual	0.0498*** (0.016)	0.0300*** (0.0055)	0.0580** (0.026)
G. Nonlinear & Incl. Younger	0.0761*** (0.0099)	0.0328*** (0.0037)	0.0972*** (0.017)
H. Adding Controls	0.0749*** (0.0070)	0.0340*** (0.0031)	0.0922*** (0.012)
I. Dummy 1999 Reform	0.0746*** (0.0070)	0.0344*** (0.0032)	0.0978*** (0.012)
J. Period 2008–2013	0.0758*** (0.012)	0.0459*** (0.0068)	0.104*** (0.025)
K. 2008–2013 & VERP Eligible	0.0674*** (0.011)	0.0434*** (0.0063)	0.104*** (0.024)
L. Retirement Flow Variable	0.0573*** (0.0053)	–	–

Notes: This table reports the scaled effect estimates from replicating our main analysis over different sample definitions and over different specifications of the estimation models (equations 1 and 3). Row A reproduces results from our baseline specification, which correspond to those reported in Table 1. Row B replicates the analysis over a sample extended to include reference individuals of ages 55 and 56. Row C excludes reference individuals aged 59. Row D keeps couples with partners that are more than 8 years apart from each other. Row E keeps reference individuals who turn 60 in December. Row F allows for a nonlinear counterfactual by adding a second order polynomial. Row G implements the two changes applied in B and F. Row H controls for predetermined region and education of reference individuals and spouses. Row I adds a dummy for individuals born after 1939, who are affected by the 1999 reform. Row J estimates the effect over the period 2008–2013. Row K estimates the effect over the same period as J and restricts the sample to reference individuals who have contributed to VERP at least once between ages 50 and 59. Row L reports the estimate for retirement defined as a flow variable, allowing individuals to retire multiple times. Bootstrapped standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 3: Heterogeneity in the Effect of Reaching Pension Eligibility Age on Retirement by Age Difference and Gender

Reference Individual	Young	Old	Female	Male	Male (w)	Male (w)
Spouse	Old	Young	Male	Female	Female (w)	Female (w)
	(1)	(2)	(3)	(4)	(5)	(6)
A. Retirement						
Reference Individual	0.2562*** (0.002)	0.1588*** (0.002)	0.2668*** (0.002)	0.1479*** (0.001)	0.1616*** (0.002)	0.1795*** (0.003)
Spouse	0.026*** (0.002)	0.004*** (0.001)	0.020*** (0.001)	0.011*** (0.002)	0.021*** (0.002)	0.024*** (0.003)
Scaled Effect	0.0994*** (0.010)	0.0287*** (0.011)	0.0745*** (0.008)	0.0751*** (0.013)	0.130*** (0.017)	0.132*** (0.019)
B. Claiming						
Reference Individual	0.4307*** (0.002)	0.28*** (0.002)	0.4567*** (0.002)	0.2544*** (0.002)	0.2632*** (0.002)	0.295*** (0.003)
Spouse	0.021*** (0.002)	0.000 (0.000)	0.017*** (0.001)	0.008*** (0.001)	0.018*** (0.002)	0.020*** (0.003)
Scaled Effect	0.0495*** (0.006)	0.0035*** (0.001)	0.0374*** (0.004)	0.0301*** (0.004)	0.0674*** (0.009)	0.0683*** (0.011)
C. Earnings						
Reference Individual	-9,558*** (93.657)	-7,970*** (97.971)	-9,081*** (81.417)	-8,408*** (104.024)	-9,035*** (140.987)	-9,160*** (162.506)
Spouse	-1,763*** (111.106)	-490*** (75.159)	-1,105*** (106.718)	-570*** (64.957)	-589*** (97.229)	-726*** (150.946)
Scaled Effect	0.186*** (0.020)	0.0633*** (0.017)	0.122*** (0.020)	0.0678*** (0.014)	0.0661*** (0.018)	0.0802*** (0.026)
N. of clusters	297,686	334,966	302,589	330,172	330,172	330,172
Observations	1,038,096	1,167,948	1,054,359	1,151,685	1,151,685	1,151,685

Notes: This table reports the effect of the reference individuals reaching pension eligibility age on their own retirement and on their spouses' retirement, distinguishing heterogeneous responses by gender and age differences within the couple. Each column shows results for a different subsample. The subsample in column (5) is reweighted to have the same distribution of age differences as the subsample from column (3) and the subsample in column (6) is further reweighted to have the same distribution of earnings shares as (3). Each panel reports results for a different outcome variable. Within each panel, the first row reports the full-exposure effect of pension eligibility on own retirement. The second row reports the full-exposure effect on spouses of their partners being eligible for retirement pension. The third row reports the scaled effect resulting from dividing the spouse full-exposure effect by the reference individual full-exposure effect. Robust standard errors in parentheses, clustered at the couple level. Bootstrapped standard errors for scaled effects. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 4: Heterogeneity in the Effect of Reaching Pension Eligibility Age on Retirement by Relative Earnings

A. By Gender				
Reference Individual	Female Primary	Female Sec. (w)	Male Primary	Male Secondary (w)
Spouse	Male Secondary	Male Primary (w)	Female Secondary	Female Primary (w)
	(1)	(2)	(3)	(4)
Reference Individual	0.2475*** (0.004)	0.2745*** (0.002)	0.1434*** (0.002)	0.1426*** (0.004)
Spouse	0.0111*** (0.003)	0.025*** (0.002)	0.0117*** (0.001)	0.003 (0.003)
Scaled Effect	0.0434** (0.021)	0.0909*** (0.011)	0.0816*** (0.016)	0.0225 (0.028)
N. of clusters	58,311	201,541	229,321	53,949
Observations	191,681	713,870	800,843	185,860
B. By Age Differences				
Reference Individual	Young Primary	Young Sec. (w)	Old Primary	Old Second. (w)
Spouse	Old Secondary	Old Prim. (w)	Young Secondary	Young Prim. (w)
	(1)	(2)	(3)	(4)
Reference Individual	0.2096*** (0.003)	0.2651*** (0.003)	0.1413*** (0.002)	0.1541*** (0.004)
Spouse	0.0198*** (0.003)	0.028*** (0.002)	0.0072*** (0.002)	0.004 (0.002)
Scaled Effect	0.0931*** (0.022)	0.104*** (0.012)	0.0535*** (0.018)	0.0265 (0.024)
N. of clusters	94,735	161,573	193,106	93,917
Observations	321,527	571,978	670,997	327,752

Notes: The table reports the effect of the reference individuals reaching pension eligibility age on their own retirement and on their spouses' retirement, distinguishing heterogeneous responses by primary earner status within the couple. Panel A further distinguish by gender and Panel B by age differences. Each column contains results for a subsample of the population. In Panel A, the subsamples in columns (2) and (4) are reweighed to have the same distribution of age differences as columns (1) and (3), respectively. In Panel B the subsamples in columns (2) and (4) are reweighed to have the same distribution of gender and age differences as columns (1) and (3), respectively. Within each panel, the first row reports the full-exposure effect of pension eligibility on own retirement. The second row reports the full exposure-effect on spouses of their partners being eligible for retirement pension. The third row reports the scaled effect resulting from dividing the spouse full-exposure effect by the reference individual full-exposure effect. Robust standard errors in parentheses, clustered at the couple level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

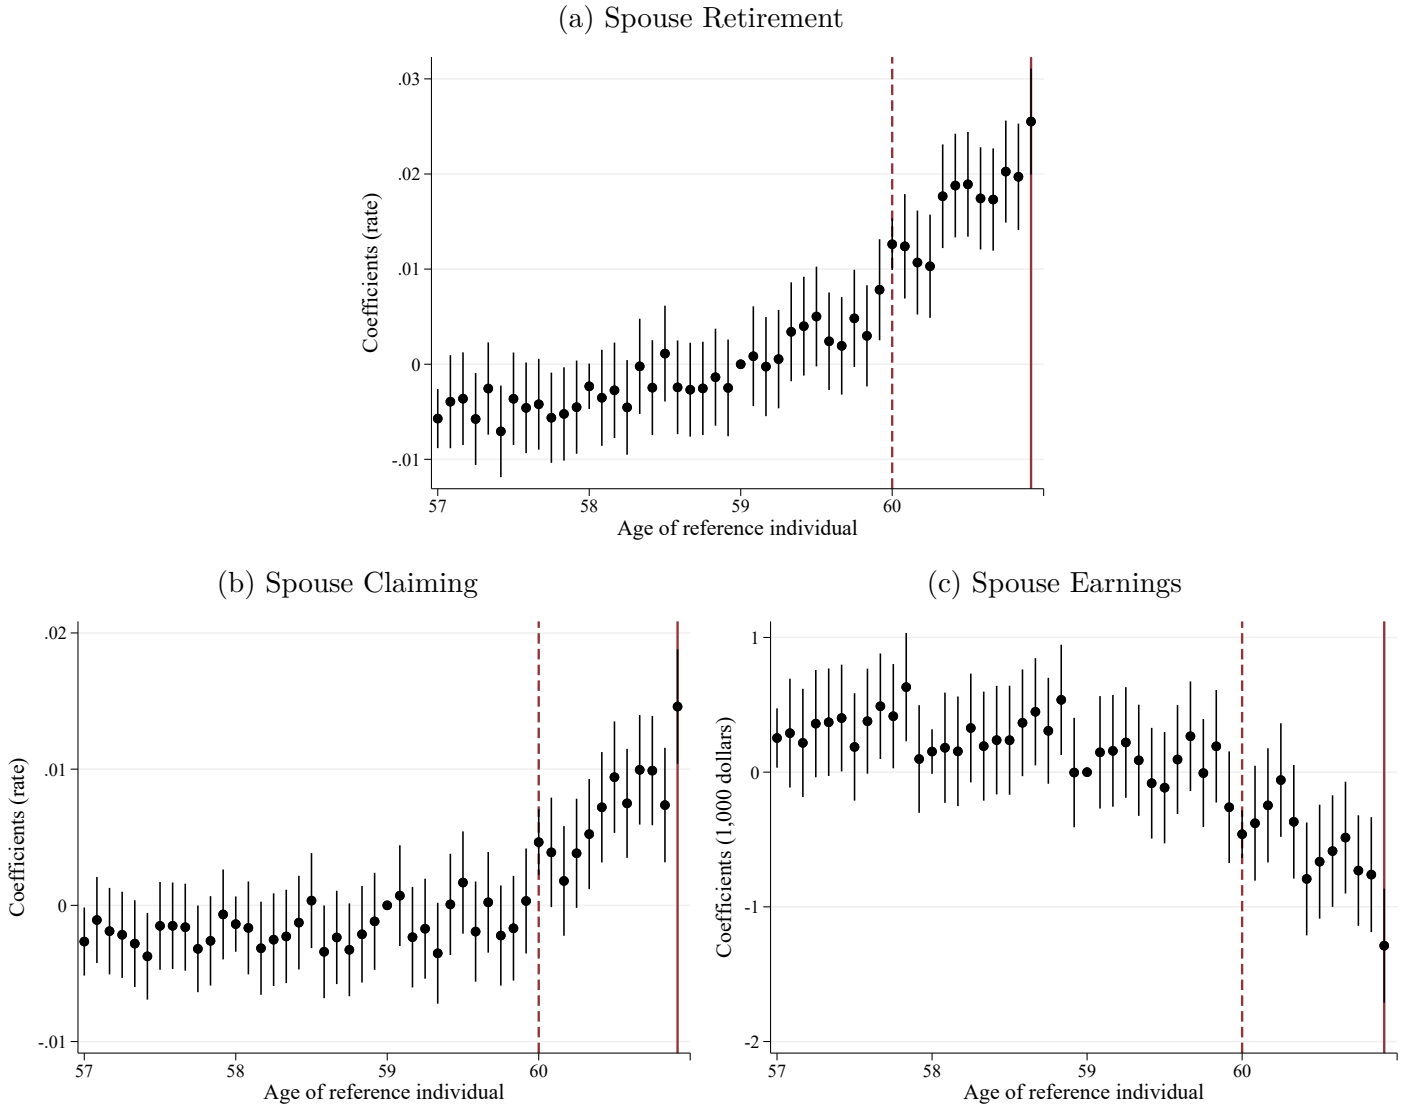
Table 5: The Effect of the Reform Increasing Pension Eligibility Age

	Retirement	Claiming	Earnings
Reference Individ.	-0.191*** (0.007)	-0.264*** (0.007)	8,140*** (479.0)
Spouse	-0.0172** (0.007)	-0.0110* (0.006)	690 (532.1)
Scaled Effect	0.0902** (0.038)	0.0418* (0.024)	0.0847 (0.065)
F-test instr.	662.3	1,643.6	288.8
N. of clusters	10,321	10,321	10,321
Observations	73,395	73,395	73,395

Notes: This table reports the effect of the 2011 reform, which increased the pension eligibility age. Each column reports results for a different outcome. The first row reports the effect on the individuals affected by the reform (the first stage) and the second row reports the spillover effect to their spouses (the reduced-form effect), which are estimated using equation (5). The third row reports the scaled effect resulting from the 2SLS model estimated in equation (6). Robust standard errors in parentheses, clustered at the couple level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Appendix A Age Discontinuity Design


Figure A.1: Alternative Graphical Evidence of the Effect of Pension Eligibility Age on Spouses




Notes: These figures show an alternative approach to obtain nonparametric evidence on spouses behavior around the pension eligibility age of their partner. They plot the δ_a^r coefficients from estimating the regression $y_{it}^s = \alpha + \sum_{a=57}^{62} \delta_a^r \cdot D_a^r + \sum_{a=49}^{69} \delta_a^s \cdot D_a^s + \sum_{a=49}^{69} \gamma_a \cdot D_a^s \cdot D_g + \sum_{c=1991}^{2013} \kappa_c \cdot D_c + \epsilon_{st}$, where y_{it}^s are the different outcomes plotted in each figure, D_a^r are age dummies for the reference individual, D_a^s are age dummies for the spouse, D_g is a gender dummy for the spouse, and D_c are calendar year dummies.

Figure A.2: Placebo Test Assigning Fake Spouses of Similar Age for the Effect of Reaching Pension Eligibility Age


(a) Spouse Retirement



(b) Spouse Claiming




(c) Spouse Earnings



Notes: These figures plot results from replicating the analysis over a placebo sample where the reference individuals are the same as in the main analysis, but they are matched to fake spouses of similar age. The figures show no evidence of joint retirement, as is expected if the research design is valid: fake spouses cannot affect each other's retirement behavior, and the effect coming from the correlation between their ages is controlled for by the empirical design. For more details on the construction of this figure, see the notes of Figure 2. See Appendix Table A.4 for the placebo point estimates.


Figure A.3: The Effect of Reaching Pension Eligibility Age on Retirement Defined as Flow




Notes: These figures plot an alternative definition of the retirement outcome, defined as a flow variable that takes the value one in the year in which an individual retires and zero otherwise. For more details on the construction of these figures see notes of Figures 1 and 2. The scaled effect estimate resulting from this outcome is reported in Table 2.

Figure A.4: The Evolution of Joint Retirement Over Time


(a) Retirement



(b) Claiming

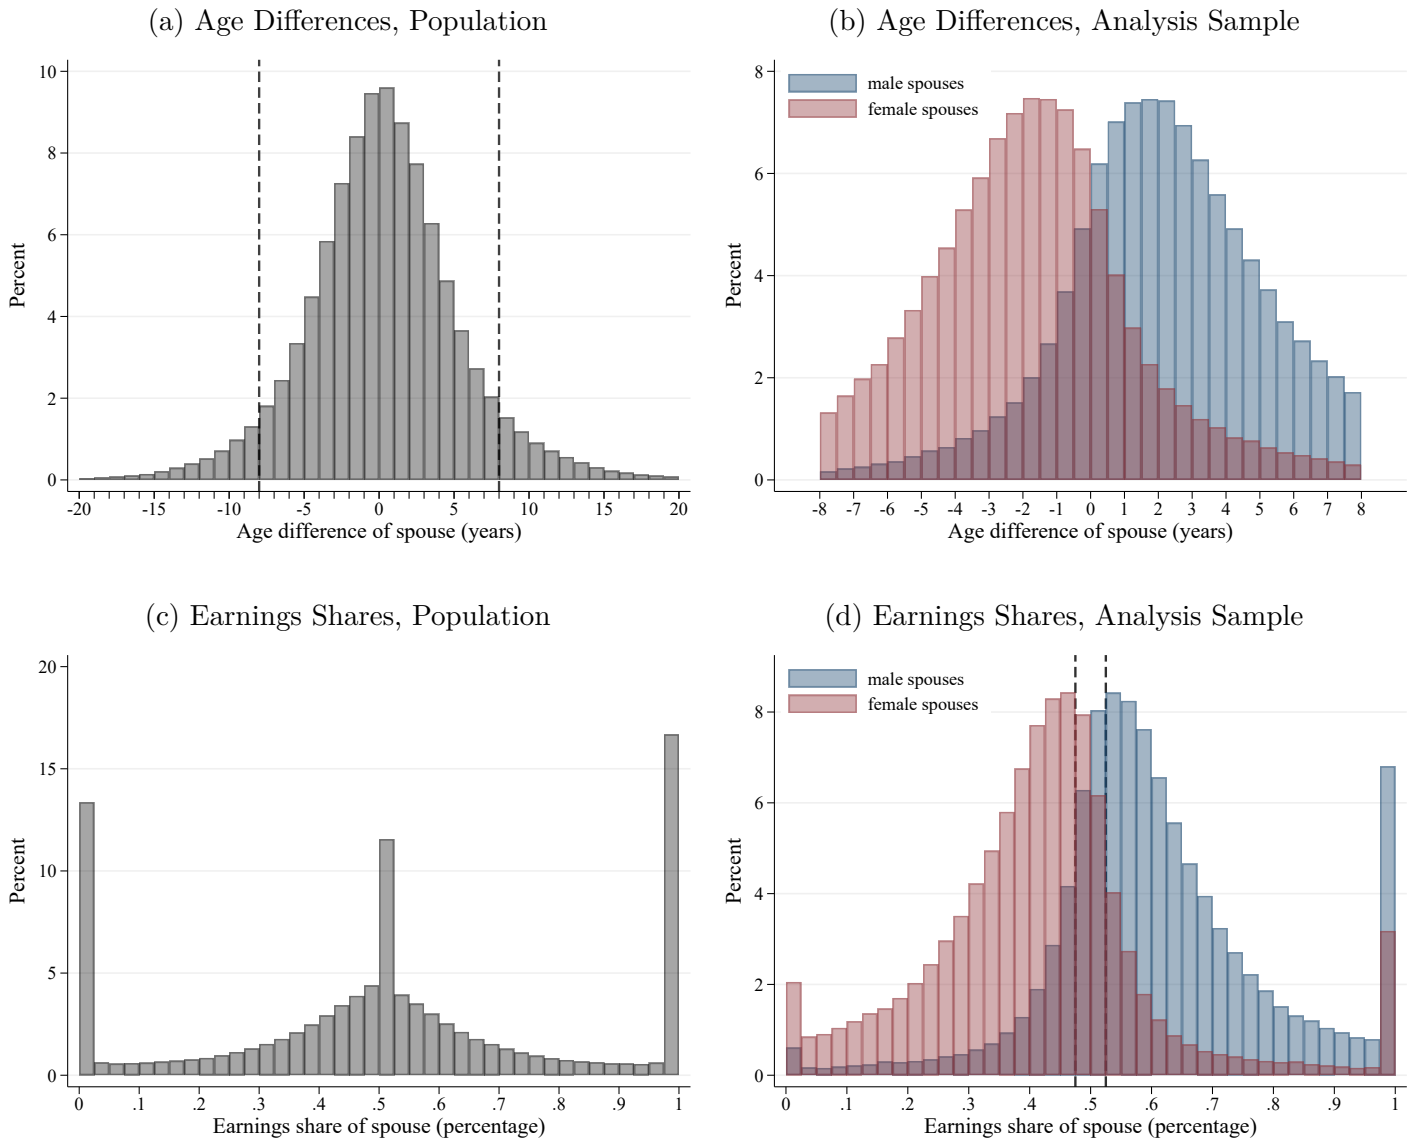


(c) Earnings



Notes: These figures plot the evolution over time of the scaled estimates of joint retirement for different outcomes. Scaled effects are estimated over a 5-year running window using the same methodology as for the full time period: first estimating models (1) and (3) to obtain full-exposure effects and then dividing the full-exposure effect on spouses by the full-exposure effect on reference individuals. We report 95% confidence intervals calculated from bootstrapped standard errors.

Figure A.5: Distribution of Spouses' Age Differences and Earnings Shares



Notes: Panel (a) plots the distribution of age differences within spouses for the population of Danish couples between 1991 and 2013, before applying any sample restrictions. The vertical dashed lines mark the tails that are excluded from the sample of analysis, corresponding to couples with more than 8 years difference in age. Panel (b) plots the distribution of age differences for the age-based sample of analysis resulting from imposing the restrictions described in Section 3.3. Panel (c) plots the distribution of earnings shares within the couple, based on average annual labor market earnings of each partner between ages 55 and 57, for the full Danish population between 1991 and 2013. Panel (d) plots earnings shares for the age-based sample of analysis. The vertical dashed lines mark the interval of couples with very similar earnings shares (between 0.475 and 0.525) who are excluded in the heterogeneity analysis that defines an indicator variable to identify which member of the couple is the primary earner.

Table A.1: Summary Statistics

	Age-Based Design Period (1991–2013)				Reform-Based Design Period (2008–2014)			
	Population		Analysis Sample		Population		Analysis Sample	
	Mean (1)	SD (2)	Mean (3)	SD (4)	Mean (5)	SD (6)	Mean (7)	SD (8)
A: Reference Individuals								
Age	58.45	1.12	58.44	1.12	57.45	2.04	57.47	2.06
Male	0.51	0.50	0.52	0.50	0.50	0.50	0.47	0.50
Dane	0.98	0.15	1.00	0.00	0.97	0.18	1.00	0.00
Copenhagen region	0.26	0.44	0.27	0.44	0.25	0.43	0.22	0.41
Educ. Primary	0.37	0.48	0.29	0.45	0.30	0.46	0.25	0.43
Educ. Secondary	0.41	0.49	0.45	0.50	0.41	0.49	0.45	0.50
Educ. Tertiary	0.03	0.18	0.04	0.19	0.04	0.20	0.04	0.20
Educ. Bachelor	0.14	0.34	0.17	0.37	0.18	0.39	0.20	0.40
Educ. Master	0.05	0.22	0.05	0.22	0.07	0.26	0.05	0.23
Earnings age 55-57	45,268	41,165	60,289	35,186	55,582	41,780	64,156	32,218
Retired by age 57	0.20	0.40	0.09	0.29	0.25	0.43	0.12	0.32
Retired by age 58	0.22	0.41	0.11	0.31	0.26	0.44	0.13	0.34
Retired by age 59	0.24	0.43	0.14	0.35	0.29	0.45	0.16	0.37
Retired by age 60	0.39	0.49	0.34	0.47	0.43	0.49	0.35	0.48
B: Spouses								
Age difference (years)	0.34	5.23	0.25	3.46	0.19	5.26	-0.10	3.50
Age	58.12	5.36	58.19	3.64	57.26	5.62	57.57	4.04
Male	0.49	0.50	0.48	0.50	0.50	0.50	0.53	0.50
Dane	0.99	0.08	1.00	0.06	0.98	0.12	0.99	0.08
Copenhagen region	0.26	0.44	0.27	0.44	0.25	0.43	0.22	0.41
Educ. Primary	0.37	0.48	0.29	0.46	0.28	0.45	0.23	0.42
Educ. Secondary	0.41	0.49	0.44	0.50	0.42	0.49	0.45	0.50
Educ. Tertiary	0.03	0.18	0.04	0.19	0.04	0.21	0.05	0.23
Educ. Bachelor	0.14	0.35	0.17	0.38	0.18	0.39	0.20	0.40
Educ. Master	0.05	0.22	0.05	0.22	0.07	0.26	0.06	0.24
Earnings age 55-57	43,510	40,250	57,770	35,070	52,069	44,725	66,224	34,921
Retired by age 57	0.20	0.40	0.12	0.33	0.26	0.44	0.15	0.36
Retired by age 58	0.21	0.41	0.13	0.34	0.26	0.44	0.14	0.35
Retired by age 59	0.22	0.42	0.15	0.35	0.26	0.44	0.15	0.35
Retired by age 60	0.34	0.48	0.30	0.46	0.35	0.48	0.27	0.44
Number of Observations	4,366,996		2,206,044		166,554		73,395	

Notes: This table reports means and standard deviations of relevant variables for different samples of interest. The first four columns correspond to the age-based period of analysis (1991–2013) where the pension eligibility age remained stable, and it includes individuals of age 57 to 60. The last four columns correspond to the reform-based period of analysis (2008–2014) where the pension eligibility age was increased starting in 2014, and it includes individuals born between July 1, 1953 and June 30, 1954. Columns denoted “Population” correspond to the full population without applying any sample restriction. Columns denoted “Analysis sample” correspond to our baseline samples of analysis, after applying the restrictions described in subsection 3.3.

Table A.2: Heterogeneity in the Effect of Reaching Pension Eligibility Age on Retirement
Alternative to Reweighting: Split by Age Differences and Gender

Reference Individ. Spouse	Young Female Old Male (1)	Young Male Old Female (2)	Old Female Young Male (3)	Old Male Young Female (4)
A. Retirement				
Reference Individ.	0.2801*** (0.002)	0.1765*** (0.004)	0.2257*** (0.004)	0.1409*** (0.002)
Spouse	0.025*** (0.002)	0.030*** (0.003)	0.002 (0.002)	0.005*** (0.001)
Scaled Effect	0.0872*** (0.010)	0.167*** (0.030)	0.00954 (0.016)	0.0359*** (0.015)
B. Claiming				
Reference Individ.	0.4758*** (0.002)	0.2793*** (0.004)	0.3975*** (0.004)	0.2482*** (0.002)
Spouse	0.020*** (0.002)	0.025*** (0.004)	0.000 (0.001)	0.000 (0.000)
Scaled Effect	0.0428*** (0.0055)	0.0878*** (0.017)	0.0041* (0.0023)	0.0033*** (0.0013)
C. Earnings				
Reference Individ.	-9,579*** (93.052)	-9,740*** (248.045)	-7,571*** (166.764)	-8,076*** (114.302)
Spouse	-1,881*** (131.71)	-1,200*** (191.19)	-284 (226.94)	-583*** (75.48)
Scaled Effect	0.197*** (0.023)	0.127*** (0.033)	0.0450 (0.051)	0.0725*** (0.016)
N. of clusters	228,199	69,596	74,390	260,576
Observations	797,667	240,429	256,692	911,256

Notes: This table reports the effect of the reference individuals reaching pension eligibility age on their own retirement and on their spouses' retirement, distinguishing heterogeneous responses by gender and age composition of the couple. Each column contains results for a different subsample. Each panel reports results for a different outcome variable. Within each panel, the first row reports the full exposure effect of pension eligibility on own retirement as estimated in equation (1). The second row reports the full exposure effect on the spouses of their partners being eligible for retirement pension estimated in equation (3). The third row reports the scaled effect resulting from dividing the spouse effect by the own effect. Robust standard errors in parentheses, clustered at the couple level. Bootstrapped standard errors for scaled effects. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table A.3: Descriptive Statistics by Gender and Age Differences

	Female				Male			
	Younger		Older		Younger		Older	
	Mean (1)	SD (2)	Mean (3)	SD (4)	Mean (5)	SD (6)	Mean (7)	SD (8)
Earnings age 55-57	48,213	23,886	50,393	24,445	74,823	45,510	72,220	39,940
College education	0.22	0.42	0.28	0.45	0.25	0.43	0.21	0.41
Retired by age 57	0.11	0.32	0.12	0.32	0.08	0.27	0.07	0.25
Copenhagen region	0.26	0.44	0.30	0.46	0.30	0.46	0.26	0.44
Numer of Observations	213,862		69,661		65,431		240,733	

Notes: This table reports means and standard deviations of relevant variables for all reference individuals in the sample of analysis used for the age-based empirical design. Column (1) corresponds to females who are younger than their partner, whereas column (2) corresponds to females that are older than their partners. Columns (3) and (4) do the same for males. Labor market earnings are computed as the average between ages 55 and 57. Retirement, education, and whether they live in the capital region, are measured at age 57.

Table A.4: Placebo Test with Fake Spouses
for the Effect of Reaching Pension Eligibility Age

	Retirement	Claiming	Earnings
Reference Individual	0.2034*** (0.001)	0.3496*** (0.001)	-8,642*** (69.431)
Spouse	-0.001 (0.001)	-0.002 (0.001)	-30 (78.79)
Scaled Effect	-0.00405 (0.0077)	-0.00481 (0.0035)	0.00349 (0.017)
N. of clusters	367,585	367,585	367,585
Observations	2,206,044	2,206,044	2,206,044

Notes: This table reports the results of replicating the analysis over a placebo sample where the reference individuals are the same as in the main analysis, but they are matched to fake spouses of similar age. The placebo test finds no evidence of joint retirement, as should be expected if the empirical strategy is valid. Fake spouses cannot affect each other's retirement behavior, and the effect coming from the correlation between their ages is controlled for by the empirical design. See the notes of Table 1 for a detailed explanation of the content of the table. Robust standard errors in parentheses, clustered at the couple level. Bootstrapped standard errors for scaled effects.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$


Appendix B Reform-Based Discontinuity Design

Figure B.1: Graphical Depiction of the 2011 Reform



Notes: This figure depicts the 2011 reform that increased retirement ages in 6-month steps contingent on birth date. Cohorts born before January 1, 2014 were unaffected by the reform. Cohorts born between January 1, 2014 and July 1, 2014 experienced an increase of 6 months in their pension eligibility ages. Their early pension eligibility age increased from 60 to $60\frac{1}{2}$, their incentivized early pension eligibility age increased from 62 to $62\frac{1}{2}$ and their full retirement pension increased from 65 to $65\frac{1}{2}$. The red square marks the discontinuity that we exploit in our reform-based research design, where we study the effect of increasing pension eligibility ages. Later cohorts experienced larger increases.

Figure B.2: Birth Date of Spouses by Treatment Group for the Reform Sample



Notes: This graph plots the kernel density function and the probability distribution of the birth date of spouses in the treatment and control groups. Spouses in the treatment group are slightly younger than those in the control group, as a consequence of defining the treatment and control groups based on whether the reference individual was born, respectively, after or before January 1, 1954. Spouses that are born between January 1 and June 30, 1954 (indicated by the solid and dashed vertical lines) are directly impacted by the reform in 2014. We can see from the probability distribution, which is depicted by the dots, that spouses in the treatment group are 1.3 percentage points more likely to be born within those dates than the spouses from the control group (6.5% against 5.2%). Spouses born after June 30, 1954 (dashed vertical line) are impacted by the reform only after 2014. Spouses in the treatment group are 2.2 percentage points more likely to be born after June 30, 1954 (44.3% against 42.1%).

Table B.1: Heterogeneity in the Effect of the Reform Increasing Pension Eligibility Age

Reference Individual	Young	Old	Female	Male
Spouse	Old	Young	Male	Female
	(1)	(2)	(3)	(4)
A. Retirement				
Reference Individual	-0.259*** (0.011)	-0.118*** (0.0098)	-0.258*** (0.011)	-0.118*** (0.0097)
Spouse	-0.0315** (0.013)	-0.00354 (0.0069)	-0.0232** (0.011)	-0.0107 (0.0090)
Scaled Effect	0.122** (0.049)	0.0301 (0.058)	0.0898** (0.044)	0.0907 (0.076)
B. Claiming				
Reference Individual	-0.327*** (0.010)	-0.197*** (0.009)	-0.338*** (0.010)	-0.184*** (0.008)
Spouse	-0.0219* (0.012)	-0.000220 (0.001)	-0.0185* (0.011)	-0.00297 (0.007)
Scaled Effect	0.0669* (0.038)	0.00112 (0.0034)	0.0546* (0.031)	0.0162 (0.036)
C. Earnings				
Reference Individual	10,885*** (667.6)	5,381*** (695.5)	10,678*** (612.7)	5,381*** (743.1)
Spouse	1,366 (905.9)	-29.97 (546.9)	928.1 (872.6)	381 (567.6)
Scaled Effect	0.126 (0.083)	-0.0056 (0.10)	0.0869 (0.081)	0.0707 (0.11)
N. of clusters	5,385	5,161	5,541	5,008
Observations	37,541	35,854	38,542	34,853

Notes: This table reports the effect of the 2011 reform, which increased the pension eligibility age, distinguishing heterogeneous responses by age composition and gender of the couple. Each column contains results for a different subsample. Each panel reports results for a different outcome variable. Within each panel, the first row reports the effect on the individuals affected by the reform and the second row reports the spillover effect on their spouses, which are both estimated in equation (5). The third row reports the scaled effect resulting from the 2SLS model estimated in equation (6). F-tests for the strength of the instruments are all well above 10. Robust standard errors in parentheses, clustered at the couple level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table B.2: Robustness to Alternative Specifications
for the Effect of the Reform Increasing Pension Eligibility Age

	Retirement	Claiming	Earnings
A. Baseline	0.0902** (0.038)	0.0418* (0.024)	0.0847 (0.065)
B. Reweight Spouses Birth	0.0966** (0.040)	0.0324 (0.026)	0.0638 (0.068)
C. Donut Affected Spouses	0.0954** (0.039)	0.0380 (0.025)	0.0895 (0.067)
D. Without Anticipation	0.0858*** (0.032)	0.0432** (0.021)	0.0871 (0.053)
E. Smaller Bandwidth	0.101** (0.042)	0.0494* (0.027)	0.0932 (0.073)
F. Larger Bandwidth	0.0590* (0.035)	0.0320 (0.022)	0.0847 (0.065)
G. Not Balancing	0.0932** (0.039)	0.0462* (0.024)	0.0664 (0.069)
H. Adding Controls	0.0901** (0.038)	0.0415* (0.024)	0.0822 (0.065)
I. No VERP restriction	0.104*** (0.034)	0.0317 (0.021)	0.107** (0.055)

Notes: This table reports the scaled effect estimates (2SLS estimates) from replicating our main analysis using different sample definitions and different specifications of the estimation model (equation 6). Row A reproduces results from our baseline specification, which correspond to those reported in Table 5. Row B reweighs the observations so that the treated and control group have the same distribution of spouses' birth date. Row C excludes spouses born in the first half of 1954. Row D does not estimate the anticipation period separately. Row E reduces the bandwidth by 2 weeks. Row F extends the bandwidth by 2 weeks. Row G does not balance the sample. Row H controls for region and education of reference individuals and their spouses. Row I extends the sample to include individuals who did not contribute to the VERP program between ages 50-59. Robust standard errors in parentheses, clustered at the couple level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table B.3: The Effect of the Reform Increasing Pension Eligibility Age.
Replication Over Sample of Spouses At Least 3 Months Older

	Retirement	Claiming	Earnings
Reference Individual	-0.258*** (0.011)	-0.326*** (0.0099)	10,718*** (680.8)
Spouse	-0.0280** (0.013)	-0.0179 (0.013)	1,480 (938.1)
Scaled Effect	0.109** (0.051)	0.0550 (0.039)	0.138 (0.087)
F-test instr.	523.4	1,078.6	247.8
N of clusters	5,096	5,096	5,096
Observations	35,511	35,511	35,511

Notes: This table replicates the analysis for a subsample where spouses are at least 3 months older than their partners. This ensures that all spouses are born before January 1, 1954, and therefore are totally unaffected by the 2011 reform. This rules out the possibility that the spillover effect to spouses is driven by spouses in the treated and control groups being differentially impacted by the reform. See Table 5 for notes on the construction of this table.