Network for Studies on Pensions, Aging and Retirement

Netspar DISCUSSION PAPERS

Bo Honoré and Áureo de Paula

Joint Retirement in Europe

DP 10/2014-052

Joint Retirement in Europe*

Bo E. Honoré †

Princeton University

Áureo de Paula ‡

UCL, São Paulo School of Economics-FGV, CeMMAP, IFS and CEPR

This Version: October 2014

Abstract

In this project we study joint retirement of couples in Europe. We characterize various empirical regularities and use a model for simultaneous duration variables developed in Honoré and de Paula (2014). Whereas conventionally used duration models cannot account for joint retirement, our model admits joint retirement with positive probability, allows for simultaneity and nests the traditional proportional hazards model. In contrast to other statistical models for simultaneous durations, it is based on Nash bargaining and it is therefore interpretable in terms of economic behavior. We focus on the Survey on Health, Ageing and Retirement in Europe (SHARE) and the English Longitudinal Survey of Ageing (ELSA).
JEL Codes: J26, C41, C3.

*We acknowledge financial support from Netspar. Bo Honoré also acknowledges financial support from the National Science Foundation and from the Gregory C. Chow Econometric Research Program. Áureo de Paula also acknowledges financial support from the Economic and Social Research Council through the ESRC Centre for Microdata Methods and Practice grant RES-589-28-0001 and the European Research Council through the ERC Starting Grant 338187. Priscilla Fialho provided very able research assistance.

†Princeton University (Department of Economics), Princeton, NJ 08542. E-mail: `honore@princeton.edu`

‡University College London (Department of Economics), São Paulo School of Economics-FGV, C*e*MMAP and IFS, London, UK. E-mail: `apaula@ucl.ac.uk`

1 Introduction and Related Literature

This project investigates the determinants of joint couples' retirement decisions in Europe. A majority of retirees are married and many studies indicate that a significant proportion of individuals retire within a year of their spouse. Many articles have documented the joint retirement of couples, mostly in the United States and United Kingdom. A few notable examples (and datasets employed) include Hurd (1990) (New Beneficiary Survey); Blau (1998) (Retirement History Study); Gustman and Steinmeier (2000) (National Longitudinal Survey of Mature Women); Michaud (2003) and Gustman and Steinmeier (2004) (Health and Retirement Study); and Banks, Blundell, and Casanova Rivas (2007) (English Longitudinal Study of Ageing, ELSA). In this research, we investigate the joint retirement of couples in Europe using data from the Survey on Health, Ageing and Retirement in Europe (SHARE) and from ELSA.

In the Health and Retirement Study for example, 55% of respondents report that they expect to retire at the same time as their spouses.[1] There are at least two distinct reasons why couples might retire simultaneously. One is that the partners receive correlated shocks (observable or not) inducing retirement at similar times. The other is that retirement is jointly decided, reflecting the taste and budget interactions of both members of the couple. For example, institutional reasons related, say, to pensions and social security features, health insurance, etc., will provide constraints and incentives under which retirement decisions are undertaken. Those will affect the timing of retirement not only through embedded incentives that are common across potential retirees in the population, but also provide guidance in the decision making of partners whose preference for leisure may complement each other's. Other sources of commonalities include assortative mating, representing a source of common unobserved heterogeneity.

The distinction between these two drivers of joint retirement (which are not mutually exclusive) is similar to the motivation for studying linear simultaneous equation models, and it parallels the categorization by Manski (1993) of correlated and endogenous (direct) effects in social interactions. There, as in this article, discerning these two sources of correlation in outcomes is relevant for analytical and policy reasons. For example, if the estimated model does not allow for the joint decision by the couple, then the estimate of the effect of a retirement-inducing policy shock will be

[1] The figure corresponds to those who answer either YES or NO to the question: "Do you expect your spouse to retire at about the same time that you do?" (`R1RETSWP`). It excludes those whose spouse was not working.

biased if the retirement times are indeed chosen jointly. Furthermore, the multiplier effect induced by the endogenous direct effect of husband on wife or vice-versa is an important conduit for policy. The quantification of its relative importance is then important for both methodological and substantive reasons.

Unfortunately, standard econometric duration models are not suitable for the analysis of joint durations with simultaneity of the kind we have in mind. One tempting estimation strategy is to include the spouse's retirement date or, in the case of a hazard model, a time-varying variable indicating his or her retirement date. Because this variable is a choice that is potentially correlated with the unobservable variables determining a person's own retirement, standard estimators are bound to be inconsistent. Essentially, this approach would amount to including an endogenous variable from a simultaneous equation model in the right-hand side of a regression. To address this issue we employ an econometric duration model that allows for simultaneity developed in Honoré and de Paula (2014). As in the linear simultaneous equation model, identification is obtained using exclusion restrictions and, in our particular case, using the timing patterns in the data. We briefly describe the methodology below and refer to Honoré and de Paula (2014) for futher details.

The broader literature on retirement is abundant, and a number of papers focusing on retirement decisions in a multi-person household have appeared in the last decades. Hurd (1990) presents one of the early documentations of the joint retirement phenomenon. Later papers confirming the phenomenon and further characterizing the correlates of joint retirement are Blau (1998); Michaud (2003); Coile (2004a); and Banks, Blundell, and Casanova Rivas (2007). Gustman and Steinmeier (2000) and Gustman and Steinmeier (2004) work with a dynamic economic model in which the husband's and wife's preferences are affected by their spouse's actions, but the couple makes retirement decisions individually.[2] These papers focus on Nash equilibria to the joint retirement decision, i.e. each spouse's retirement decision is optimal given the other spouse's timing and vice-versa.[3] More recently, Gustman and Steinmeier (2009) present a richer (non-unitary) economic model with a solution concept that differs from a Nash equilibrium and is guaranteed to exist and be unique. Michaud and Vermeulen (2011) estimate a version of the "collective" model introduced by Chiap-

[2]In the family economics terminology, their model is a non-unitary model in which people in the household make decisions individually. In unitary models, the household is viewed as a single decision-making unit. A characterization of unitary and non-unitary models can be found in Browning, Chiappori, and Lechene (2006).

[3]When more than one solution is possible, they select the Pareto dominant equilibrium, i.e., for all other equilibria at least one spouse would be worse off. If no equilibrium is Pareto dominant, the equilibrium where retirement by at least one household member happens earliest is assumed (see, e.g., Gustman and Steinmeier (2000), pp. 515, 520).

pori (1992) in which (static) labor force participation decisions by husband and wife are repeatedly observed from a panel (i.e., the Health and Retirement Study). Casanova Rivas (2010) suggests a detailed unitary economic dynamic model of joint retirement. Coile (2004b) presents statistical evidence on health shocks and couples' retirement decisions and Blau and Gilleskie (2004) present an economic model that also focuses on health outcomes and couples' retirement decisions. To the best of our knowledge, ours is the first exploration of joint retirement using SHARE data.

In our estimated model, we assume that retirement decisions are made through Nash bargaining on the retirement date. The Nash solution corresponds to a set of behavioral axioms on the bargaining outcomes (essentially Pareto efficiency, independence of irrelevant alternatives and symmetry), and it is widely adopted in the literature on intra-household bargaining. It can be shown that this solution approximates the equilibrium outcome of a situation in which husband and wife make offers to each other in an alternating order, and the negotiation breaks down with a certain probability. As this probability goes to zero, the equilibrium converges to the Nash solution (see Binmore, Rubinstein, and Wolinsky (1986)). Though this solution also leads to Pareto efficient outcomes, it imposes more structure than Casanova Rivas (2010) or Michaud and Vermeulen (2011) [see Chiappori (1992) and Chiappori, Donni, and Komunjer (2012)].

The econometric model used in our study builds on the developments in Honoré and de Paula (2014). Although admittedly stylized, it extends traditional duration models that are well understood and reasonably straightforward to estimate. The model extends the usual statistical framework in a way that allows for *joint* termination of *simultaneous* spells with positive probability. In the usual hazard modeling tradition, this property does not arise. One could appeal to existing statistical models (e.g., Marshall and Olkin (1967)) to address the joint termination alone as done by An, Christensen, and Gupta (2004) in the analysis of joint retirement in Denmark, but parameter estimates cannot be directly interpretable in terms of the couples' simultaneous decision process. The model also extends simultaneous duration models differently from Honoré and de Paula (2010): whereas that paper suggests a non-cooperative game theoretic framework, the use of a cooperative framework is much more appealing for the application to joint retirement that we address here. This brings in new features as well as new challenges. The framework presented in this paper directly corresponds to an economic model of decision-making by husband and wife, and it can consequently be more easily interpreted in light of such a model. To estimate the model, we resort to indirect inference (Smith (1993); Gourieroux, Monfort, and Renault (1993); and Gallant and Tauchen (1996)), following Honoré and de Paula (2014).

The remainder of this paper proceeds as follows. Section 3 describes our model and the empirical strategy for its estimation. In Section 2 we briefly describe the data and subsequently discuss our results in Section 4. We conclude in Section 5.

2 Preliminary Analysis of the Data

We break our analysis between the Survey on Health, Ageing and Retirement in Europe (SHARE) and the English Longitudinal Study of Ageing (ELSA).

2.1 Survey of Health, Ageing and Retirement in Europe

We base our analysis on four waves of the Survey of Health, Ageing and Retirement in Europe (SHARE) data (collected on 2004, 2006/7, 2008/9 and 2011/2). The survey focuses on individuals aged 50 or over and collected information on more than 60,000 people. In its first wave, data was collected on households from 11 countries: Austria, Germany, Sweden, Netherlands, Spain, Italy, France, Denmark, Greece, Switzerland and Belgium (plus Israel in 2005/6). Data from the Czech Republic, Poland and Ireland was also gathered in wave 2 in addition to the original eleven countries. In wave 3, data was collected in 13 countries (those from wave 2 minus Israel). finally in wave 4, 4 additional countries were added: Hungary, Portugal, Slovenia and Estonia. A fifth wave is currently being collected. To analyze retirement, we use the self-reported retirement date given by individuals.[4] Retirement is observed at least at an yearly frequency and, for many individuals, at a monthly frequency. In this document we focus either on yearly data or use imputed months for those who do not report their month of retirement. Throughout our analysis, we exclude households with multiple spouses and/or couples throughout the period of analysis, couples with conflicting information over marital status or other joint variables, and couples of the same gender.

To gain familiarity with the patterns of joint retirement, Table 1 documents wives' retirement behaviour given the retirement year of the husband among couples where both partners have retired by their last interview. In this initial analysis, we use only couples with uncensored

[4] We also generated other retirement variables, based on time at work and income source. Although those variables also indicate very similar patterns to the one we use here, we believe our measure of retirement is more suitable to our exercise.

retirement years for both spouses where both partners retire between 2005 and 2012.[5]

TABLE 1 HERE

It is apparent that a substantial proportion of individuals retire in the same year as their spouses. An interesting pattern arising from the tabulation is that while official retirement rates were constant in most countries in the sample, joint retirement seem more prevalent in 2005-2007 than in 2008-2012[6]. This can probably be ascribed to the deterioration of economic conditions in the second half of the period depicted in the table. Using data on *expected retirement dates* from ELSA, for instance, Crawford (2011) finds some evidence that individuals positively revised their expected retirement ages in direct response to asset value deterioration. Although that response is relatively imprecisely estimated, a general increase in retirement age is observed around the crisis period is detected though. As explained in the paper, this is can be ascribed to a tendency to upwardly revise retirement expectations as one ages (which is observed in previous waves of the survey), but also potentially an increase in risk aversion and/or uncertainty.

Breaking up the sample above by countries, we can also compare the relative frequencies of observations where both partners are not censored with the overall frequencies from each country. This is done in Table 2. As displayed, some countries with smaller representation and/or which started being surveyed only in later waves in our SHARE sample have censoring in the retirement date for at least one partner (i.e., Estonia, Hungary, Ireland, Portugal and Slovenia). The proportion of joint retirees per country is positive in all remaining countries (except for Italy), though the number of observations per country is then relatively small.

TABLE 2 HERE

In our subsequent statistical analysis, to avoid left-censoring, we select households that had both partners in the labor force when the oldest in the couple was around 50 years-old. Right-censoring occurs if the couple splits, when someone dies or has his or her last interview before the end of the survey. This leaves us with 4,083 couples. Of those, 131 couples have both the husband's

[5]Only 2 additional couples with both partners uncensored has a partner retiring before 2005. (In our subsequent analysis we use many more couples as we also incorporate censored observations.)

[6]When absolute frequencies are used, the 2005-2007 sub-matrix is clearly "dominant diagonal" whereas the 2008-2012 is clearly not.

and the wife's uncensored retirement dates.[7] We condition covariates on the first "household year": when the oldest partner reaches 50 years-old.[8] Because additional variables would reduce our sample substantially (even using imputed versions of those variables), we decided to focus our study on a relatively small subset of variables. The covariates we use are:

1. the age difference in the couple (husband's age minus wife's age in years);

2. dummies for country; and

3. self-reported health dummies (good health, very good health, and poor health);

Table 3 presents summary statistics for these variables. Note that we observe potential censoring months even for the observations that are uncensored in the data. As pointed out previously, this allows us to impose the same censoring process in the simulations as used to generate the data. It should be noted that the censoring rate is quite elevated relative to similar surveys such as the Health and Retirement Study. This is partly ascribed to SHARE being a "younger" survey, where many individuals are still in the labor force by their latest interview, and to its higher attrition rates.

TABLE 3 HERE

Figure 1 presents the Kaplan-Meier estimates for the retirement behavior in our sample (measured in months since the oldest partner turned 50 years old). We should note that, since age difference (husband's age minus wife's age) is positive, at on average 2.58 years (approximately, 30 months), the estimated distribution for women is shifted to the right with respect to their calendar age. As seen from the graphs, two dates stand out: 120 and 180 months. Those two dates correspond to ages 60 and 65, which are focal points for retirement decisions and coincide with the official retirement ages for the countries in our sample during the period under study.[9]

FIGURE 1 HERE

[7] The sample used in Tables 1 and 2 focuses instead on couples retiring between 2005 and 2012. This reduces the overall number of couples observed relative to the sample used in the analysis that follows. The number of couples where both spouses are uncensored is smaller in the subsequent analysis because there we restrict ourselves to couples where both partners were in the labor force when the oldest spouse was around 50-years old.

[8] We take the measurements from the first interview after the oldest spouse turns 50.

[9] See, for example, `http://www.oecd.org/els/emp/ageingandemploymentpolicies-statisticsonaverageeffectiveageofretirement.htm`.

As a preliminary assessment of how one's retirement behaviour is affected by his or her spouse's retirement timing, we have also estimated a simple discrete duration analysis of a couple's retirement choices. On Table 4, we display the results for a logit model of retirement at a given year as a function of dummies for the individual's age, whether his or her spouse retires in that particular year and whether his or her spouse has already retired. Columns labeled I display the probability for retirement at a particular age given that the spouse has not retired nor retires at the same year. Those labeled II provide the probabilities for retirement at a particular age given that the spouse retires at the same year. We notice that the effect of contemporaneous spousal retirement on a husband's retirement timing is not statistically significant, but does appear significantly for the wives, almost doubling the probability of retirement at certain ages. Columns III then show the probabilities of retirement at given ages conditional on the spouse having already retired. Previous retirement by the spouse is significant for both men and women.

TABLE 4 HERE

Finally, we estimate simple Weibull proportional hazards models for husbands and wives separately on Tables 5 using some of the covariates enumerated above. Because of the high fraction of censored observations, parameters are rather imprecisely estimated for many of the covariates though we reject the hypothesis that they are jointly equal to zero in any of the specifications. The parameter θ_1 is the Weibull parameter that controls duration dependence. Both male and female results display strong positive duration dependence: the probability of retirement increases with age. Age difference, defined as husband's age minus wife's age, affects positively the hazard rate of men and negatively the hazard rate of women. Since the oldest in the couple tends to be the husband, this partly captures the fact that larger age differences will imply that women are typically younger, lower their hazard of retirement. Improvements in health seem to affect both men and women positively at first and then decay for men though not for women.

TABLE 5 HERE

2.2 English Longitudinal Study of Ageing

Our analysis of ELSA parallels that of the SHARE data. ELSA is a large scale longitudinal panel study of people aged 50 and over, and their partners, living in private households in England. The same group of respondents have been interviewed at two-yearly interviews. The sample was drawn from households that had previously responded to the Health Survey for England (HSE). We use five different waves from ELSA. Wave 1 was conducted between 2002 and 2003 and the sample was selected from three years of HSE: 1998, 1999 and 2001. Wave 2 was conducted between 2004 and 2005 (there was no refreshment sample). Wave 3 was collected between 2006 and 2007 and a refreshment sample included new people from HSE 2001-2004, who were previously too young to join ELSA. Wave 4 was taken between 2008 and 2009 and then again a refreshment sample included new people from HSE 2006. Wave 5 was collected between 2010 and 2011 without a refreshment sample. Finally, a sixth wave was collected between 2012 and 2013 together with a refreshment sample of individuals aged between 50 and 55 from the HSE 2009, 2010 and 2011 surveys. This allows us to observe the behaviour of retirees at the onset of reforms to the State Pension Age for women (supposed to increase from 60 years old in April 2010 to 65 years old in 2020) and leverage our identification strategy.

As with our analysis of SHARE, we select households that had both partners in the labor force when the oldest in the couple was around 50 years-old. Right-censoring occurs if the couple splits, when someone dies or has his or her last interview before the end of the survey. This leaves us with 1,391 couples. Of those, 237 couples have both the husband's and the wife's uncensored retirement dates and 75 of those (32%) retire on the same year. We condition covariates on the first "household year" and use the same variables as in the SHARE analysis. The retirement timing is selected to the date of the last job for those who qualify themselves are retired (or semi-retired) (and, when that was not available, their self-reported retirement date). For ELSA, the timing variable used is observed at an yearly frequency (and expressed in months for congruence with the SHARE analysis).

Table 6 presents summary statistics for the variables we used. The censoring rate is also higher in comparison to similar surveys such as the Health and Retirement Study but lower than the one we found with SHARE.

TABLE 6 HERE

Figure 2 gives the Kaplan-Meier estimates for the retirement behavior in our sample (measured in months since the oldest partner turned 50 years old). As seen from the graph, one date (180 months) stands out in the estimates for men. This corresponds to the State Pension Age in

England. No such discontinuity is noticed in the female chart.

FIGURE 2 HERE

We have also estimated a simple discrete duration analysis of a couple's retirement choices using ELSA. On Table 7, we display the results for a logit model of retirement at a given year as a function of dummies for the individual's age, whether his or her spouse retires in that particular year and whether his or her spouse has already retired. The description of each one of the columns is as in the analogous table presented in our preliminary analysis of SHARE. We notice that the effect of contemporaneous spousal retirement is not only more statistically significant than in our presentation on SHARE, but also numerically more salient (for both genders). *The probability of retirement is often three to four times higher when the spouse retires in the same year. Previous retirement by the spouse is also significant for both men and women.*

TABLE 7 HERE

Finally, we also estimate simple Weibull proportional hazards models for husbands and wives separately on Table 8 using some of the covariates enumerated above. Because of the high fraction of censored observations, parameters are rather imprecisely estimated for many of the covariates though we reject the hypothesis that they are jointly equal to zero in any of the specifications. The parameter θ_1 is the Weibull parameter that controls duration dependence. Both male and female results display strong positive duration dependence: the probability of retirement increases with age. Age difference appears to be a much more tenuous determinant of the hazard rate for men than in the SHARE data. Here, improvements in health seem to affect both men and women positively at first and then decay.

TABLE 8 HERE

3 Model and Empirical Strategy

3.1 Basic Setup

In this section we describe a simple economic model that captures the features that spouses may decide jointly when to retire and that the optimal decision can be to retire at the same time. The framework is developed in Honoré and de Paula (2014) and we refer the reader to that paper for additional details. The model is explicitly designed to have the proportional hazard model as a special case. As is usual in choice models, the retirement decision depends only on the difference in utility between being retired or not. The levels of the utilities do not matter. This implies that many of the seemingly arbitrary assumptions made below are mere normalizations with no behavioral implications.

In order to simplify the notation, we measure time in terms of "household age". As mentioned earlier, in our empirical analysis family age is set to zero when the oldest partner in the couple reaches age 50. Throughout, we use $i = 1, 2$ to denote the two spouses in a married couple. n is used to index couples.

In our model, individual i with observable characteristics, $\mathbf{x}_i$, and whose spouse retires at time t_j, receives a utility flow of $K_i > 0$ before retirement. The vector (K_1, K_2) is the source of randomness in our econometric model. It is drawn from a joint distribution and its elements are potentially correlated due to, e.g., sorting or other commonalities. After retirement, the utility flow at time s is given by the deterministic function, $Z_i(s)\varphi_i(\mathbf{x}_i)\, D(s, t_j)$. The function $D(s, t_j)$ is defined as $(\delta - 1)1(s \geq t_j) + 1$ with $\delta \geq 1$ and it captures the idea that there can be complementarities in retirement. These complementarities can be either ascribed to taste or to institutional features such as Social Security rules that may promote coordination in retirement timing among husband and wife. (Whereas this parameter would not be invariant to changes in such regulations, it may be taken as fixed with respect to other counterfactuals.) In the calculations and exposition below, we restrict our attention to the case where δ is greater than or equal to 1. The parameter δ could in principle be less than one. However, this would not generate a positive probability that the individuals retire at the same time (as observed in the data). Given these, the discounted utility for individual i, who retires at t_i and whose spouse retires at t_j, is given by

$$U^i(t_i, t_j, \mathbf{x}_i, k_i) \equiv \int_0^{t_i} k_i e^{-\rho s} ds + \int_{t_i}^{\infty} Z_i(s)\varphi_i D(s, t_j) e^{-\rho s} ds.$$

We assume that the function $Z_i(\cdot)$ is increasing with $Z_i(0) = 0$. This simplifies the algebra, and it implies that, at the time the retirement decision is made, the couples expect retirement to be an absorbing state. As mentioned above, only the difference in utilities matters, so the assumption is that retirement becomes relatively more attractive over time. In particular, we are not assuming that some absolute measure of happiness increases with age. The multiplicative structure for $Z_i(s)\varphi_i(\mathbf{x}_i)\, D(s,t_j)$ is imposed because we want the resulting model to have the same structure as the familiar proportional hazard model. Except for that, it could easily be relaxed. In principle, it is possible to allow for kinks or discontinuities in $Z_i(\cdot)$. In a model without interdependence, those would correspond to discontinuities in the hazard rate in the case of kinks in $Z_i(\cdot)$ or, in the case of discontinuities in $Z_i(\cdot)$, positive probability of retirement at the discontinuity date to capture, for example, the positive probability of retirement at ages 60 and 65.

The δ could be made spouse-specific as well, but we focus on homogeneous δ and leave this degree of heterogeneity for future analysis. The reason for this is simplicity, and the fact that while the probability of joint retirement will be driven by δ, it is difficult to think of features of the data that would allow us to separately identify a δ for husbands and wives.

Given a realization (k_1, k_2) for the random vector (K_1, K_2), we assume that retirement timing is obtained as the solution to the Nash bargaining problem [Nash (1950); see also Zeuthen (1930)]:

$$\begin{aligned}\max_{t_1,t_2} \Bigg(\int_0^{t_1} k_1 e^{-\rho s} ds \quad &+ \int_{t_1}^{\infty} Z_1(s)\varphi_1 D(s,t_2) e^{-\rho s} ds - A_1 \Bigg) \times \\ &\Bigg(\int_0^{t_2} k_2 e^{-\rho s} ds + \int_{t_2}^{\infty} Z_2(s)\varphi_2 D(s,t_1) e^{-\rho s} ds - A_2 \Bigg)\end{aligned} \tag{1}$$

where A_1 and A_2 are the threat points for spouses 1 and 2, respectively. In the estimation, we set A_i equal to a fraction of the maximum utility individual i would obtain without the increased utility from the externality from the spouse's retirement. This specification of the threat points makes economic sense, and it also saves us from having to deal with the possibility that there are parameter values for which the factors in (1) cannot be made positive. In the general setting there may also be asymmetric bargaining weights that appear as exponents in the objective function.

The Nash bargaining solution concept is widely used in economics (see, for example, Chiappori, Donni, and Komunjer (2012)). It can be derived from a set of behavioral axioms on the bargaining outcomes (essentially Pareto efficiency, independence of irrelevant alternatives and symmetry) and it is widely adopted in the literature on intra-household bargaining. While it does not pin down a particular negotiation protocol between the parties involved, it can be motivated by the observation that it approximates the equilibrium outcome of a situation where husband and wife make offers to each other in an alternating order and the negotiation breaks down with a certain probability. As this probability goes to zero, the equilibrium converges to the Nash solution (see Binmore, Rubinstein, and Wolinsky (1986)).

As pointed out in Honoré and de Paula (2014), one alternative to the Nash bargaining framework used here would be a utilitarian aggregation of the utility functions in the household (i.e., the collective model of Chiappori (1992)). In that case, the retirement dates would solve:

$$\max_{t_1,t_2} cU^1(t_1,t_2;\mathbf{x}_1,K_1)+U^2(t_2,t_1;\mathbf{x}_2,K_2),$$

where c stands for the relative weight of agent 1's utility. This leads to the following first-order conditions:

$$c\times\frac{\partial U^1(t_1,t_2;\mathbf{x}_1,K_1))}{\partial t_i}+\frac{\partial U^2(t_2,t_1;\mathbf{x}_2,K_2)}{\partial t_i}=0,\quad i=1,2.$$

The setting we propose focuses instead on maximizing $(U^1(t_1,t_2;\mathbf{x}_1,K_1)-A_1)\times(U^2(t_2,t_1;\mathbf{x}_2,K_2)-A_2)$. This leads to the following first-order conditions:

$$\frac{U^2(t_2,t_1;\mathbf{x}_2,K_2)-A_2}{U^1(t_1,t_2;\mathbf{x}_1,K_1)-A_1}\times\frac{\partial U^1(t_1,t_2;\mathbf{x}_1,K_1))}{\partial t_i}+\frac{\partial U^2(t_2,t_1;\mathbf{x}_2,K_2)}{\partial t_i}=0,\quad i=1,2.$$

Consequently, the two are equivalent only if

$$c=\frac{U^2(t_2,t_1;\mathbf{x}_2,K_2)-A_2}{U^1(t_1,t_2;\mathbf{x}_1,K_1)-A_1}.$$

In this sense, the Nash bargaining setup imposes further restrictions on the model, as indicated in Chiappori, Donni, and Komunjer (2012). That paper also establishes identification results when a common set of covariates $\overline{\mathbf{x}}$ affects both the threat points $A_i, i=1,2$ and utilities $U^i, i=1,2$. Point-identification is achieved using spouse-specific covariates that affect the threat points $A_i, i=1,2$, but are excluded from $U^i, i=1,2$. In our empirical investigation we rely instead on spouse-specific covariates in $U^i, i=1,2$ and no excluded variables in the threat point functions $A_i, i=1,2$. Moreover, Chiappori, Donni, and Komunjer (2012) assume that latent variables (i.e., $K_i, i=1,2$) are additively separable, which is not our case.

As discussed in Honoré and de Paula (2014), the above model allows for positive probabilities of joint and sequential retirement when $\delta>1$. On the other hand, when $\delta=1$ the optimal retirement dates will correspond to

$$\log Z_i(t_i) = -\log \varphi_i + \log K_i, \quad i = 1, 2.$$

K_i following a unit exponential distribution gives a proportional hazard model. For a general distribution of K_i, this yields the generalized accelerated failure time model of Ridder (1990). This is the sense in which the approach discussed in this section can be thought of as a simultaneous equations version of a generalized accelerated failure time model.

3.2 Estimation: Indirect Inference

Because the likelihood for this model is not easily computed in closed form, we resort to simulation-assisted methods. One potential strategy would be to use simulated maximum likelihood (SML), where one non-parametrically estimates the conditional likelihood via kernel methods applied to simulations of T_1 and T_2 at particular parameter values and searches for the parameter value that maximizes the (simulated) likelihood. We opt for a different strategy for two main reasons. First, our likelihood displays some non-standard features. For example, there is a positive probability for the event $\{T_1 = T_2\}$. Second, consistency of the SML estimator requires a large number of simulations, which can be computationally expensive.

To estimate our model we therefore employ an indirect inference strategy (see Gourieroux, Monfort, and Renault (1993); Smith (1993); and Gallant and Tauchen (1996)). Rather than estimating the maximum likelihood estimator for the true model characterized by parameter θ, one estimates an approximate (*auxiliary*) model with parameter β. Let $n = 1, \ldots, N$ index a sample of households (couples). Then, under the usual regularity conditions,

$$\widehat{\beta} = \arg\max_b \sum_{n=1}^{N} \log \mathcal{L}_a\left(b; z_n\right) \xrightarrow{p} \arg\max_b E_{\theta_0}\left[\log \mathcal{L}_a\left(b; z_n\right)\right] \equiv \beta_0\left(\theta_0\right) \tag{2}$$

where $\mathcal{L}_a$ is a pseudo–likelihood function (parameterized by b) for the auxiliary model, z_n is the data for observation n, and the expectation E_{θ_0} is taken with respect to the true model. $\beta_0\left(\theta_0\right)$ is known as the pseudo–true value and the key is that it depends on the true parameters of the data–generation process (θ_0). The basic idea, then, is that if one knew the pseudo–true value as a function of θ_0, it could be used to solve the equation

$$\widehat{\beta} = \beta_0\left(\widehat{\theta}\right)$$

and obtain an estimator for θ_0. In our case, we do not know $\beta_0(\theta)$, but we can easily approximate this function using simulations. For a particular value of the parameters of the structural model, θ, we generate R draws

$$\{(z_{1r}(\theta), z_{2r}(\theta), \ldots, z_{Nr}(\theta))\}_{r=1}^{R}$$

from our structural model. In practice this is done by transforming uniform random variables. These are then kept fixed as one varies θ. The parameter, θ, enters through the transfromation of these uniform random variables. We can then estimate the function

$$\beta_0(\theta) \equiv \arg\max_b E_\theta\left[\log \mathcal{L}_a(b; z_n)\right]$$

by

$$\widetilde{\beta}_R(\theta) = \arg\max_b \frac{1}{R}\sum_{r=1}^{R}\frac{1}{N}\sum_{n=1}^{N}\left(\log \mathcal{L}_a(b; z_{nr}(\theta))\right).$$

This suggests finding $\widehat{\theta}$ such that the generated data set using $\widehat{\theta}$ gives the same estimate in the auxiliary model as we got in the real sample, $\widehat{\beta} = \widetilde{\beta}_R\left(\widehat{\theta}\right)$. When the dimensionality of β is greater than the dimension of θ, this is not possible, and one then estimates θ by a minimum distance approach that makes the difference between $\widehat{\beta}$ and $\widetilde{\beta}_R(\theta)$ as small as possible.

While this approach is conceptually straightforward, it requires one to estimate β for each potential value of θ. This can be computationally burdensome and we therefore adopt a slightly different version based on the first-order conditions from estimating the auxiliary model. The expression (2) implies that

$$\frac{1}{N}\sum_{n=1}^{N}\mathcal{S}_a\left(\widehat{\beta}; z_n\right) = 0,$$

and that $\widehat{\beta}$ converges to the solution to $E_\theta\left[\mathcal{S}_a(b; z_n)\right] = 0$, where S_a is the pseudo–score associated with $\mathcal{L}_a$. Of course, the solution $E_\theta\left[\mathcal{S}_a(b; z_n)\right] = 0$ is just $\beta_0(\theta_0)$ from equaton (2). As before, we estimate $E_\theta\left[\mathcal{S}_a\left(\cdot; z_n\right)\right]$ as a function of θ using

$$\frac{1}{R}\sum_{r=1}^{R}\frac{1}{N}\sum_{n=1}^{N}\mathcal{S}_a\left(\cdot; z_{nr}\left(\theta\right)\right)$$

and θ_0 is estimated by making it as close to zero as possible. Specifically, if $\dim\left(\mathcal{S}_a\right) > \dim\left(\beta\right)$, we minimize

$$\left(\frac{1}{R}\sum_{r=1}^{R}\frac{1}{N}\sum_{n=1}^{N}\mathcal{S}_a\left(\widehat{\beta}; z_{nr}\left(\theta\right)\right)\right)^{\top} W\left(\frac{1}{R}\sum_{r=1}^{R}\frac{1}{N}\sum_{n=1}^{N}\mathcal{S}_a\left(\widehat{\beta}; z_{nr}\left(\theta\right)\right)\right) \tag{3}$$

over θ. The weighting matrix W is a positive definite matrix performing the usual role in terms of estimator efficiency. The optimal W can be calculated using the actual data (before estimating θ) and the asymptotic properties follow from standard GMM arguments (see Gourieroux and Monfort (1996) for details) . This strategy is useful because we only estimate the auxiliary model once using the real data. After that, we evaluate its first-order condition for different values of θ.

The outcome variable in our empirical analysis is censored. To use simulation-based inference we must be able to simulate data that have been censored by the same process. In practice this means that we must either model the censoring process parametrically or observe the censoring times even for those observations that are uncensored in the data. As discussed below, our application falls into the second category.

3.2.1 Auxiliary Model

Our auxiliary model is composed of four reduced-form models that are chosen to capture the features of the data that are our main concern: the duration until retirement for each of the two spouses, the idea that some married couples choose to retire jointly and the possibility that unobserved shocks may be correlated. For the first two, we use a standard proportional hazard model for each spouse with a Weibull baseline hazard and the usual specification for the covariate function. For the third, we use an ordered logit model as suggested by our paper Honoré and de Paula (2010). Finally, we use the correlation of residuals from a regression of time-to-retirement for husband and wife on relevant covariates to capture any association in unobservables. We present the models in more detail below.

3.2.2 Weibull Proportional Hazard Model

For each spouse i, the hazard for retirement conditional on x is assumed to be $\lambda_i(t|\,x) = \alpha_i t^{\alpha_i - 1}\exp\left(x^\top \beta_i\right)$. The (log) density of retirement for spouse i conditional on x, $\log f_i(t|x)$, is then given by:

$$\log\left\{\lambda_i(t)\exp\left(x^\top\beta_i\right)\exp\left(-Z_i(t)\exp\left(x^\top\beta_i\right)\right)\right\} = \log\alpha_i + (\alpha_i - 1)\log t + x^\top\beta_i - t^{\alpha_i}\exp\left(x^\top\beta_i\right)$$

The (conditional) survivor function can be analogously obtained and is given by:

$$\log S_i(t|x) = \log\left\{\exp\left(-Z_i(t)\exp\left(x^\top\beta_i\right)\right)\right\} = -t^{\alpha_i}\exp\left(x^\top\beta_i\right)$$

Letting $c_{i,n} = 1$ if the observed retirement date for spouse i in household n is (right-)censored, and $= 0$ otherwise, we obtain the log-likelihood function:

$$\log\mathcal{L}_i = \sum_{n=1}^{N}(1 - c_{i,n})\left(\log\alpha_i + (\alpha_i - 1)\log(t_{i,n}) + x'_{i,n}\beta_i\right) - \sum_{n=1}^{N} t_{i,n}^{\alpha_i}\exp\left(x'_{i,n}\beta_i\right)$$

3.2.3 Ordered Logit Model Pseudo MLE

In the spirit of the estimation strategy suggested in Honoré and de Paula (2010), we also use an ordered logit model as an auxiliary model. Whereas the Weibull model will convey information on the timing of retirement, this second auxiliary model will provide information on the pervasiveness of joint retirement and help identify the taste interactions leading to this phenomenon (i.e., δ). Define

$$y_n = \begin{cases} 1, & \text{if } t_1 > t_2 + 1 \\ 2, & \text{if } |t_1 - t_2| \leq 1 \\ 3, & \text{if } t_2 > t_1 + 1 \end{cases}$$

Incorrectly assuming an ordered logit model for y_n yields

$$P(y_n = 1 \text{ or } y_n = 2) = \Lambda\left(x_n^\top\gamma_1\right) \quad \text{and} \quad P(y_n = 2) = \Lambda\left(x_n^\top\gamma_1 - \gamma_0\right)$$

where $\Lambda(\cdot)$ is the cummulative distribution function for the logistic distribution. This allows us to construct the following pseudo-likelihood function:

$$\mathcal{Q} = \sum_{y_n=0} \log\left(1-\Lambda\left(x_{0n}^{\top}\gamma\right)\right) + \sum_{y_n\neq 0} \log\left(\Lambda\left(x_{0n}^{\top}\gamma\right)\right) + \sum_{y_n\neq 2} \log\left(1-\Lambda\left(x_{1n}^{\top}\gamma\right)\right) + \sum_{y_n=2} \log\left(\Lambda\left(x_{1n}^{\top}\gamma\right)\right)$$

where

$$x_{0n} = \left(x_n^{\top} \vdots \mathbf{0}\right)^{\top} \qquad x_{1n} = \left(x_n^{\top} \vdots \mathbf{1}\right)^{\top} \qquad \gamma = \left(\gamma_1^{\top} \vdots -\gamma_0\right)^{\top}$$

The explanatory variables in the three parts of the auxiliary model need not be the same, and they need not coincide with the explanatory variables in the model to be estimated. In the empirical section below, the covariates in the Weibull auxiliary models are each spouses's own values of the explanatory variables in the model of interest. We use a constant only as an explanatory variable in the ordered logit model. This leaves the number of overidentifying restrictions constant across specifications.

3.2.4 Covariance in Failure Times

To allow for correlation in the unobservables variables K_1 and K_2, we use copula functions. To perform the estimation, we augment our auxiliary models with the covariance in failure times (including censored observations in both data and simulation moments). Specifically, we calculate the covariance between the residuals from a regression of (log) failure time on all covariates for husband and wife. An alternative is to use the residuals from regressions on spouse-specific variables and/or to define generalized residuals from a proportional hazard model estimated by maximum likelihood. The reason why we did not choose those approaches is that the asymptotic distribution for the covariance would then depend on nuisance parameters (i.e., the regression coefficients). This is not the case if we use the same set of covariates for husband and wife and estimate the model by OLS.

4 Results

We now present estimation results for our simultaneous duration model using monthly data on couples' retirement. The discount rate ρ is set to 0.004 per month (i.e., 5% per year) and the threat points are set at 0.6 times the utility level they would have obtained if his or her partner never retired.[10] The number of simulations in each set of estimates is $R = 10$. Here we assume that $Z_i(\cdot)$ is smooth: $Z_i(t) = t^{\theta_{1i}}$. As pointed out earlier, it is possible to allow for non-smoothness in $Z_i(\cdot)$ and we hope to incorporate this to account for positive probabilities of retirement at ages 60 and 65. In our baseline specifications, utility flows while in the labor force are drawn from independent unit exponentials, $K_i \sim \exp(1)$. To allow for positive correlation between the unobserved variables K_1 and K_2 (induced, e.g., by sorting), we use a Clayton-Cuzick copula function (see Clayton and Cuzick (1985)). More precisely, we model the joint cummulative distribution function of K_1 and K_2 as:

$$F_{K_1,K_2}(k_1, k_2; \tau) = K(1 - \exp(-k_1), 1 - \exp(-k_2); \tau),$$

where

$$K(u, v; \tau) = \begin{cases} \left(u^{-\tau} + v^{-\tau} - 1\right)^{-1/\tau} & \text{for} \quad \tau > 0 \\ uv & \text{for} \quad \tau = 0. \end{cases}$$

The unobservables are independent when $\tau = 0$. When $\tau > 0$, there is positive dependence between variables K_1 and K_2. Specifically, Kendall's rank correlation for the Clayton-Cuzick copula is equal to $\tau/(2 + \tau)$ (see, for example, Trivedi and Zimmer (2006)). This copula is commonly used to introduce dependence in the duration literature. Finally, we take $\varphi_i(x_i) = \exp(\theta_{2i}^{\top} x_i)$. This implies that when $\delta = 1$ and $\tau = 0$, the durations follow simple independent proportional hazard Weibull models (Lancaster (1990), p.44). This is the sense in which our approach generalizes simple standard econometric duration models. We also note that, even if $\delta = 1$, the copula used here plays the dual role of introducing correlation between the unobserved variables K_1 and K_2 and allowing for unobserved heterogeneity in the hazard rates to retirement.

In this paper, we only present results for a very simple specification of our model, using only age difference and self-reported health as additional controls (plus country dummies in the case of SHARE). Table 9 presents our estimates using the SHARE data. There is positive duration dependence: retirement is more likely as the household ages. Age differences tend to increase the retirement hazard for men and decrease it for women as in our independently estimated duration models. Since men are typically older and we count "family age" from the 50th year of the older partner, a larger age difference implies that the wife is younger at time zero and less likely to retire

[10]Our experience with the HRS nevertheless suggests that this choice will not matter for the estimated parameters of interest.

at any “family age” than an older woman (i.e., a similar wife in a household with a lower age difference).

We estimate versions of the model allowing for correlation in the unobservable component and not. As expected, the interaction parameter is higher when correlation in unobservables is precluded: any correlation in retirement timing rests more heavily on that parameter. Perhaps surprisingly, when correlation is allowed for, the interaction parameter is estimated at very close to one. In terms of our model, this means that the monthly utility flow of retirement increases very little (if at all) when one’s partner retires. We should nevertheless highlight that our inference on the effect of one spouse’s retirement on the other spouse retirement choise is severely complicated by the attrition levels in the data. We also note that the copula parameter is around 0.8, yielding a Kendall’s rank correlation coefficient of about 0.3. This is higher than our findings for the analysis with the HRS (see Honoré and de Paula (2014)). As explained previously, this correlation is potentially due to, e.g., sorting or other commonalities, and indicates that such phenomena are more prevalent in European data.

TABLE 9 HERE

The results for ELSA are much more marked. When no correlation between unobservables is imposed, the interaction parameter is high and estimated at 1.46 when only age difference is used and at 1.36 when self-reported health is allowed for. When we allow for correlation between the unobservables, the parameter is estimated at 1.03 and 1.01, which is slightly below our previous findings for the HRS. In comparison with both the HRS and SHARE, the correlation in unobservables, encompassing potential commonalities, is much stronger with an implied Kendall’s rank correlation of about 0.53. This parameter is nevertheless very noisily estimated.

TABLE 10 HERE

With ELSA we are able to examine the behaviour of retirees at the onset of reforms to the State Pension Age for women (supposed to increase from 60 years old in April 2010 to 65 years old in 2020) and leverage our identification strategy. To study the effect of the reform, we run the specifications above (without unobserved heterogeneity) controlling for whether the wife was exposed to the reform (i.e., born after 1950). The results are displayed on Table 11. As intuition would suggest, the reform affects directly the wife’s inclination to retire, lowering the hazard rate of retirement. The direct effect on husband's hazard rate is positive, but numerically small and statistically insignificant. Consequently, the reform's effect on husbands navigates mostly via its effects on women, delaying retirement. Interestingly, if one estimates a duration model with the wife's exposure to the reform, a statistically significant and positive coefficient is obtained for the direct effect of the reform on a husband's hazard rate of retirement.

TABLE 11 HERE

5 Concluding Remarks

The presents an analysis of joint retirement in Europe. Whereas high rates of attrition and the fewer number of waves lead to rather tenuous effets in the Survey of Health, Ageing and Retirement in Europe (SHARE), the longer running English Longitudinal Study of Ageing (ELSA) provides strong evidence of complementarities. Despite these shortcomings, both the reduced form and structural analysis yield evidence that retirement is best modeled as a joint economic decision. This is potentially important for estimating the effect of policy interventions. In particular, we are able to assess the effect of a reform to the State Pension Age for women (supposed to increase from 60 years old in April 2010 to 65 years old in 2020). The reform directly affect women's propensity to retire and, indirectly, the retirement of men.

References

An, M., B. Christensen, and N. Gupta (2004): "Multivariate Mixed Proportional Hazard Modeling of the Joint Retirement of Married Couples," *Journal of Applied Econometrics*, 19(6), 687–704.

Banks, J., R. Blundell, and M. Casanova Rivas (2007): "The Dynamics of Retirement Behavior in Couples: Reduced-Form Evidence from England and the U.S.," Working Paper, University College London (UCL).

Binmore, K., A. Rubinstein, and A. Wolinsky (1986): "The Nash Bargaining Solution in Economic Modelling," *Rand Journal of Economics*, 17, 176–188.

Blau, D. (1998): "Labor Force Dynamics of Older Married Couples," *Journal of Labor Economics*, 16(3), 595–629.

Blau, D., and D. Gilleskie (2004): "Health Insurance and Retirement of Married Couples," Working Paper, University of North Carolina at Chapel Hill.

Browning, M., P.-A. Chiappori, and V. Lechene (2006): "Collective and Unitary Models: A Clarification," *Review of Economics of the Household*, 4(1), 5–14.

Casanova Rivas, M. (2010): "Happy Together: A Structural Model of Couples' Joint Retirement Choices," Working Paper, University of California at Los Angeles (UCLA).

Chiappori, P.-A. (1992): "Collective Labor Supply and Welfare," *Journal of Political Economy*.

Chiappori, P.-A., O. Donni, and I. Komunjer (2012): "Learning from a Piece of Pie," *Review of Economic Studies*, 79(1), 162–195.

Clayton, D., and J. Cuzick (1985): "Multivariate Generalizations of the Proportional Hazards Model," *Journal of the Royal Statistical Society Series B*, 34, 187–220.

Coile, C. (2004a): "Health Shocks and Couples' Labor Supply Decisions," NBER Working Paper 10810.

——— (2004b): "Retirement Incentives and Couples' Retirement Decisions," *Topics in Economic Analysis and Policy*, 4(1).

Crawford, R. (2011): "The effect of the financial crisis on the retirement plans of older workers," MSc Dissertation, Department of Economics, University College London.

Gallant, R., and G. Tauchen (1996): "Which Moments to Match?," *Econometric Theory*, 12(4).

Gourieroux, C., and A. Monfort (1996): *Simulation-Based Econometric Methods*. Oxford: Oxford University Press.

Gourieroux, C., A. Monfort, and E. Renault (1993): "Indirect Inference," *Journal of Applied Econometrics*, 8, S85–S118.

Gustman, A., and T. Steinmeier (2000): "Retirement in Dual-Career Families: A Structural Model," *Journal of Labor Economics*, 18(3), 503–545.

——— (2004): "Social Security, Pensions and Retirement Behaviour with the Family," *Journal of Applied Econometrics*, 19, 723–737.

——— (2009): "Integrating Retirement Models," NBER Working Paper 15607.

Honoré, B., and A. de Paula (2010): "Interdependent Durations," *Review of Economic Studies*, 77, 11381163.

Honoré, B. E., and A. de Paula (2014): "Interdependent Durations in Joint Retirement," Working Paper, Princeton University and University College London (UCL).

Hurd, M. (1990): "The Joint Retirement Decisions of Husbands and Wives," in *Issues in the Economics of Aging*, ed. by D. A. Wise. Chicago: University of Chicago Press.

Lancaster, T. (1990): *Econometric Analysis of Transition Data*. Cambridge University Press.

Manski, C. (1993): "Identification of Endogenous Social Effects: The Reflection Problem," *Review of Economic Studies*, 60(3), 531–542.

Marshall, A., and I. Olkin (1967): "A Multivariate Exponential Model," *Journal of the American Statistical Association*, 62, 30–44.

Michaud, P., and F. Vermeulen (2011): "A Collective Labor Supply Model with Complementarities in Leisure: Identification and Estimation by Means of Panel Data," *Labour Economics*, 18, 159–167.

Michaud, P.-C. (2003): "Labor Force Participation Dynamics and Social Security Claiming Decisions of Elderly Couples," Working Paper, CentER Tilburn University and IZA.

Nash, J. (1950): "The Bargaining Problem," *Econometrica*, 18, 155–162.

Ridder, G. (1990): "The Non-Parametric Identification of Generalized Accelerated Failure-Time Models," *Review of Economic Studies*, 57.

Smith, A. (1993): "Estimating Nonlinear Time Series Models Using Vector-Autoregressions: Two Approaches," *Journal of Applied Econometrics*, 8, S63–S84.

Trivedi, P. K., and D. M. Zimmer (2006): "Copula Modeling: An Introduction for Practitioners," *Foundations and Trends in Econometrics*, 1(1), 1–111.

Zeuthen, F. (1930): *Problems of Monopoly and Economic Warfare.* London: Routledge.

Appendix: Figures and Tables

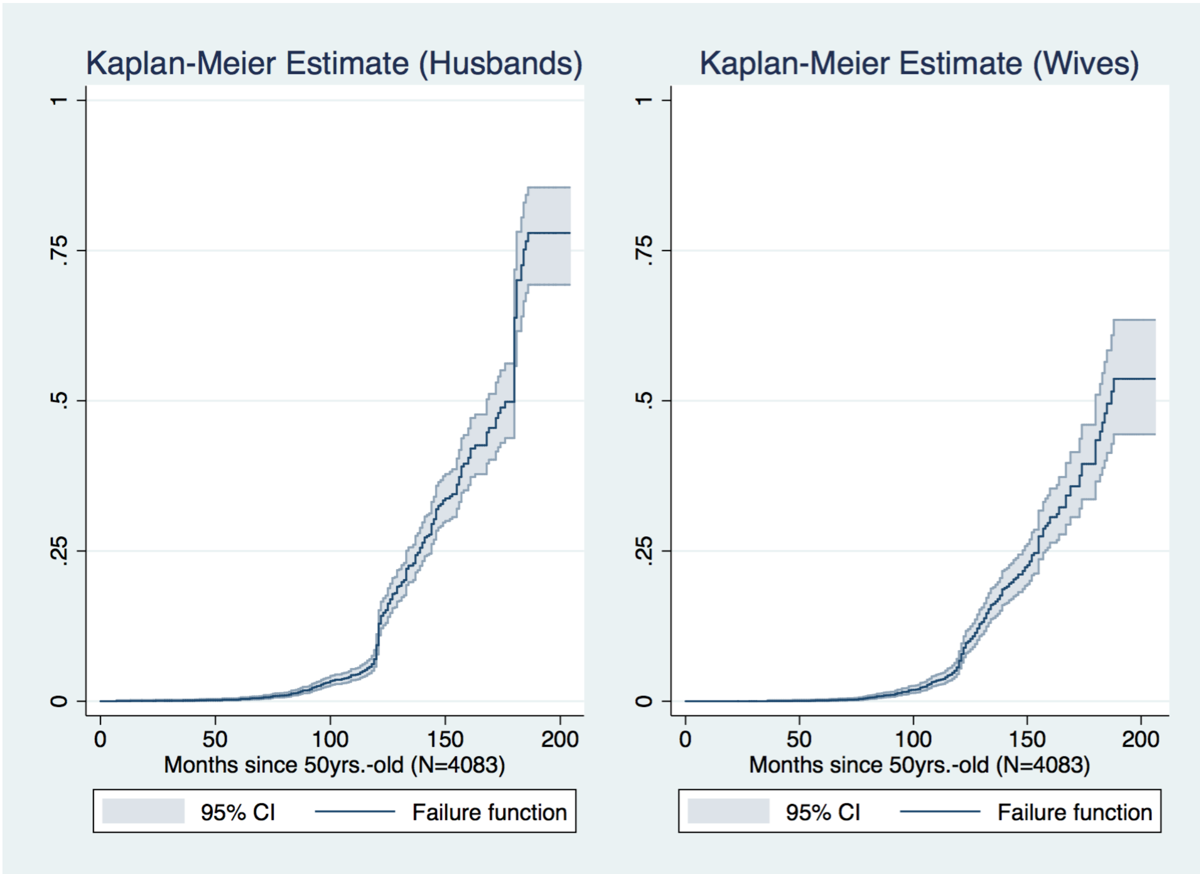

Figure 1: Kaplan-Meier Estimates (SHARE): Husband and Wife

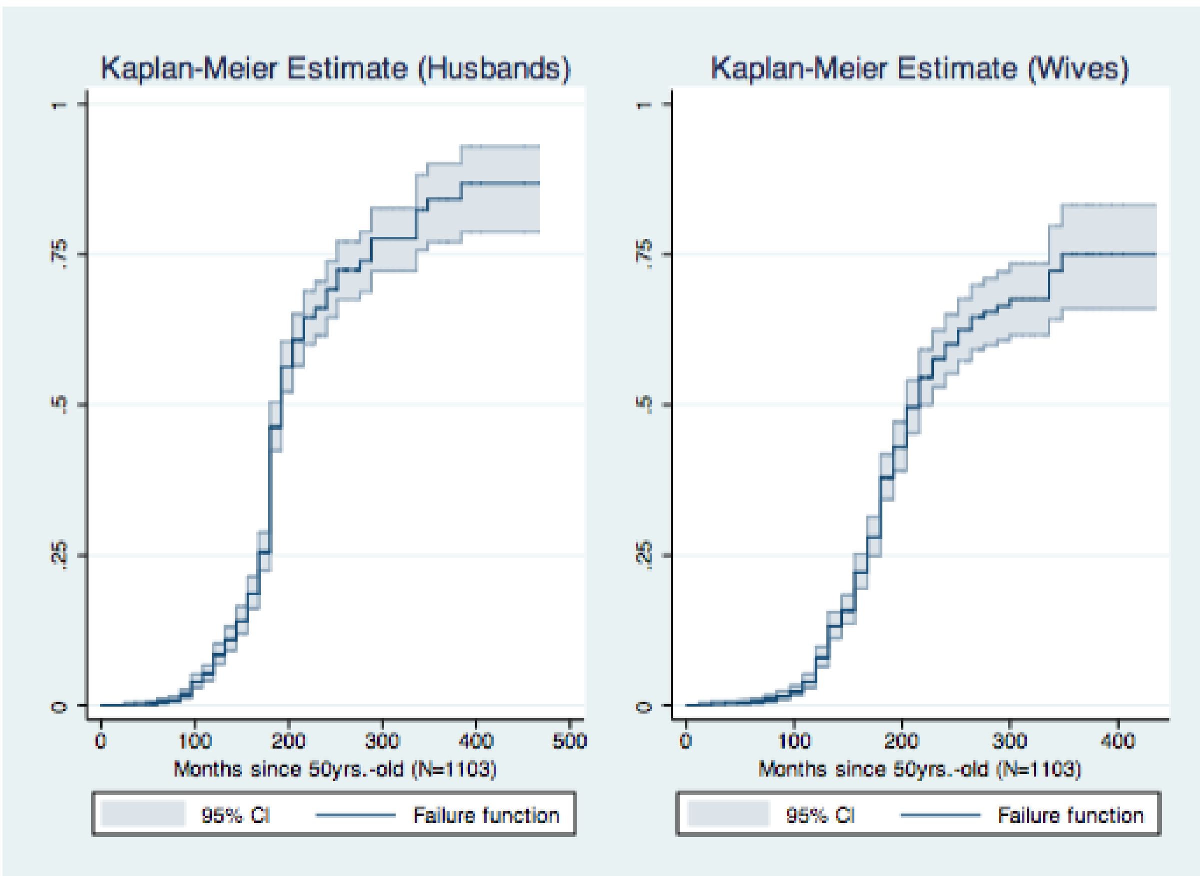

Figure 2: Kaplan-Meier Estimates (ELSA): Husband and Wife

Table 1: Conditional Retirement Year Frequency (in %) (SHARE)

	Wife							
Husbands	2005	2006	2007	2008	2009	2010	2011	Total
2005	**42.86**	14.29	0.00	7.14	21.43	7.14	7.14	100.00
2006	13.33	**36.67**	6.67	10.00	16.67	13.33	3.33	100.00
2007	8.33	0.00	**41.67**	8.33	20.83	20.83	0.00	100.00
2008	2.86	5.71	17.14	**22.86**	20.00	22.86	8.57	100.00
2009	0.00	11.76	8.82	17.65	**26.47**	32.35	2.94	100.00
2010	3.12	3.12	6.25	31.25	21.88	**21.88**	12.50	100.00
2011	6.06	12.12	18.18	6.06	15.15	24.24	**18.18**	100.00

Source: SHARE. The total number of couples with uncensored retirement years for both spouses is 202. There are many more observations where either husband or wife retire before 2005, but only 2 of those are uncensored for both spouses.

Table 2: Frequency by Country (SHARE)

	Overall (%)	Both Uncensored (%)
Austria	9.46	7.92
Belgium	16.94	13.86
Czech Republic	5.51	7.43
Denmark	3.64	5.45
Estonia	1.46	0.00
France	16.84	22.28
Germany	5.93	7.92
Greece	2.29	0.50
Hungary	1.77	0.00
Irland	0.10	0.00
Italy	7.90	6.44
Netherlands	9.25	11.39
Poland	2.39	5.45
Portugal	2.18	0.00
Slovenia	1.98	0.00
Spain	5.93	2.97
Sweden	4.05	5.94
Switzerland	2.39	2.48
N	962	202

Source: SHARE. These summary statistics pertain to the analysis documented on Table 1. The analysis in subsequent tables uses a different sample selection.

Table 3: Summary statistics (SHARE)

	All Observations		Uncensored		Censored	
Variable	**Mean**	**N**	**Mean**	**N**	**Mean**	**N**
Gender	0.50	8166	0.59	552	0.49	7614
Failure Month	90.50	8166	122.97	552	88.15	7614
Censored	0.93	8166	0	552	1	7614
Censoring Month[a]	92.41	8166	151.16	552	88.15	7614
Age Dif.	2.58	8166	1.85	552	2.63	7614
Excellent Health	0.15	7474	0.14	523	0.15	6951
Very Good Health	0.28	7474	0.24	523	0.29	6951
Good Health	0.38	7474	0.47	523	0.38	6951
Fair Health	0.16	7474	0.13	523	0.16	6951
Poor Health	0.03	7474	0.02	523	0.03	6951

a. For those uncensored, the censoring month is either the last interview or death date, which ever is the earlier date. It is used in the simulations for indirect inference.

Table 4: Probability of Retirement at Different Ages (SHARE)

	Husbands				**Wives**			
Variable	**Coef.**[a] (Std. Err.)	**I**[b] (in %)	**II**[c] (in %)	**III**[d] (in %)	**Coef.**[a] (Std. Err.)	**I**[b] (in %)	**II**[c] (in %)	**III**[d] (in %)
Spouse Retires at t	-0.09 (0.27)				0.57 † (0.30)			
Spouse Retired < t	1.75 ** (0.16)				1.19 ** (0.21)			
Age 54	0.25 (0.80)	0.06	0.05	0.32	2.03 ** (0.59)	0.27	0.48	0.89
Age 55	2.44 ** (0.45)	0.50	0.46	2.79	3.15 ** (0.51)	0.83	1.46	2.68
Age 56	2.50 ** (0.46)	0.53	0.49	2.98	3.08 ** (0.52)	0.77	1.36	2.49
Age 57	3.19 ** (0.43)	1.05	0.96	5.74	3.74 ** (0.50)	1.49	2.60	4.73
Age 58	3.18 ** (0.44)	1.04	0.96	5.70	3.98 ** (0.50)	1.88	3.28	5.92
Age 59	3.87 ** (0.42)	2.05	1.88	10.72	4.62 ** (0.49)	3.52	6.05	10.68
Age 60	5.43 ** (0.40)	9.12	8.42	36.49	5.87 ** (0.47)	11.28	18.34	29.42
Age 61	4.69 ** (0.42)	4.57	4.20	21.50	4.99 ** (0.52)	5.01	8.53	14.75
Age 62	4.95 ** (0.43)	5.83	5.37	26.17	5.59 ** (0.52)	8.78	14.54	24.00
Age 63	4.75 ** (0.48)	4.81	4.42	22.44	5.08 ** (0.63)	5.47	9.28	15.96
Age 64	4.62 ** (0.54)	4.28	3.94	20.38	3.88 ** (1.12)	1.71	2.98	5.39
Age $\geq$ 65	6.65 ** (0.48)	25.37	23.75	66.07	6.59 ** (0.71)	20.72	31.58	46.14
N of Individuals	4083				4083			

Significance levels : † : 10% $*$: 5% $**$: 1%.
a. Coefficient estimates are for a logit regression of retirement at year t on age dummies, whether the spouse retires that year and a dummy for whether the spouse is already retired. Observations correspond to an individual-year.
b. Probability of retirement at given age given that spouse has not yet retired.
c. Probability of retirement at given age given that spouse retires that year.
d. Probability of retirement at given age given that spouse has already retired.

Table 5: Weibull Duration Model (SHARE)

	HUSBANDS			WIVES		
Variable	**Coef.** (Std. Err.)	**Coef.** (Std. Err.)	**Coef.** (Std. Err.)	**Coef.** (Std. Err.)	**Coef.** (Std. Err.)	**Coef.** (Std. Err.)
θ_1	5.57 (0.19)	5.75 (0.21)	5.46 (0.22)	5.87 (0.24)	6.38 (0.27)	6.55 (0.29)
Age Dif./10	0.34 ** (0.13)	0.44 ** (0.13)	0.48 ** (0.14)	-1.35 ** (0.12)	-1.44 ** (0.12)	-1.45 ** (0.12)
Excellent or Very Good Health			-0.37 ** (0.13)			-0.10 (0.16)
Fair or Poor Health			-0.39 * (0.19)			0.18 (0.19)
Country Controls	NO	YES	YES	NO	YES	YES
Number of Obs.	4803	4083	3737	4803	4803	3737

Significance levels : † : 10% $*$: 5% $**$: 1%. θ_1 is the Weibull parameter.
Significance levels are not displayed for θ_1. Ommitted category is `Good Health`.

Table 6: Summary statistics (ELSA)

Variable	All Observations Mean	N	Uncensored Mean	N	Censored Mean	N
Gender	0.50	2782	0.54	777	0.49	2005
Failure Month	145.63	2782	161.93	777	139.31	2005
Censored	0.72	2782				
Censoring Month	156.22	2782	199.85	777	139.31	2005
Age Dif.	2.82	2778	2.35	777	3.01	2001
Excellent Health	0.19	2381	0.19	716	0.19	1665
Very Good Health	0.34	2381	0.32	716	0.35	1665
Good Health	0.34	2381	0.35	716	0.34	1665
Fair Health	0.10	2381	0.12	716	0.10	1665
Poor Health	0.03	2381	0.02	716	0.03	1665

a. For those uncensored, the censoring month is either the last interview or death date, which ever is the earlier date. It is used in the simulations for indirect inference.

Table 7: Probability of Retirement at Different Ages (ELSA)

	Husbands				Wives			
Variable	**Coef.**[a] (Std. Err.)	**I**[b] (in %)	**II**[c] (in %)	**III**[d] (in %)	**Coef.**[a] (Std. Err.)	**I**[b] (in %)	**II**[c] (in %)	**III**[d] (in %)
Spouse Retires at t	1.19 ** (0.21)				1.472 ** (0.215)			
Spouse Retired before t	1.01 ** (0.16)				0.736 ** (0.168)			
Age 54	0.29 (1.16)	0.07	0.23	0.20	-0.746 (1.055)	0.08	0.34	0.16
Age 55	2.25 ** (0.69)	0.51	1.65	1.39	1.568 ** (0.462)	0.78	3.29	1.60
Age 56	2.15 ** (0.71)	0.46	1.49	1.25	1.886 ** (0.443)	1.06	4.47	2.19
Age 57	2.88 ** (0.65)	0.95	3.06	2.57	1.664 ** (0.474)	0.85	3.61	1.76
Age 58	3.70 ** (0.61)	2.13	6.69	5.66	2.432 ** (0.416)	1.82	7.47	3.73
Age 59	3.01 ** (0.65)	1.08	3.46	2.91	2.667 ** (0.407)	2.29	9.27	4.67
Age 60	4.00 ** (0.61)	2.85	8.82	7.48	4.556 ** (0.355)	13.42	40.32	24.45
Age 61	3.62 ** (0.63)	1.97	6.21	5.25	3.501 ** (0.392)	5.12	19.05	10.13
Age 62	3.67 ** (0.63)	2.05	6.46	5.46	3.626 ** (0.394)	5.76	21.05	11.32
Age 63	4.34 ** (0.61)	3.96	11.95	10.20	3.746 ** (0.401)	6.45	23.11	12.59
Age 64	4.50 ** (0.61)	4.58	13.66	11.69	3.314 ** (0.448)	4.29	16.33	8.55
Age $\geq$ 65	5.34 ** (0.59)	10.07	26.95	23.59	3.489 ** (0.372)	5.06	18.86	10.02
N of Individuals	1391				1391			

Significance levels : † : 10% $*$: 5% $**$: 1%.
a. Coefficient estimates are for a logit regression of retirement at year t on age dummies, whether the spouse retires that year and a dummy for whether the spouse is already retired. Observations correspond to an individual-year.
b. Probability of retirement at given age given that spouse has not yet retired.
c. Probability of retirement at given age given that spouse retires that year.
d. Probability of retirement at given age given that spouse has already retired.

Table 8: Weibull Duration Model (ELSA)

	HUSBANDS		WIVES	
Variable	**Coef.** (Std. Err.)	**Coef.** (Std. Err.)	**Coef.** (Std. Err.)	**Coef.** (Std. Err.)
θ_1	3.24 (0.10)	3.30 (0.11)	3.07 (0.11)	3.14 (0.11)
Age Dif.	-0.05 (0.08)	-0.07 (0.08)	-0.50 ** (0.07)	-0.50 ** (0.08)
Excellent or Very Good Health		0.03 (0.11)		0.07 (0.12)
Fair or Poor Health		-0.06 (0.16)		0.01 (0.18)
Number of Obs.	1389	1166	1389	1215

Significance levels : † : 10% $*$: 5% $**$: 1%. θ_1 is the Weibull parameter.
Significance levels are not displayed for θ_1. Ommitted category is `Good Health`.

Table 9: Simultaneous Duration (SHARE)

Variable	Wife Coef. (Std. Err.)	Husb. Coef. (Std. Err.)	Wife Coef. (Std. Err.)	Husb. Coef. (Std. Err.)	Wife Coef. (Std. Err.)	Husb. Coef. (Std. Err.)	Wife Coef. (Std. Err.)	Husb. Coef. (Std. Err.)
δ	1.08		1.01		1.08		1.00	
	(0.18)		(0.29)		(0.51)		(0.02)	
θ_1	6.44	5.84	6.53	5.84	6.55	5.95	6.73	5.95
	(0.45)	(0.24)	(2.40)	(0.32)	(1.90)	(0.87)	(0.36)	(0.78)
Age Diff.	-1.58 **	0.31	-1.50 **	0.30	-1.76 **	0.34 †	-1.52 **	0.34
	(0.57)	(0.14)	(0.80)	(0.22)	(0.42)	(0.19)	(0.36)	(0.44)
$\geq$ V.G. Health					-0.16	-0.45 **	-0.13	-0.45 *
					(0.23)	(0.15)	(0.22)	(0.18)
$\leq$ Fair Health					0.21	-0.57 **	0.06	-0.48 *
					(0.27)	(0.21)	(0.33)	(0.22)
Country Controls	YES		YES		YES		YES	
τ			0.81				0.77	
			(2.58)				(1.00)	
N	4083		4083		3715		3715	
Function Value	6.09		1.68		7.42		2.02	

Significance levels :† : 10% $*$: 5% $**$: 1%. Significance levels are not displayed for θ_1 or δ. $\rho = 0.004$ and $R = 10$.

Table 10: Simultaneous Duration (ELSA)

Variable	Wife Coef. (Std. Err.)	Husb. Coef. (Std. Err.)	Wife Coef. (Std. Err.)	Husb. Coef. (Std. Err.)	Wife Coef. (Std. Err.)	Husb. Coef. (Std. Err.)	Wife Coef. (Std. Err.)	Husb. Coef. (Std. Err.)
δ	1.46		1.03		1.36		1.01	
	(0.12)		(0.32)		(0.18)		(0.13)	
θ_1	2.82	2.85	2.94	3.18	3.01	3.29	3.11	3.38
	(0.17)	(0.11)	(0.34)	(1.11)	(0.36)	(0.27)	(0.19)	(0.18)
Age Diff.	-0.74 **	0.16	-0.56 **	0.01	-0.56 **	0.12	-0.53	-0.08
	(0.26)	(0.16)	(0.26)	(0.20)	(0.42)	(0.19)	(0.38)	(0.23)
$\geq$ V.G. Health					0.16	0.05	0.12	0.11
					(0.24)	(0.20)	(0.21)	(0.20)
$\leq$ Fair Health					0.29	0.14	0.22	0.17
					(0.24)	(0.34)	(0.36)	(0.27)
τ			2.26				2.81	
			(5.13)				(2.56)	
N	1389		1389		1110		1110	
Function Value	1.91		0.01		3.72		0.11	

Significance levels :† : 10% $*$: 5% $**$: 1%. Significance levels are not displayed for θ_1 or δ. $\rho = 0.004$ and $R = 10$.

Table 11: Simultaneous Duration: Reform (ELSA)

Variable	Wife Coef. (Std. Err.)	Husb. Coef. (Std. Err.)	Wife Coef. (Std. Err.)	Husb. Coef. (Std. Err.)
δ	1.33 (0.13)		1.26 (0.22)	
θ_1	2.79 (0.19)	3.06 (0.21)	2.93 (0.27)	3.28 (0.18)
Reform	-0.71 ** (0.26)	0.05 (0.18)	-0.69 ** (0.31)	0.08 (0.22)
Age Diff.	-0.44 ** (0.36)	0.14 (0.19)	-0.31 (0.43)	0.13 (0.23)
$\geq$ V.G. Health			0.15 (0.18)	0.04 (0.19)
$\leq$ Fair Health			0.25 (0.24)	0.12 (0.29)
N	1389		1110	
Function Value	3.50		3.65	

Significance levels :† : 10% $*$: 5% $**$: 1%. Significance levels are not displayed for θ_1 or δ. $\rho = 0.004$ and $R = 10$.