Spousal Spillovers in Retirement: A Structural Assessment

Sigurd M. Galaasen* and Herman Kruse†

July 2021

Abstract

This paper provides a quantitative analysis of the impact of joint retirement on aggregate labor supply elasticities. We extend a standard life-cycle retirement model by incorporating dual-earner households with heterogeneous age gaps and non-separable leisure preferences. To structurally estimate the degree of non-separability we exploit recent quasi-experimental reduced form evidence of spousal retirement spillover effects, from a major Norwegian pension reform. We find evidence of considerable leisure complementaries. Despite large age disparity between spouses, the estimated model importantly generates a high degree of joint retirement, in line with the pattern observed in the data. By solving for the long-run impact of the Norwegian reform we show that leisure complementarity is also quantitatively important. Comparing our estimated model to a model with separability in leisure we find that complementarity accounts for one-third of the reform-induced long-run increase in old-age labor supply.

Keywords: Joint retirement, couples, life-cycle, pension reform, leisure complementarity, spousal spillover

JEL classifications: D15, E2, E6, E65, H55, J08, J2

*Norges Bank. E-mail: sigurd-molster.galaasen@norges-bank.no. This paper should not be reported as representing the views of Norges Bank. The views expressed are those of the authors and do not necessarily reflect those of Norges Bank.

†University of Oslo and Statistics Norway. E-mail: herman.kruse@ssb.no. All the data belongs to Statistics Norway and are used in compliance with the rules given by the Norwegian Data Inspectorate. I am grateful towards the Norwegian Research Council for the financial support (The evaluation of the pension reform, grant number 248868). Thanks are due to NORAM for travel funds. The authors would also like to thank Kjetil Storesletten for invaluable comments and suggestions.

1 Introduction

Across the OECD, generous public pension combined with increasing longevity put pressure on public finances (OECD, 2017, 2019). Social security reforms are thus on the agenda, and a key objective is to stimulate older workers to postpone retirement. Alongside population aging, the composition of households facing retirement decisions has also changed, from traditional "breadwinner" male providers into dual-earner households. Individuals now approaching retirement ages, born in the late 1940s or the early 1950s, are among the first cohorts in which women more commonly have participated in paid labor with full-length careers. As couples, despite large age gaps, tend to retire jointly,[1] this compositional shift may affect how aggregate labor supply responds to pension reforms. A growing body of empirical evidence indeed reveals large spillover effects in retirement decisions among couples; the retirement of one spouse increases the likelihood of the partner also retiring (Kruse (2020); Lalive and Parrotta (2017); Stancanelli and van Soest (2012a,b)). Failure to account for the implied causal link between partners' choices when making policy evaluations can potentially lead to biases in estimated reform impacts (Coile (2003)).

The reduced form approach is, however, unable to separately identify the income and substitution effects underlying the spousal spillover channel. In order to make policy predictions it is crucial to disentangle these two forces. The retirement of one spouse generates two (opposing) effects on the other spouse's labor supply incentives. First, due to the retirement of one spouse, total household income is lower, indirectly reducing the marginal utility of leisure for the working spouse. However, if leisure are complements the working spouse marginal utility of leisure is directly affected, increasing the incentive to retire. In this paper we interpret the reduced form evidence through the lens of a structural model, in order to disentangle the two forces and evaluate the importance of spousal spillover effects for pension reform outcomes.

[1]see (Hurd (1990); Banks et al. (2010) Hospido (2015);)

Our framework is a standard life-cycle retirement model extended by incorporating dual-earner households with heterogeneous age gaps and non-separable leisure preferences. Individuals enter the model as married couples. Each period, households consume, save and make extensive margin labor supply decisions, facing idiosyncratic shocks to productivity, marital status and mortality.[2] At age 62, workers are eligible for old-age pension through an early retirement program. Following Heathcote et al. (2010), we adopt a unitary model of household decision making. In this setup, the household unanimously decides the allocation of each spouse's time, and pool resources into a common budget constraint.

The model is calibrated to the Norwegian economy, using a panel of cohorts (1944–1952) covering the time period 2007–2015. Key preference parameters are estimated using a combination of simulated method of moments and indirect inference. Importantly, the time period encompasses the implementation of the 2011 Norwegian pension reform which was designed to increase work incentives. In our calibration, we exploit this reform to quantify the magnitude and direction of non-separability in leisure. In particular, we target the reduced form point estimate in Kruse (2020), who finds that women (men) aged 63–66 are approximately 17 (6) percentage points more likely to work if their partner also works. Key to the identification is the quasi-experimental setting in which the reform affected only the private sector early-retirement program. We pin down non-separability in leisure by replicating the empirical design on a similar simulated panel, and find that a large degree of complementarity is needed for the model to be consistent with the reduced form evidence. To illustrate the identification of complementarity, we also consider a calibration with separable leisure in which we drop the reduced form target. The implied spillover effect is then highly negative, indicating that the reduced form moment is particularly informative for the complementarity parameter.

Importantly, the estimated model is able to replicate the frequency of joint

[2] Upon separation, either through divorce or death, individuals remain single for the remainder of their lives.

retirement seen in the data. A key feature in the data is that, despite large age gaps, couples tend to retire at the same time (within the same year). Our estimated model, featuring the empirical age gap distribution,[3] generates similar patterns. Consequently, the degree of complementarity consistent with the reduced form spillover effects of the Norwegian pension reform, is simultaneously able to account for the observed patterns of joint retirement. By contrast, in the model with separable leisure, couples tend to not retire at the same time. The fact that leisure complementarity is needed to explain the incidence of joint retirement is consistent with other structural models of couples labor supply decisions (e.g Gustman and Steinmeier (2004); Casanova (2010); Michaud et al. (2018)). We contribute to this literature by showing that leisure complementary can jointly explain the quantitative magnitude of spillover effects and the occurrence of couples retiring together.

Finally, we quantify the aggregate importance of leisure complementarity for old-age labor supply elasticities. To do so, we simulate the long-run impact of the 2011 Norwegian pension reform. The importance of complementarity is derived by comparing the long-run employment response in the model with and without leisure complementarity. Our quantitative model exercise shows that 28 percent of the increase in average retirement age among couples, in the Norwegian reform context, is due to leisure complementarity. To derive these results we simulate an economy in which the share of treated couples are taken from the data.

However, the importance of leisure complementarity for the aggregate employment is dependent on whether either or both of the spouses are directly targeted by the policy change. If we consider only couples in which both spouses are targeted, meaning that both spouses are subject to some policy change that affects the individual incentives to work, leisure complementarity accounts for 18 percent of the aggregate employment response. On the other hand, in couples

[3] It is crucial to incorporate age gaps in the model in order to compare to the data along this dimension. In models with only same-age couples it is difficult to separate between preferences and age-effects as explanations behind synchronized decisions.

where only one spouse is targeted directly, leisure complementarity accounts for 34 percent of the aggregate employment response. Consequently, our results indicate that spillover effects are nearly twice as important in targeted versus universal reforms.[4] Many labor market reforms are effectively targeted as opposed to universal (e.g. policies aimed at reducing early retirement through welfare benefits such as disability pension). Our findings indicate that such reforms may have substantial positive indirect effects, through the spousal spillover channel.

1.1 Related literature

Our paper relates to the literature on quantitative life-cycle models, and in particular to the increasing literature on dual-earner households and social security reforms.

A growing body of recent contributions are emphasizing the importance of including women, and the interlinking of spouses, to life-cycle models. Fernández and Wong (2011) study female labor force participation in the US over time, and points to the increase from a 40% participation rate among women born in 1935 to a 70% participation rate among women born in 1955. She studies this in a life-cycle model and finds that increased divorce rates and changes to wage structures can account for 60% of this observed change in participation rates. Eckstein and Lifshitz (2011) also study female labor force participation, and compares married women to unmarried women. They show that the rise in education levels can account for 33% of the increase in female employment, and another 20% by increased wage levels and reduced gender-gap. However, they cannot account for the remaining 40%, but empirically attributes it to cohort-specific changes in preferences, child-care cost and household-maintenance.

Guner et al. (2012) study gender-based taxes, and find that lower and proportional income taxes on women will improve output, labor participation and welfare among women. Fernández and Wong (2014) look at how assortative mar-

[4] The intuition is that without leisure complementarity, the non-targeted spouse reduces labor supply due to a negative income effect caused by the targeted spouse increasing labor supply.

riage, fertility, divorce and remarriage probabilities and women's return to work affect women's labor force participation. Using a dynamic life-cycle model with certain similarities to our framework, they are able to account for about 85 percent of the change in married women's LFP from the 1935 cohort to the 1955 cohort. Blundell et al. (2016) use a dynamic life-cycle model for women in the UK to study employment and human capital accumulation. They exploit a tax and benefit reform to calibrate their parameters, and find substantial labor elasticities (overall Frisch elasticity of 0.64 on the extensive margin, and 0.24 on the intensive margin), in particular for lone mothers. They also find large and significant returns to experience for full-time work, in particular for high educated women. Eckstein et al. (2016) study the life-cycle decisions of five cohorts of American men and women born between 1930 and 1970. They focus on explaining differences in education, work, divorce and fertility between these cohorts, and find that married women of the more recent cohorts have significantly higher observed and unobserved skills compared to earlier cohorts.

A recent contribution is that of Borella et al. (2018). Their paper investigates the entire life-cycle of couples in an attempt to emphasize that models disregarding gender and marriage will not match key economic aggregates particularly well. They show in particular that a model of an economy consisting of "no marriage, only men" will overestimate aggregate outcomes of labor participation, earnings, and hours worked, while it will underestimate retirement saving.

Our paper relates closely to the literature on joint retirement in life-cycle models. van der Klaauw and Wolpin (2008) focus on poor households and their retirement and savings behavior. They use the first three waves of the Health and Retirement Survey (HRS) to calibrate a model incorporating various labor market moments to simulate responses to reforms of social security systems, for instance removal of earnings-tested pension benefits and elimination of early retirement. Their model predicts sharp increases in annual hours of work and in full-time employment among elderly workers (62–69 years), and smaller responses among younger workers. Interestingly, they find that singles respond more than married couples, and men respond more than women. Casanova (2010) estimates joint retirement in the US using the HRS, and find that leisure complementarity accounts for 8 percent of observed joint retirements, while another 13 percent is accounted for by social security spousal benefits. Martín and Marcos (2010) exploit a Spanish pension reform to analyze saving and labor market decisions at the household level in an OLG framework, and show inter alia that traditional single-earner households are inadequate and may result in underestimates of the future financial burden of pension systems. Later, Michaud et al. (2018) used a model with intra-household bargaining to show strong positive correlations between partners' preferences for joint leisure, and showed that leisure complementarity accounts for a substantial part of joint retirement.

Our paper adds to this set of papers on joint retirement by providing new evidence on leisure complementarity in structural models of dual earner households. We use a structural model consistent with reduced form evidence Kruse (2020), and show that preferences towards shared leisure in the household utility function can perfectly match the estimated a reduced form spillover coefficient. Moreover, our model explores the effects of age gaps in marriages: even though age gaps in marriages may be rather disperse, couples tend to synchronize their retirement timing, and empirical evidence shows that a significant fraction of couples retire within the same calendar year (Kruse, 2020; Hospido, 2015). We pin our age gap distribution to Norwegian register data, and show in an out-of-sample test that our model replicates this characteristic "spike" in synchronized retirement. We also evaluate an extensive reform of the Norwegian pension system from 2011 in light of leisure complementarity in the long-run, and show that leisure complementarity accounts for a substantial part of aggregate labor market responses.

The remainder of this paper is organized as follows. First, in Section 2, we present background information on a recent major reform of the Norwegian pension system, as well as introduce the reduced form evidence on which we base the calibration of our model. Then, in Section 3, we present our life-cycle model. In Section 4, we present the calibration of our model, while in Section 5 we present the results. Finally, we conclude in Section 6.

2 Background

In this section, we describe institutional details and empirical background that will be key to calibrate and evaluate our structural model. Norway reformed its pension system in 2011, where a key target was to motivate elderly workers to extend their working lives. We describe the main features of the Norwegian pension system before and after the reform in Section 2.1. The main reason for why we introduce the Norwegian pension reform, is to exploit exogenous variation in work incentives to pin down leisure complementarity, as the reform changed the financial returns to work for certain worker groups, while others remained largely unaffected. Therefore, in Section 2.2.1 we present reduced form evidence of spousal retirement spillover effects from a study by Kruse (2020), as well as stylized facts regarding the extent of joint retirement. The reduced form evidence will be employed as a calibration target to structurally pin down leisure complementarity, while the stylized facts will be used as an out-of-sample test of the model.

2.1 The 2011 Norwegian pension reform

The Norwegian reform targeted work incentives among the elderly workers by the removal of a specific confiscatory earnings test. We divide this section into two parts. First, we present the *public old-age pension system*, which is common to all workers. This is the basic foundation of the pension system, for which all workers eventually become eligible. In terms of work incentives, the system was largely unchanged, but it saw some changes that nevertheless will be documented. Second, we present the *early retirement scheme*, known by it's acronym AFP.[5] This system is available to the entire public sector, and about half of the private sector workers

[5] From Norwegian "Avtalefestet pensjon".

(working in firms that have opted in to the scheme). It is in this system, and in particular *only for the private sector workers*, that the major changes to work incentives occurred in the 2011 reform.

2.1.1 Public old-age pension

The Norwegian public old-age pension is a "pay-as-you-go" system with three main pillars: (*i*) the old-age guarantee pension granted after forty years of residence and "pension-awarding income"[6], (*ii*) an income-related pension based on a mapping of the 20 best years of income, and (*iii*) a mandatory defined-contribution system for employers. The sum of the three pillars is the *old-age pension* with a universal access age at 67.

This system was largely unchanged in terms of work incentives in the 2011 reform. The old-age pension access age was reduced to 62 years, but is now subject to actuarial fair adjustment and life-expectancy adjustment. This means that claiming benefits earlier leads to a cut in annual benefits, so that the access age reduction is *incentive neutral, but provides flexibility.* However, this change was universal across private and public sector, and has been shown by Hernæs et al. (2016) and Hernæs et al. (2019) to have minor implications for work incentives, at least on the extensive margin.[7]

2.1.2 Early retirement scheme

Since it's launch January 1, 1989, about two-thirds of all workers (the entire public sector and about half of the private sector) have had access to an *early retirement option*, known by its acronym AFP. Prior to the 2011 reform, the AFP was identical in the public and private sector, but the reform only covered the private sector workers, meaning that the systems are currently different. In the pre-reform system (which is still at force in the public sector), the AFP offered the equivalent of a *full old-age pension claim* already at age 62 for workers who had sufficient accrual.

[6]Defined as annual income above a baseline income level, set by the government and indexed annually. In 2019, the baseline annual income level was 98,866 NOK or roughly 11,500 USD. In the following, we will denote this unit G.

[7]In particular Hernæs et al. (2019) show that there is a small negative effect of the eligibility reduction on the intensive margin, mostly among high-income earners, but they find no effect on the extensive margin.

The benefits could be claimed between ages 62–66, after which the pensioner would be transferred to the old-age pension system.

Pre-2011, and still in the public sector, workers who claimed AFP benefits faced an *earnings test* on continued work. The earnings test was a *pro rata* benefit cut if the individual decided to combine claiming benefits with working.[8] In the reformed system, applying to private sector workers born after 1949 the benefits were recalculated into *lifelong annuities* and the *earnings test was removed entirely*.[9] This meant that the retirement *timing* and the *benefit claiming* was decoupled, and private sector workers with AFP are now free to combine early retirement claiming with continued work. In expectation, the lifelong annuities have the same net present value as the old AFP claims, and are paid as a *top-up* of the old-age pension.

2.2 Model targets for leisure complementarity

Empirical evidence using Norwegian register data shows that, even though age gaps in couples are rather disperse, we observe that there is a bunching around same-year exits from the labor market within couples (Figure 1).[10] This labor market exit synchronization could in principle be explained by at least two different mechanisms: correlated labor market shocks or synchronized decisions arising from leisure complementarities. We will, in Section 3, build a model incorporating the latter mechanism. Importantly, we will present empirical evidence on leisure complementarity in Section 2.2.1 which is independent of the former mechanism. This is the target evidence we use to pin down the degree of leisure complementarity. Moreover, the striking "spike" in Figure 1a, which indicates synchronized

[8]This implied a one-to-one cut in benefits for every dollar earned.

[9]There was transitional rules applying to the 1948 cohort, who were given the possibility to be transferred into the new system if they postponed claiming until after January 1, 2011. In addition, the cohorts born in 1947 and earlier were given the same opportunity, but with only a partial transfer to the new system based on the number of years of claiming remaining in the old system. For instance, the 1947 cohort was 64 years old in 2011, and therefore was given 3/5 (60 percent) of the new system, while the 1946 cohort was then given 2/5 (40 percent) of the new system.

[10]This phenomenon is not unique to Norway. For instance Hospido (2015) showed that the same pattern can be found using the Survey of Health, Ageing, and Retirement in Europe, and Banks et al. (2010) showed similar patterns using both UK and US data.

labor market exit, is replicated in the model as an out-of-sample test.[11]

Figure 1
Joint exit and age gap distribution in couples. In panel (a), left tail are couples where the wife retires last. In panel (b), left tail are couples where the wife is older.

(a) Joint exit **(b)** Age gap distribution

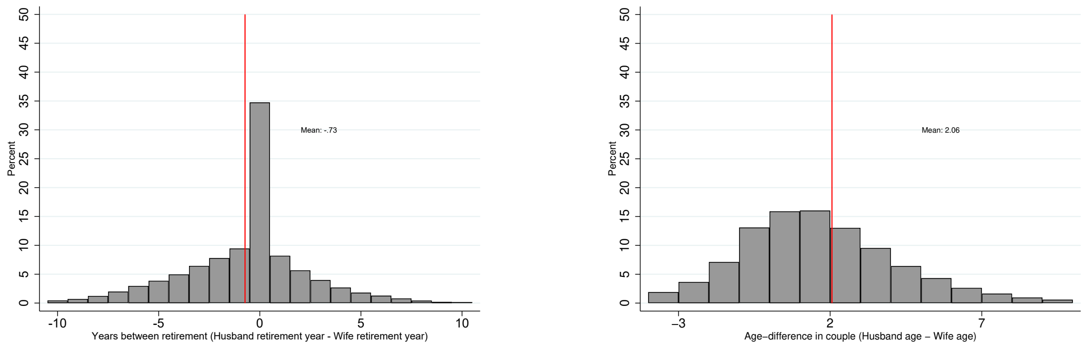

Notes: Each bar in panel (a) represents the fraction of couples in which both spouses retire within a specific time gap. Panel (b) displays the distribution of age gaps among couples. The red lines show the mean of each distribution. Data source: Norwegian register data on all married couples where at least one partner is born in 1944–1952. Both partners must work at age 60.

2.2.1 Reduced form evidence

To pin down complementarity in leisure we target evidence from a recent study by Kruse (2020) exploiting the 2011 pension reform to investigate spousal spillovers in retirement. The author exploits the fact that the reform affected private sector AFP workers differently from public sector AFP workers. The heterogeneous treatment implies that the author can compare the labor supply responses to the reform, for two different types of individuals who themselves are unaffected by the reform (i.e. who worked in the public sector at the time of the reform). The first type has a spouse who is affected by the reform (the spouse works in the private sector with AFP) while the other type has a spouse who is unaffected (the spouse works in the public sector). The *focal partner* the author estimates the spillover effect onto (individual j in the following empirical setup), is always a public sector worker and therefore not directly affected by the changes to the AFP. The spouse (individual i in the empirical setup), can either be a private sector worker with access to AFP (the treatment group) or a public sector worker (the control group).

[11]In Appendix A, we show that the latter mechanism, correlated earnings shocks, is essentially unable to generate synchronized retirement even at a very high degree of correlation.

The empirical setup is then a two-stage least squares model:

$$E_{i,t} = \alpha_0 + \alpha_1 D_i^{private} + \alpha_2(D_i^{private} \times D_i^{post}) + \alpha_3 X_{ij} + \lambda_t + \varepsilon_{i,t} \quad (1)$$

$$E_{j,t} = \phi_0 + \phi_1 D_i^{private} + \phi_2 E_{i,t} + \phi_3 X_{ij} + \lambda_t + \varepsilon_{j,t} \quad (2)$$

where equation 1 is the first-stage and equation 2 is the second-stage. Here λ_t is time fixed effects, $D_i^{private}$ is the sector affiliation of spouse i, and X_{ij} is time-invariant control variables for the couple. Essentially, $D_i^{private} \times D_i^{post}$ is an instrumental variable for the sector affiliation of the i-spouse and the spouse being born in the treated cohorts, i.e. born after 1949. A focal partner (j) married to a spouse with these characteristics, should in presence of leisure complementarity experience an increased employment probability if the i-spouse is also employed. This is governed by the coefficient ϕ_2. Kruse (2020) shows that women aged 63–66 in this setup has a ϕ_2 of 0.17, which can be interpreted as a 17 percentage points increase in the probability of remaining employed if the husband remains employed. For men, the effect is imprecisely estimated, but with a much lower point estimate of 0.06. In our calibration, we target this empirical finding to pin down leisure complementarity. Importantly, we use our model to generate a similar panel data set, and use the exact same empirical setup.

3 Model

We consider a standard consumption-saving life cycle model, extended with multi-person households and endogenous retirement. There are J overlapping generations, and each model period corresponds to a calendar year (age). Individuals enter the model at $j_g = 22$, born with gender g (either male m or female f), and married. Over the life-cycle, marriages may dissolve either due to divorce or mortality. Households derive utility from consumption and leisure, and make decisions on consumption, saving, and labor market participation. Financial markets are incomplete and agents face three sources of idiosyncratic risk: labor earnings, marital status (divorce), and mortality. The household can save in a risk-free asset subject to a borrowing constraint.

In addition to heterogeneity generated by the realization of idiosyncratic risk over the life cycle, individuals are ex-ante heterogeneous along two dimension: (*i*) pension system eligibility and (*ii*) spousal age gaps. First, upon reaching age 62, agents are eligible for old-age pension, either from the public sector or the private sector early retirement (AFP) system. We define an individual as being ex-ante born as either a private or public sector worker. For the purpose of the model, the distinction between public and private sector is to distinguish the individuals that were affected by the reform presented in Section 2.1. Since the evolution of labor earnings in the data differs substantially between public and private sector worker types, we estimate a flexible earnings process which may differ across types.

Second, couples are ex-ante heterogeneous in the within-couple age disparity. In the data, most individuals live in marriages in which one spouse is older than the other. In the presence of leisure complementarity, such age gaps may have important implications for individual retirement patterns.[12]

We now turn to the household decision problems. First we lay out the couple's environment before we describe how the environment changes when living as a single individual household, in our model interpreted as a divorcé(e) or a widow(er).

3.1 Couples

We follow Heathcote et al. (2010) and adopt a unitary model of multi-person households. When explicitly modeling households as dual-earners, one must consider how the household coordinate their decisions. In this framework, the household unanimously decides the allocation of each individual's time, and they pool their income into a common budget constraint. One interpretation of this approach is that both consumption and leisure are public goods within the household, or

[12] Age gaps means that the older spouse will enter the model before the younger spouse. When solving the model, we make the (innocuous) assumption that the older spouse knows that marriage will be formed once the younger spouse enters the model at age 22.

that couples are perfectly altruistic towards each other. In contrast to households in Heathcote et al. (2010), our multi-person household also cares about future non-shared states. To handle this we assume a Pareto problem with fixed Pareto weights, as in Hong and Ríos-Rull (2012).[13]

In the following sections we will lay out the the households choices and preferences, exogenous shock processes, and budget constraints. To ease the notation, we suppress the subscript denoting which sector (type) the individual belongs to (private or public). Since each household member can can be one of two types (private/public), a multi-person household can be one of four types. All parameters are common to all individual types, except for the pension scheme and labor earnings process, which are type specific.

3.1.1 Choices and preferences

Couples consist of a male (m) and a female (f) spouse, of ages (j_m, j_f). They have preferences over joint consumption and leisure. Each period the couple makes joint decisions on how much the household should consume c and save in the safe asset a, and whether to work or not $(l_f, l_m) \in \{0, 1\}$. The objective is to maximize expected present value of utility, taking into account the probability of separation (due to divorce or death).

We adopt an iso-elastic instantaneous joint-utility function, separable in consumption and household total leisure, but non-separable in spousal leisure. For a couple of age $j = (j_m, j_f)$ joint utility is given by:

$$u_j(c, l_f, l_m) = \frac{\left(\frac{c}{\zeta}\right)^{1-\gamma}}{1-\gamma} + u_j^l(l_f, l_m) \tag{3}$$

where the last term captures the utility of *effective* leisure, given by:

$$u_j^l(l_f, l_m) = \frac{\kappa}{1-\phi} \left[\left(\eta (1 - \theta l_f - \chi_{j_f})^{\rho} + (1-\eta)(1 - \theta l_m - \chi_{j_m})^{\rho} \right)^{\frac{1}{\rho}} \right]^{1-\phi} \tag{4}$$

[13] An alternative approach is to assume that each household member maximize its own utility, with household allocations determined through bargaining. In this paper, we choose our approach due to its tractability and computational simplicity.

The parameter γ is the inverse of the intertemporal elasticity of substitution (IES) of consumption, and ϕ the inverse IES of leisure, while ζ is the equivalence scale of shared consumption determining the returns to scale in joint consumption. The parameter κ measures the relative weight on leisure in utility, η the relative weight on female leisure, ρ the substitution elasticity between the male and female leisure, and θ time spent working full time.

We assume that effective leisure is decreasing in age, captured by χ:

$$\chi_{j_g} = \begin{cases} 0 & \text{if} \quad j_g \leq j^\chi \\ 1 - \exp(-\chi_0(j_g - j^\chi)) & \text{if} \quad j_g > j^\chi \end{cases} \tag{5}$$

where j^χ is the age at which the decline starts. This effectively makes working more costly as the agent approaches older ages, as less effective time is available. This is consistent with life-cycle models in French (2005), French and Jones (2011) and Capatina (2015), in which individual time endowment is on average convexly decreasing in age, due to deteriorating health. The latter paper also shows that an individual's expected time endowment decreases non-linearly in age. We do not model health explicitly, and instead adopt a reduced form approach.[14]

In the parameterization of the model, we will take values for (γ, ϕ, ζ) from standard estimates in the literature, and calibrate $(\rho, \eta, \kappa, \chi_0)$ internally by matching retirement patterns in the data. Specifically, κ determines the average retirement age in the model, while η targets average level difference in retirement ages between males and females, and χ_0 the slope of the aggregate age-participation rate among married couples. Finally, a key feature of this specification of the instantaneous utility function is the tractable interpretation of leisure complementarity.

[14]More generally, this modeling assumption is broadly in line with several studies that adopt a convex (in age) time cost of working, to address the sharp rise in non-participation among old-age workers, see e.g. Kitao (2014), Cooley and Henriksen (2018) and Cooley et al. (2019).

Specifically, if:

$$1-\phi-\rho\begin{cases} >0 & \text{then} \quad \text{leisure is complementary} \\ \to 0 & \text{then} \quad \text{separable utility of leisure} \\ <0 & \text{then} \quad \text{leisure is a substitute} \end{cases} \tag{6}$$

This simple relationship comes from the fact that the sign of the cross-partial derivative of utility with respect to leisure depends only on the sign of $1-\phi-\rho$. The interpretation is that the leisure of spouses are complements when the marginal utility of leisure is increasing in the spouse's leisure.

3.1.2 Mortality and divorce

Since our focus is mainly on the later stages of the life-cycle, we adopt a simple process for the family structure. Individuals are born into a couple, and remain married until divorce or widowhood. We abstract from remarriage, assuming that once the relationship ends, individuals remain single for the remainder of their lives.

For a couple, there are five possible next-period demographic outcomes: they survive and remain married, they survive and divorce, the female spouse dies, the male spouse dies, or both die. We denote the probability of transitioning from currently married at age composition (j_m, j_f) to these five next-period states as $\xi^z_{j_m,j_f}$, where $z = \{M, d, w_m w_f\}$ denotes "remain married", "divorce", "male widower" and "female widow", respectively. These probabilities are derived from the survival probabilities s_{g,j_g} and the divorce probabilities $d(j_m, j_w)$ which are exogenously given.[15]

Divorce is modeled as an exogenous shock. One could argue that divorce in reality is an endogenous decision, often related to job loss, bad income shocks, health or other factors such as child-care or loss of affection. However, apart from income shocks, these demographic and economic factors are not in the model, or are unobservable. Hence, from a model perspective, they should be treated

[15] For example, the probability of becoming divorced next period is $\xi^d_{j_m,j_f} = s_{m,j_m+1} s_{f,j_f+1} d(j_m, j_f)$.

as random. We therefore follow Cubeddu and Ríos-Rull (2003) and Hong and Ríos-Rull (2012) and model divorce as a stochastic shock.

3.1.3 Labor earnings process

At the beginning of a period, each member $g \in \{m, f\}$ receives an endowment of ω_{g,j_g} log units of labor efficiency. The individual labor income is then given by:

$$y_{g,j_g} = \exp(\omega_{g,j_g}) l_g \tag{7}$$

The endowment $\omega_{g_j g} = q_{g,j_g} + e_{g,j_g}$ consists of two parts. The first term, q, captures the deterministic age-component of the labor earnings profile, common to all agents. The latter term, e, captures the idiosyncratic shock, which follows an AR(1) process

$$\begin{aligned} e_{g,j_g} &= \nu_g e_{g,j_g-1} + \varepsilon_{g,j_g}, \\ \varepsilon_{g,j_g} &\sim N(0, \sigma^2_{\varepsilon_g}) \end{aligned} \tag{8}$$

The earnings shock process is assumed to be independent across spouses.[16] When taking the labor earnings relationship to the data, we estimate it separately for the different worker types (public/private sector)

3.1.4 Pension system

To keep the model tractable, we adopt a simplified pension system in the model. Pension benefits are pre-determined and assumed to be a constant fraction of the deterministic part of labor earnings at age 64:

$$\hat{b}_g = 0.6 \exp(q_{g64})\theta \tag{9}$$

[16]Note that correlated earnings processes could generate correlated retirement behavior across spouses, in much the same way as leisure complementarity. However, since we pin down the complementarity parameter ρ by relying on exogenous variation in retirement incentives embedded in the Norwegian Pension Reform, our estimate of ρ is robust to potential correlated income processes in the data.

Although highly stylized, the pension system approximates the strong progressivity of Norwegian pension system. The fraction 0.6 is roughly in line with the average gross replacement rate in the Norwegian system. Note that despite the fact that the pension benefit level is homogeneous (within worker type/gender), heterogeneous labor earnings generate progressive replacement rates.

Workers are eligible for pensions when they reach age 62. However, the pension benefit is tested against labor earnings. Importantly, consistent with the Norwegian pension system prior to 2011, the lost pension benefit does not lead to an upwards re-adjustment of future pensions, causing a substantial implicit tax on labor. Consequently, the annual net pension benefit individuals receive is:

$$b_{g,j_g} = f_{j_g}(\hat{b}_g, y_{g,j_g}), \tag{10}$$

where the function f is zero until age 62, and decreasing in y.

3.1.5 Constraints

Non-participation (i.e $l_g = 0$) is assumed to be an absorbing state, which means that once retired, the worker cannot return to the labor market.[17] Formally, let $\mathbb{I}_{g,j_g}$ denote a 0-1 indicator function taking the value 1 if the individual did not work in the previous period. The current period labor supply choice set is then defined as:

$$l_{g,j_g} \in \begin{matrix} \{0\} & \text{if} \;\; \mathbb{I}_{g,j_g} = 1 \\ \{0,1\} & \text{if} \;\; \mathbb{I}_{g,j_g} = 0 \end{matrix} . \tag{11}$$

The household evolution of joint assets follows from the sequence of resource

[17] Our primary focus is on the retirement pattern of older workers and our data sample consists of workers between the ages 60-69, who have not yet retired at age 60. The assumption of absorbing non-participation is consistent with the very low transition rate back to work among these individuals.

constraints:

$$a_{j+1} = (1 + r(1 - \tau_a))a_j + T(y_{f,j_f}, b_{f,j_f}) + T(y_{m,j_m}, b_{m,j_m}) - c_j \tag{12}$$
$$a_{j+1} \geq 0,$$

where the last constraint rules out borrowing. The asset yields an exogenous safe return r taxed at rate τ_a and the function T maps pre-tax labor and pension income into after tax income.

3.2 Divorced or widowed individuals

Upon separation an individual becomes a single-person household, which is an absorbing state. The only thing that changes relative to the couples' environment is the utility function, which for a single household is:

$$u^S_{j_g}(c_g, l_g) = \frac{c_g^{1-\gamma}}{1-\gamma} + \frac{\kappa^S}{1-\phi}(1 - \theta l_g - \chi^S_{j_g})^{1-\phi} \tag{13}$$
$$\chi^S_{j_g} = \begin{cases} 0 & \text{if} \quad j_g \leq j^\chi \\ 1 - \exp(-\chi^S_0(j_g - j^\chi)) & \text{if} \quad j_g > j^\chi \end{cases}$$

where the superscript S denotes single household. We allow the weight on leisure κ^S and the rate of decline in time endowment χ^S_0 to potentially differ to that of a couple household. When calibrating the model, these parameters will target the retirement patterns of separated individuals.

The evolution of assets is given by:

$$a_{g,j_g+1} = (1 + r(1 - \tau_a))a_{g,j_g} + T(y_{g,j_g}, b_{g,j_g}) - c_{g,j_g} \tag{14}$$
$$a_{g,j_g+1} \geq 0,$$

The separated agent faces the same individual mortality and earnings process, and pension system as when in a couple. Hence, labor earnings evolves according to (7)-(8) and pension benefit is determined by (9)-(10). Labor supply is still constrained by absorbing retirement following (11).

3.3 Recursive formulation

Let the vector of state variables for a couple be:

$$\tilde{\Omega} = \{a, e_m, e_f, j_m, j_f, \mathbb{I}_m, \mathbb{I}_f\} \tag{15}$$

where m denotes the male spouse and f denotes the female spouse, and a is the shared assets, e is the current income shock, j is the current age. The corresponding state vector for a single person is:

$$\Omega = \{a_g, e_g, j_g, \mathbb{I}_g\} \tag{16}$$

In recursive form, the couple's problem is given by:

$$\begin{aligned} &V(\tilde{\Omega}) = \max_{a', c, l_f, l_m} \{u(c, l_f, l_m) + \\ &\beta \left[\xi^M \mathbb{E}V(\tilde{\Omega}'|\tilde{\Omega}) + \xi^d((1-\lambda)\mathbb{E}V^S(\Omega'_f|\tilde{\Omega}) + \lambda \mathbb{E}V^S(\Omega'_m|\tilde{\Omega})) + \right. \\ &\left. \xi^{w_f}(1-\lambda)\mathbb{E}V^S(\Omega'_f|\tilde{\Omega}) + \xi^{w_m}\lambda \mathbb{E}V^S(\Omega'_m|\tilde{\Omega})\right]\Big\} \end{aligned} \tag{17}$$

subject to (7)-(12), where β is the discount factor. The first term in the square brackets represents the continuation value if remaining married, which happens with probability ξ^M. The next four continuation values represents states that are not shared. The first two corresponds to divorce, occurring with probability ξ^d, and the final two corresponds to the death of either the female or male spouse. In the event of both spouses dying, the continuation value is zero, implying no bequest motive. The parameter λ reflects the weight the couple puts on the male continuation value in future states that are not shared. We set this parameter to $\lambda = 0.5$, implying equal weight on future non-shared states, consistent with the altruistic interpretation of the instantaneous utility function. The expectation operator $\mathbb{E}$ is over next period earnings shock realization, conditional on current period state.

In the event of death, the entire joint asset a' goes to the widow(er). In the event of divorce, household assets are divided among the spouses and a share π is allocated to the female, implying $a'_f = \pi a'$ and $a'_m = (1-\pi)a'$

For a single individual, the corresponding problem is:

$$V_g^S(\Omega_g) = \max_{a'_g, c_g, l_g} \left\{ u^S(c_g, l_g) + \beta s_{g,j+1} \mathbb{E}\left[V_g^S(\Omega'_g | \Omega_g) \right] \right\} \tag{18}$$

subject to (7)–(11) and (14).

For the older spouse in a couple, we also need to solve the decision problem prior to the younger spouse entering the model. We do so by assuming that the older spouse adopts the single household decision problem, with perfect foresight about the transition into a couple when the younger spouse enters the model at age 22.

3.4 Solution method

The model is solved by backwards recursion. We use straightforward discretization of the state and control space. The auto-regressive part of the earnings process is approximated by a two-state first order Markov process, following Tauchen and Hussey (1991). We consider a partial equilibrium framework, meaning that interest, wages, and tax rates are constant and exogenous.

4 Calibration

We estimate the model using a combination of external and internal calibrations. In the external calibration presented in Section 4.1, we find the mortality, divorce and labor earnings processes using Norwegian register data. The tax and pension system rules are set to mimic key features of the Norwegian system, both before and after the 2011 pension reform. Finally, a set of preference parameters are taken from the literature. Table 1 provides an overview of the external parameters, values, and data sources. In the internal calibration, presented in Section 4.2, we pin down parameters by matching simulated moments to corresponding data moments. In particular, the remaining parameters $\{\rho, \eta, \beta, \kappa, \kappa^S, \chi_0, \chi_0^S\}$ targets the retirement patterns of men and women, wealth to income ratios, as well as spousal spillover effects in retirement, cf. Section 2.2.1. Table 3 provides an overview of the parameters and targets.

Table 1
External parameters of the model.

Parameter	**Description**	**Value/source**
Demographics		
s_{g,j_g}	Gender specific conditional survival probabilities	*
$d(j_m, j_w)$	Divorce probabilities	*
J	Maximum age	102 (81 model periods)
R	Forced latest retirement age	70 (48 model periods)
j^χ	Age when time endowment starts to decline	61 (40 model periods)
$\max(j_m - j_f)$ and $\min(j_m - j_f)$	Age gap extrema	$\in \{-3, 7\}$
$j_m - j_f$	Age gap distribution	*
Preferences		
ϕ	Intertemporal elasticity of substitution, total leisure	3, Heathcote et al. (2010)
γ	Intertemporal elasticity of substitution, consumption	1.5, Low and Pistaferri (2015); Low et al. (2018)
ζ	Equivalence scale couple utility	1.3, Hong and Ríos-Rull (2012)
π	Sharing rule, joint assets	0.5, Hong and Ríos-Rull (2012)
λ	Weight on male continuation value	0.5
θ	Time spent working (full-time)	0.4 (40 percent of time endowment)
Labor productivity process		
ν_m	Wage persistency, men, private sector	0.916, *
ν_f	Wage persistency, women, private sector	0.852, *
$\sigma_{\varepsilon_m^2}$	Wage innovation, men, private sector	0.187, *
$\sigma_{\varepsilon_f^2}$	Wage innovation, women, private sector	0.174, *
ν_m	Wage persistency, men, public sector	0.819, *
ν_f	Wage persistency, women, public sector	0.794, *
$\sigma_{\varepsilon_m^2}$	Wage innovation, men, public sector	0.131, *
$\sigma_{\varepsilon_f^2}$	Wage innovation, women, public sector	0.106, *
q_{gj_g}	Gender specific age-earnings profiles	Table 2, *
Pension system		
f_{j_g}	Earnings test on pension benefits	Norwegian social security rules
Government		
r	Pre-tax interest rate	0.04
τ_a	Tax on interest income	0.28, *

Notes: *Source: Norwegian register data. Survival probabilities, divorce probabilities and the labor productivity processes are life-cycle moments which must be computed for all ages. They are therefore calculated on the entire Norwegian population for the time period 1993–2015, except the labor productivity process which only includes monthly data in the period 1997–2014. The age gap distribution is calculated based on the 1944–1952 cohorts and are similar to the estimation sample used in Kruse (2020). Norwegian social security rules regarding earnings test on pension benefits are largely invariant separately before and after the 2011 reform (but changes over the reform margin), at least for the time period for which we estimate (until 2015).

4.1 External calibration

4.1.1 Demographics

Mortality and divorce The mortality and divorce data is taken from a longitudinal panel of all Norwegian individuals from 1993–2015. The only initial restriction we make when computing the demographic statistics is that the individuals must be above age 22 and be Norwegian residents, which means that we exclude emigrants (but include immigrants).

We compute a gender-specific mortality table for each gender at each age, giving us an age-gender-specific risk of death. We set the maximum age to $J = 102$ (which corresponds to 81 model periods). The mortality risk is simply found by computing the ratio of individuals who die at age $j + 1$ to those who are alive at age j. This means that our estimate of mortality risk includes a (weighted[18]) average of mortality at age j over the years 1993–2015.

Divorce probabilities are estimated for every age from 22–70. Similarly to the mortality risk, the divorce risk is the (weighted) average of married individuals who divorce at age $j + 1$ to the number of individuals married at age j.

Age gap and sector distributions The distributions are computed on the sample of 1944–1952 cohorts of married individuals who are working at age 60.[19] The age gap distribution is computed directly from the data sample. On average, men are about 2 years older than their wives, and the distribution is skewed heavily towards men being the older spouse (Figure 1b).

We calibrate the fraction of men and women in each sector to the corresponding data distribution. Since all individuals in the model are covered by AFP (the early retirement option) either from the private or the public sector, we also exclude non-covered individuals in the data when we find the ratios. We then get an average of the four different couple compositions: both have AFP in private

[18]By the sample size for each year, to population growth causing a decrease in the risk measures over time.

[19]Since the spillover coefficient in Kruse (2020) is a key target in the internal calibration, we choose the same sample of individuals as in that paper. See Section 4.2.

sector, both have AFP in public sector and one spouse is covered by private AFP while the other is covered by public AFP. There is a significant difference in the fraction of men and women who work in the two sectors, which means that the group where the husband is covered by private AFP and the wife by public AFP is larger than vice versa.[20]

Labor earnings process The labor earnings process is estimated on monthly earnings for all Norwegian individuals (both married and unmarried) above age 22 from 1997–2014. The data also contains information about registered contracted hours, from which we obtain a measure of hourly wages. The equation map to the data is the empirical analog to the endowment ω in (7), where we have approximated the age effect by a fourth degree polynomial:

$$\ln(w_{ijt}) = \alpha + \beta_1 j + \beta_2 j^2 + \beta_3 j^3 + \beta_4 j^4 + \psi_t \mathbb{D}_t + e_{ij} \tag{19}$$

The subscripts (i, j, t) refer to an individual i of age j in year t. The dependent variable is hourly wages, $\mathbb{D}_t$ a set of year t dummies and ϵ_{ij} represents the idiosyncratic shock component of earnings. We estimate the coefficients by running OLS on (19), separately for men and women, and for public and private sector workers with access to AFP. The estimated age-coefficients from these four regressions are reported in Table 2. Our deterministic age-earnings profiles in (7) then follow:

$$q_{g,j_g} = \hat{\alpha} + \hat{\beta}_{g1} j + \hat{\beta}_{g2} j^2 + \hat{\beta}_{g3} j^3 + \hat{\beta}_{g4} j^4 \tag{20}$$

As in the life-cycle model with endogenous retirement in İmrohoroğlu and Kitao (2012), we set labor efficiency $\omega_{g,j_g} = 0$ for individuals older than 69, implicitly assuming that agents do not work past age 70. In the data, very few work

[20] In the data, the group where the husband has private AFP, while the wife does not is 16 percent of the couples, while the group where the wife has private AFP and the husband does not is about 10 percent.

after reaching that age.[21] The resulting polynomial coefficients from estimating (19) using OLS are reported in Table 2.

Table 2
Income profile polynomial regression. Coefficients from OLS on (19).

	α	β_1	β_2	β_3	β_4
Private sector, men	2.953	0.129	-0.003	0.00004	$-1.96{\cdot}10^{-7}$
Public sector, men	1.632	0.261	-0.008	0.0001	$-4.77{\cdot}10^{-7}$
Private sector, women	2.100	0.201	-0.005	0.00006	$-2.89{\cdot}10^{-7}$
Public sector, women	1.892	0.160	-0.005	0.00006	$-2.91{\cdot}10^{-7}$

Notes: Collected from Norwegian register data, all Norwegian tax payers from 1997–2014. Regressions include year dummies.

The stochastic earnings process in AR(1) process in (8) is estimated on the residuals, $\hat{e}_{ij}$, obtained from the OLS regressions on (19). The resulting OLS estimates for the persistency and standard deviation of the shock is reported in Table 1.

$$\hat{e}_{ij} = \nu \hat{e}_{ij-1} + \varepsilon_{ij}, \tag{21}$$

$$\varepsilon_{sgj} \sim N(0, \sigma^2_{\varepsilon_{sgj}}) \tag{22}$$

From our estimations, we see that men have more persistent shocks than women, and somewhat more dispersion. Even more characterizing for our processes is that public sector workers have safer income profiles.[22]

Pension system before and after the reform When parameterizing (10) we differentiate between two types of earnings-tests broadly consistent with the pre-reform pension system. Between ages 62–66 (i.e. the early retirement scheme), workers lose their entire pension if they continue to work. Between ages 67–69, pension is reduced with 40 percent of labor earnings above $2G$, where G denotes

[21] Due to selection bias, it would be difficult to estimate the labor productivity past age 70. Additionally, under the Norwegian Working Environment Act, workers are protected from unfair dismissal, e.g. due to age. However, during the time period we calibrate the model to, workers are not protected by this act when reaching age 70.

[22] This is e.g. consistent with the fact that the public sector wages are more tightly linked to centralized wage negotiations compared to private sector wages.

the basic amount in the Norwegian pension system.[23]

$$f_{j_g}(\hat{b}_g, y_{g,j_g}) = \begin{cases} 0 & \text{if} \quad j_g < 62 \\ \hat{b}_g(1 - l_{g,j_g}) & \text{if} \quad 62 \geq j_g < 67 \\ \max(0, \hat{b}_g - 0.4\max(0, y_{gj_g} - 2G)) & \text{if} \quad 67 \geq j_g < 70 \\ \hat{b}_g & \text{if} \quad 70 \geq j_g \end{cases} \tag{23}$$

In the internal calibration, we implement the pension reform as a complete removal of the earnings test. This arrives as a shock at the beginning of 2011, and applies to workers in the private sector.[24]

Tax system We set $\tau_a = 0.28$, in line with the capital income tax during our sample period. The tax function T takes into account the progressivity of the Norwegian tax schedule. In particular, it consists of a national insurance contribution, a general income tax, and a surtax on high income earners. The national insurance contribution amounts to 7.8 and 3 percent of labor earnings and pension income, respectively. The general income tax is computed based on taxable income, which is total labor and pension income net of a deductible amount. The deductible amount varies between $0.55G$ and $1.55G$, and is increasing in share of income that comes from labor income. The taxable amount is then taxed at 28 percent. Finally, the surtax on high income earners follows a three-bracket schedule, with marginal tax rate of 0 (until $6G$), 9 (until $10G$) and 21 percent (above

[23]The basic amount G is calibrated to be 18 percent of the average labor earnings among individuals between age 40–44, consistent with the data

[24]There are two other key elements in the pension reform implicitly affecting the level of pre-earnings tested pension benefits $\hat{b}_g$, which we assume to remain unaffected by the reform: (i) pension benefits are longevity adjusted and (ii) delayed claiming increases future annual pension benefit. Regarding (i), the actual pension reform was calibrated such that individuals who were 67 in 2011 (the 1944 cohort) received the same level of life-time benefits as in the pre-reform system. Later cohorts will receive less, roughly proportional to the growth in remaining life-expectancy at age 62. Since we focus on the cohorts 1944–1952 our assumption that the pension benefit remains unchanged is broadly consistent with the actual reform. Regarding (ii) delayed claiming in the new system leads to a roughly actuarial fair re-adjustment of annual pension benefit, leaving individuals approximately indifferent with respect to the timing of pension claiming. We therefore abstract from a pension claiming decision, assuming that individuals never delay claiming.

$10G$).[25]

Interest rate, full-time work, and wealth partition The time spent working full time is $\theta = 0.4$, which means that 40 percent of the time endowment is lost to full time work. The interest rate is set to $r = 0.04$. Finally, we follow Hong and Ríos-Rull (2012) and assume that, upon divorce, household wealth is divided equally among the two spouses, $\pi = 0.5$.[26]

External preference calibration Following e.g. Low and Pistaferri (2015) and Low et al. (2018) we set $\gamma = 1.5$.[27]. For the leisure curvature parameter we set $\phi = 3$, which is the value used in Heathcote et al. (2010), in a model with separable spousal leisure in utility.[28]

The value for the equivalence scale in the couples' utility function is set to $\zeta = 1.3$, in line with the estimate in Hong and Ríos-Rull (2012). This implies that a couple would need to spend only 30 percent more than a single individual to draw the same utility from consumption.

4.2 Internal preference calibration

The remaining parameters are calibrated by matching simulated model moments to corresponding data moments. For this purpose we simulate 9 cohorts (1944–1952), and restrict the sample to individuals who are still working at age 60. As the spillover coefficient in Kruse (2020) is a key target, the sample restriction is taken directly from that paper, and applied to both the data and simulated moments.

[25]In the post-pension reform regime, workers combine full time work with untested pension income. To avoid that this would place individuals in the surtax schedule, we levy this only on labor income. This is broadly in line with adjustments made in the actual Norwegian tax system.

[26]This partition rule is in line with the Norwegian Family Law.

[27]The value $\gamma = 1.5$ is in line with estimates reported in e.g. Attanasio and Weber (1995) and Attanasio et al. (1999).

[28]$\phi = 3$ in line with commonly used values, referring to the Frisch elasticity of labor supply. In a preference specification such as in (3) with $\phi = 3$, $\rho = 1 - \phi$ (no complementarity) and $\chi = 0$, assuming continuous labor supply choice (h) , the Frisch elasticity is $\frac{1-h}{h}\frac{1}{\phi}$. The elasticity of an individual working full time hours would be 0.5.

4.2.1 Couple parameters

The parameters to be estimated are $\beta, \eta, \kappa, \chi_0$, and ρ. The first four parameters targets pre-reform data, while the latter parameter targets the spousal labor supply spillover effect of the pension reform. When simulating the model, we assume that the simulated cohorts 1944–1952 were born into the pre-reform pension system, and then invoke the pension reform as a shock in 2011.

The discount factor β targets average wealth to labor income ratio at age 60. The weights on wife's and total leisure, η and κ target average retirement rate among males and females in age group 63–65. The rate of decline in time endowment χ_0 targets the average growth rate in retirement rates between age groups 63–65 and 66–68.[29] In the data we classify an individual as retired when annual labor income is below $1G$.[30]

The remaining parameter, ρ, governing the degree of complementarity in leisure, targets the reduced form spillover effects of the Norwegian pension reform. Identification relies on the fact that public sector workers were not affected by the reform. We then estimate spousal spillover effects by comparing the employment of an individual married to spouse who works in the private sector with AFP (i.e. a treated spouse) to those married to a spouse who works in the public sector (i.e. a non-treated spouse). When calibrating ρ we implement a reduced form method on simulated data from our model.[31]

4.2.2 Single parameters

The parameters κ^S and χ_0^S targets average retirement for the age group 63–65 and the growth in retirement rate between age 63–65 to 66–68. We use individuals born in the 1944–1952 cohorts who separate through divorce or the death of their spouse between ages 50–59.

[29] Since we focus on retirement patterns for individuals who still work at age 60, we assume that the decline starts at age 61 by setting $j^\chi = 61$.

[30] This is consistent with the definition of retirement used in Kruse (2020).

[31] Specifically, we follow the empirical setup presented in Section 2.2.1.

5 Quantitative Results

We now present our main results. We invoke the structural model for three main reasons: (*i*) to structurally identify the degree of leisure complementarity using quasi-experimental data on spousal spillover effects (*ii*) to explain the synchronization of retirement choice observed in data (*iii*) to quantify the importance of joint retirement for aggregate labor supply elasticities. Sections 5.2–5.4 are organized around these three issues. Before we proceed, we present the calibrated model in Section 5.1.

5.1 Discussion of the estimated model

Table 3
Internally calibrated model parameters.

Parameter	Value	Target moment	Data value*	Model value
Couple parameters complementarity				
ρ	−3.861	Retirement spillover estimate in Kruse (2020)	0.11	0.11
κ	0.353	Mean retirement at age 63–65, women	0.38	0.38
η	0.471	Level difference of retirement between 63–65 year old men and 63–65 year old women	0.05	0.05
χ_0	0.0277	Growth in retirement from 63–65 to 66–68	1.91	1.91
Single parameters				
κ^S	0.298	Mean retirement at age 63–65, single	0.26	0.26
χ_0^S	0.0402	Growth in retirement from 63–65 to 66–68, single	2.25	2.25
Shared parameters				
β	0.965	Wealth-to-income ratio (at age 60)	1.87	1.87

*Source: Norwegian register data.
Notes: Data values are computed on a sample consisting of married couples where at least one spouse is born in 1944–1952, both work at age 60. The sample corresponds to the estimation sample used in Kruse (2020).

Table 3 reports the parameter values and model simulated moments, produced by the calibration. The average (across gender) empirical point estimate implies that individuals are 12 percentage point more likely to work, if their spouse also works. Our model calibration, replicating the reduced form data sample and experimental setting, reproduces the same spillover effect. The resulting value for the preference parameter governing non-separability is $\rho = -3.23$, implying that leisure complementarity is needed in order to fit the reduced form evidence.[32]

Next, we target the mean retirement at age 63–65 for both men and women separately. In the data, women retire earlier on average than men. These moments are combined to calibrate the weights on household total leisure κ and the relative weight on the wife's leisure component η. In practice, we set κ to match the mean retirement of women in this age group, and then adjust η to match the relative level difference of men and women. We precisely match the empirical moments using these parameters, with $\kappa = 0.297$ and $\eta = 0.478$.

We then target the growth in retirement from 63–65, which is an equally weighted average of the growth for men and women. We use this to calibrate our time endowment parameter χ_0, governing the increase in labor utility cost among elderly workers. We match this growth by setting $\chi_0 = 0.0295$, which can be interpreted as the percentage loss of time endowment for each year surpassing age 61. Figure 2 shows that our model retirement profile performs well across all ages and both genders, even though only targeting average retirement rates between 63–65 and 66–68.

Finally, we target the average wealth-to-income ratio in our empirical sample using the parameter β. We match the empirical wealth-to-income ratio with a reasonable parameter value of $\beta = 0.974$.

Parameters specific to separated individuals are calibrated to be somewhat smaller (κ^S) and larger (χ_0^S) compared to the corresponding couple parameters. These parameter values imply that separated individuals value leisure relatively less, but face a steeper decline in time endowment in old age. The parameter adjustment is needed in order to rationalize that separated individuals on average postpone retirement relative to couples.

5.2 The estimated complementarity parameter

As we have shown, our benchmark model generates a substantial leisure complementarity effect in order to be consistent with the reduced form spillover effects.

[32] As explained in Section 3.1.1 spousal leisure are complements if $1 - \phi - \rho > 0$.

Figure 2
Average retirement, data versus model

(a) Men **(b)** Women

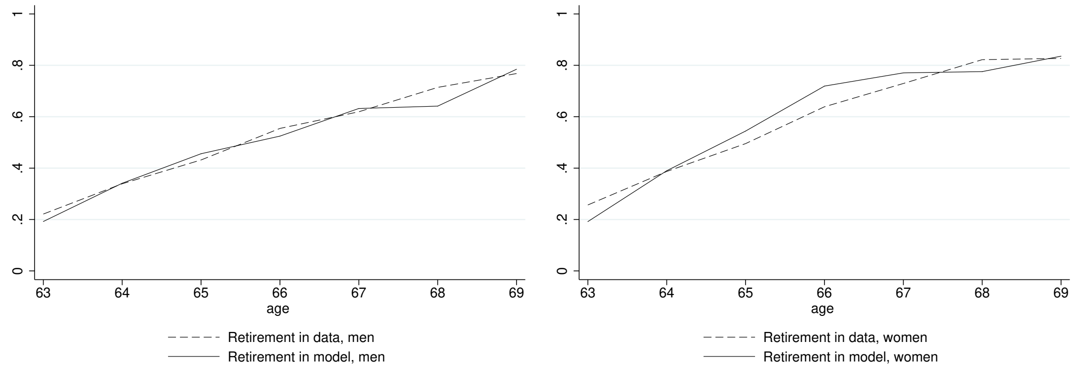

Notes: Retirement in data is calculated on the couples where at least one partner is born in 1944–1952. Both spouses must be observed working at age 60. Retirement in data is defined as having annual labor income less than one base point (G), as defined in Footnote 6. Similar restrictions are made to the model simulated couples.

To understand the role of ρ in shaping this result, we now consider an alternative calibration in which we impose separability in the couple utility function. This involves setting $\rho = -2$, and re-calibrating the model without targeting the reduced form coefficient. The resulting parameter values are reported in Table 4.

Table 4
Internally calibrated model parameters (separable leisure).

Parameter	Value	Target moment	Data value*	Model value
Couple parameters separability				
κ	0.279	Mean retirement at age 63–65, women	0.38	0.38
η	0.484	Level difference of retirement between 63–65 year old men and 63–65 year old women	0.05	0.05
χ_0	0.0406	Growth in retirement from 63–65 to 66–68	1.91	1.92
Single parameters				
κ^S	0.302	Mean retirement at age 63–65, single	0.26	0.26
χ_0^S	0.0397	Growth in retirement from 63–65 to 66–68, single	2.25	2.25
Shared parameters				
β	0.965	Wealth-to-income ratio (at age 60)	1.87	1.87

*Source: Norwegian register data.
Notes: Data values are computed on a sample consisting of married couples where at least one spouse is born in 1944–1952, both work at age 60. The sample corresponds to the estimation sample used in Kruse (2020).

We observe that with separability, the model generates a negative spillover effect of -0.29, implying that the probability of retirement is increased with 29 percentage point when your partner is working. This negative spillover comes

Figure 3
Interaction between ρ and key target moments in the calibration.

(a) Spillover coefficient **(b)** Elasticity of target moments

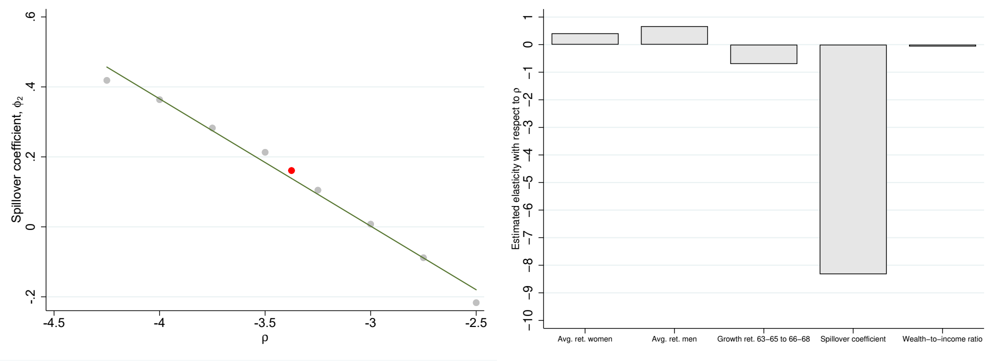

Notes: Panel 3a shows the model simulated spillover effect for different levels of ρ. Spillover effect computed as in Kruse (2020) in Equations 1 and 2. The red dot displays our calibration of $\rho = -3.23$.

Panel 3b shows the implied elasticity of each target moment with respect to a one percent change in ρ.

entirely from the income effect, as household income is higher when the partner works. Moreover, this highlights that zero spillover coefficient in the reduced form setting is not necessarily evidence of separability. Complementarity is needed to counteract the income effect.

To provide further insight into how ρ is pinned down by the spillover data moment, Figure 3a displays the sensitivity of the corresponding model moment with respect to ρ. The dots represent the simulated spillover coefficient for different values of ρ, fixing all other parameter at their benchmark values in Table 3. Clearly, the parameter is locally well identified. The figure reveals a nearly linear relationship between the parameter and the simulated moment. Furthermore, Figure 3b shows the elasticity of all moments with respect to ρ, revealing a particularly tight link between ρ and spousal spillover. The elasticity of the spillover moment more than eight times higher than elasticity of other moments.

5.3 Within-couple age gaps

As we discussed in Section 2.2, empirical evidence shows that even though the age gap distribution is rather disperse, spouses tend to synchronize retirement. Moreover, there is a striking feature of a "spike" in retirement within the same calendar year. In Figure 4, we compare the performance of our calibrated model with complementarity and with separability in leisure to the estimated distributions from Norwegian register data. Despite that this is an out-of-sample test, the model with complementarity still replicates the "spike" in synchronized retirement within the same year.

Additionally, by comparing the distribution of synchronized retirement in the model with complementarity to the model with separability in leisure, we clearly see the importance of including complementarity in leisure to our model. Without leisure complementarity, the model cannot replicate the feature of synchronized retirement at all, in fact it predicts that couples tend to avoid retiring within the same calendar year, due to a negative income effect.[33]

5.4 Aggregate implications

In this section, we evaluate the importance of leisure complementarity and joint retirement for the aggregate labor supply elasticity of elderly workers. We do so in the context of the Norwegian pension reform. As we explained in Section 2.1, Norway removed a confiscatory earnings test on continued work in 2011. In our calibration, we constructed a simulated data panel over the time period 2007–2015. When doing so, we assumed that the simulated cohorts 1944–1952 were born into the pre-reform pension system, and then invoked the pension reform as a shock in 2011. We now proceed in two steps to compute the role of complementarity. First, we generate two hypothetical simulated panels from our baseline model: (*i*) In the first panel, we simulate the full life-cycle of all cohorts without implementing the pension reform, and (*ii*) in the second panel, we simulate the life-cycles, assuming that the individuals are born into the post-reform pension system. The long-run reform-induced increase in retirement age can then be assessed by comparing the average retirement rates in the two panels. In the second step, we redo the exercise for the model in which we impose separability.

[33]The corresponding figure for the model with separable leisure, when allowing for correlated earnings shocks, is displayed in Appendix A. The joint retirement pattern is essentially unchanged when incorporating a correlation of 0.05 (in line with the empirical evidence). We also show that even at a counterfactually high correlation of 0.9, the model with separability vastly under-predicts the degree of synchronization.

Figure 4

Joint exit in model with leisure complementarity (panel a) versus model with leisure separability (panel b). Panels (c) and (d) display the joint exit in data and the age gap distribution.

(a) Joint exit, model with complementarity

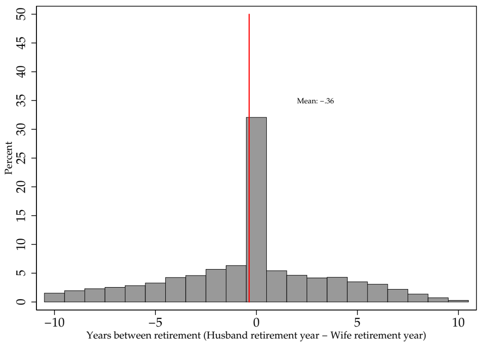

(b) Joint exit, model with separability

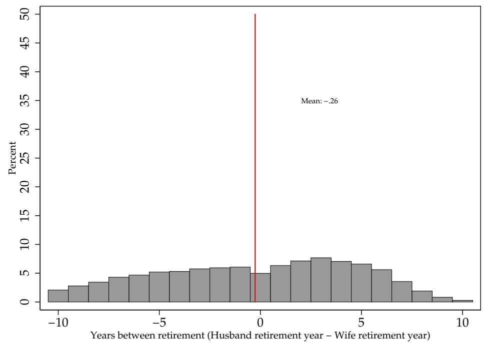

(c) Joint exit, data

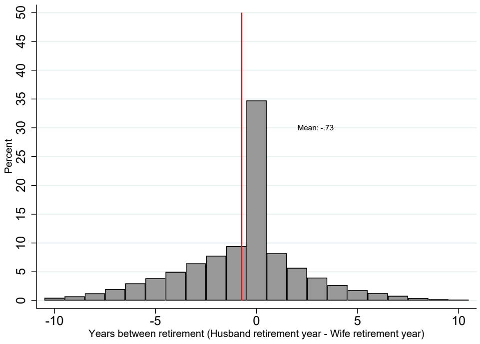

(d) Age gap distribution

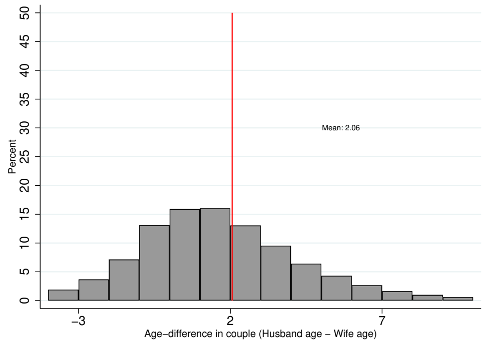

Notes: The red lines show the mean in each distribution. Data estimations are similar to Figure 1. Norwegian register data on all married couples where at least one partner is born in 1944–1952. Both spouses must work at age 60.

The increase in average retirement age in the baseline model is 28 percent larger compared to the model with separability. Consequently, 38 percent of the aggregate labor supply is due synchronization of retirement behavior arising from leisure complementarity. To derive these results we simulate an economy in which the share of treated individuals (private sector workers) are as in the 2011 reform.[34]

However, the influence of joint retirement for the aggregate employment response, is dependent on whether either or both of the spouses are directly affected by the reform. In the former case, the non-affected individual is only indirectly exposed to the reform. Intuitively, the relative importance of complementarity is

[34]These results are based on couples. Complementarity is less important for the overall labor supply effect if we also include separated individuals (it then accounts for 22 percent of the response).

larger in these couples compared to couples in which both are directly affected. Our simulation results confirm this intuition. Among couples in which both are directly exposed, complementarity accounts for 18 percent of the labor supply response. By contrast, in couples with only one directly exposed spouse, the importance of complementarity is nearly twice as large (34 percent). The reason is that without complementarity, the indirectly exposed spouse face only a negative income effect, caused by the increased labor supply of the affected spouse.

6 Conclusion

In this paper, we build a quantitative life-cycle model to assess the labor supply implications of leisure complementaries and synchronized retirement. Individuals in the model lives in a multi-person household and draw utility from consumption and non-separable leisure. Non-separability implies that household labor supply behavior is interlinked, as spouses influence each others utility valuation of leisure. If spousal leisure are complements, an individual's valuation of leisure increases if the spouse is retired. This feature generates a feedback mechanism potentially amplifying the household labor supply response to work incentives.

To pin down the magnitude and direction of non-separability, our empirical approach exploits exogenous variation in work incentives arising from a recent Norwegian pension reform. In particular, within our model we replicate the quasi-experimental setting generated by this reform. We find that a substantial leisure complementarity is need for the model to reproduce the reduced form spillover effects found in Kruse (2020). Importantly, the model is simultaneously able to account for the frequency of joint retirement observed in the data. The implied complementarity accounts for roughly one-third of the long-run aggregate labor supply response to the pension reform.

Our results show that spillover effects are much larger among couples in which only one spouse face increased work incentives. Many labor market reforms are effectively targeted as opposed to universal (e.g. policies aimed at reducing early retirement through welfare benefits such as disability pension). Our findings indicate that such reforms may have substantial positive indirect effects, through the spousal spillover channel.

Finally, our model allows us compare the differential labor supply elasticity of separated and married individuals. In fact, the reform-induced labor supply response of the latter group is 15 percent larger than for single households. Heterogeneous elasticity across household types imply that e.g. rising divorce rates may have implications for aggregate employment. A simple back-of-the-envelope calculation based on these numbers, indicate that for each percent divorce rates increase, the aggregate labor supply response drops with 0.15 percent.

References

Attanasio, O. P., Banks, J., Meghir, C. and Weber, G. (1999), Humps and Bumps in Lifetime Consumption, *Journal of Business & Economic Statistics 17*(1), 22–35.

Attanasio, O. P. and Weber, G. (1995), Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey, *Journal of Political Economy 103*(6), 1121–1157.

Banks, J., Blundell, R. and Casanova Rivas, M. (2010), The dynamics of retirement behavior in couples: Reduced-form evidence from England and the US, *UCLA Working paper (July 1, 2010)* (Unpublished).

Blundell, R., Costas Dias, M., Meghir, C. and Shaw, J. (2016), Female Labor Supply, Human Capital and Welfare Reform, *Econometrica 84*(5), 1705–1753.

Borella, M., De Nardi, M. and Yang, F. (2018), The aggregate implications of gender and marriage, *The Journal of Economics of Ageing 11*, 6–26.

Capatina, E. (2015), Life-cycle effects of health risk, *Journal of Monetary Economics 74*(September), 67–88.

Casanova, M. (2010), Happy Together: A Structural Model of Couples' Joint Retirement Choices, *UCLA Working Paper* (Mimeo).

Coile, C. (2003), Retirement incentives and couples' retirement decisions, Technical report, National Bureau of Economic Research.

Cooley, T. F. and Henriksen, E. (2018), The demographic deficit, *Journal of Monetary Economics 93*(C), 45–62.

Cooley, T. F., Henriksen, E. and Nusbaum, C. (2019), Demographic Obstacles to European Growth, *March 25*.
URL: *https://dx.doi.org/10.2139/ssrn.3356824*

Cubeddu, L. and Ríos-Rull, J.-V. (2003), Families as shocks, *Journal of the European Economic Association 1*(2-3), 671–682.

Eckstein, Z., Keane, M. and Lifshitz, O. (2016), Sources of Change in the Life-Cycle Decisions of American Men and Women: 1962–2014, (Mimeo).

Eckstein, Z. and Lifshitz, O. (2011), Dynamic Female Labor Supply, *Econometrica 79*(6), 1675–1726.

Fernández, R. and Wong, J. C. (2011), The Disappearing Gender Gap: The Impact of Divorce, Wages, and Preferences on Education Choices and Women's Work, *NBER Working Paper Series 17508*.

Fernández, R. and Wong, J. C. (2014), Divorce Risk, Wages and Working Wives: A Quantitative Lifecycle Analysis of Female Labour Force Participation, *The Economic Journal 124*(May), 319–358.

French, E. (2005), The Effects of Health, Wealth and Wages on Labour Supply and Retirement Behaviour, *Review of Economic Studies 72*, 395–427.

French, E. and Jones, J. B. (2011), The Effects of Health Insurance and Self-insurance on Retirement Behavior, *Econometrica 79*(3), 693–732.

Guner, N., Kaygusuz, R. and Ventura, G. (2012), Taxation and Household Labour Supply, *The Review of Economic Studies 79*(3), 1113–49.

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.

Heathcote, J., Storesletten, K. and Violante, G. L. (2010), The Macroeconomic Implications of Rising Wage Inequality in the United States, *Journal of Political Economy 118*(4), 681–722.

Hernæs, E., Jia, Z., Piggot, J. and Vigtel, T. C. (2019), Flexible Pensions and Labor Force Withdrawal, *ARC Centre of Excellence in Population Ageing Research Working Paper 2019/3*.

Hernæs, E., Markussen, S., Piggot, J. and Røed, K. (2016), Pension Reform and Labor Supply, *Journal of Public Economics 142*(October), 39–55.

Hong, J. H. and Ríos-Rull, J.-V. (2012), Life Insurance and Household Consumption, *American Economic Review 102*(7), 3701–30.

Hospido, L. (2015), Pension reform and couples' joint retirement decision, *IZA World of Labor 142*.

Hurd, M. D. (1990), The Joint Retirement Decision of Husbands and Wives, *in* D. A. Wise, ed., 'Issues in the Economics of Aging', University of Chicago Press, pp. 231–258.

İmrohoroğlu, S. and Kitao, S. (2012), Social Security Reforms: Benefit Claiming, Labor Force Participation, and Long-run Sustainability, *American Economic Journal: Macroeconomics 4*(3), 96–127.

Kitao, S. (2014), Sustainable social security: Four options, *Review of Economic Dynamics 17*(4), 756–779.

Kruse, H. (2020), Joint Retirement in Couples - Evidence of Complementarity in Leisure, *The Scandinavian Journal of Economics* (Forthcoming).

Lalive, R. and Parrotta, P. (2017), How does pension eligibility affect labor supply in couples?, *Labour Economics 46*, 177–188.

Low, H., Meghir, C., Pistaferri, L. and Voena, A. (2018), The aggregate implications of gender and marriage, *NBER Working Paper Series 24356*(February).
URL: *http://www.nber.org/papers/w24356*

Low, H. and Pistaferri, L. (2015), Disability Insurance and the Dynamics of the Incentive Insurance Trade-Off, *American Economic Review 105*(10), 2986–3029.

Martín, A. R. S. and Marcos, V. S. (2010), Demographic Change and Pension Reform in Spain: An Assessment in a Two-Earner, OLG Model, *Fiscal Studies 31*(3), 405–452.

Michaud, P.-C., Van Soest, A. and Bissonnette, L. (2018), Understanding joint retirement, *Journal of Economic Behavior and Organization In Press*(29 July 2019 (online)).

OECD (2017), Pensions at Glance 2017: OECD and G20 Indicators, *OECD Publishing, Paris* .

OECD (2019), Pensions at a Glance 2019: OECD and G20 Indicators, *OECD Publishing, Paris* .

Stancanelli, E. and van Soest, A. (2012a), Retirement and Home Production: A Regression Discontinuity Approach, *American Economic Review, American Economic Association 102*(3), 600–605.

Stancanelli, E. and van Soest, A. (2012b), Joint Leisure Before and After Retirement: A Double Regression Discontinuity Approach, *IZA Discussion Papers 6698*(Institute for the Study of Labor (IZA)).

Tauchen, G. and Hussey, R. (1991), Quadrature-Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models, *Econometrica 59*(2), 371–396.

van der Klaauw, W. and Wolpin, K. I. (2008), Social security and the retirement and savings behavior of low-income households, *Econometrica 79*(6), 1675–1726.

Appendix A Correlated shocks

Figure A.1
Joint exit in model with correlated earnings shocks and separable leisure.

(a) Joint exit, (calibrated) low correlation **(b)** Joint exit, high correlation

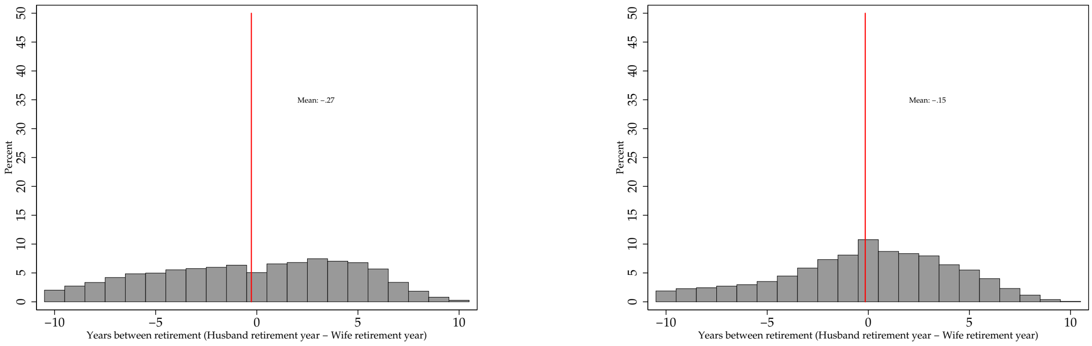

Notes: In panel a, the joint income shocks have a calibrated correlation coefficient of 0.05. In panel b, the joint income shocks have a correlation coefficient of 0.9.

In Figure A.1, we show that correlated earnings shocks within the couple cannot explain the characteristic "spike" in retirement synchronization. In panel A.1a, we have calibrated the correlated earnings shocks. We use a sample of all individuals who are aged 55 to 60 years old and their respective partners. We normalize employment at age 55, meaning that everyone starts out as employed. To avoid bias of zero-earnings individuals due to take-up of disability insurance, we disregard couples where either partner receives DI within the time-period. Apart from these restrictions, all other restrictions are similar to those made throughout the paper: only one partner may be covered by private sector AFP, and cohorts are restricted to the 1944–1952 cohorts (for the spouse who starts out aged 55). We then compute the following correlations:

$$\rho_{60,55} = \frac{cov(Y_{\Delta_{60,55},w}, Y_{\Delta_{60,55},m})}{\sigma_{\Delta_{60,55},w}\sigma_{\Delta_{60,55},m}} \quad \text{(A. 1)}$$

where $Y_{\Delta_{60,55},g}$ is the change in earnings for spouse g from age 55 to 60 and σ is the standard deviation of this, and:

$$\rho_{a+1,a} = \frac{cov(Y_{\Delta_{a+1,a},w}, Y_{\Delta_{a+1,a},m})}{\sigma_{\Delta_{a+1,a},w}\sigma_{\Delta_{a+1,a},m}} \tag{A. 2}$$

where $Y_{\Delta_{a+1,a},g}$ is the change in earnings for spouse g from age a to $a+1$ and σ is the standard deviation of this. We obtain $\rho_{60,55} = 0.054$ and $\rho_{a+1,a} = 0.039$.