cemmap
centre for microdata methods and practice

Interdependent durations in joint retirement

Bo E. Honoré
Áureo de Paula

The Institute for Fiscal Studies
Department of Economics, UCL

cemmap working paper CWP08/14

E·S·R·C ECONOMIC & SOCIAL RESEARCH COUNCIL
An ESRC Research Centre

Interdependent Durations in Joint Retirement*

Bo E. Honoré †
Princeton University

Áureo de Paula ‡
UCL, São Paulo School of Economics-FGV, C*e*MMAP, IFS

This Version: February 2014

*We thank Manuel Arellano, Richard Blundell, Stephane Bonhomme, Martin Browning, Ken Chay, Liran Einav, Eric French, Han Hong, Rafael Repullo, Geert Ridder, Jim Smith, Gerard van den Berg and Frank Wolak for useful conversations and David Card for suggesting this application. Financial support from a Steven H. Sandell grant from the Center for Retirement Research (Boston College) is gratefully acknowledged. 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 National Science Foundation, 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. Andrea Puig 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`

Abstract

This paper introduces a bivariate version of the generalized accelerated failure time model. It allows for simultaneity in the econometric sense that the two realized outcomes depend structurally on each other. The proposed model also has the feasure that it will generate equal durations with positive probability. The motivating example is retirement decisions by married couples. In that example it seems reasonable to allow for the possibility that the each partner's optimal retirement time depends on the retirement time of the spouse. Moreover, the data suggest that the wife and the husband retire at the same time for a non-negligible fraction of couples. Our approach takes as starting point a stylized economic model that leads to a univariate generalized accelerated failure time model. The covariates of that generalized accelerated failure time model act as utility-flow shifters in the economic model. We introduce simultaneity by allowing the utility flow in retirement to depend on the retirement status of the spouse. The econometric model is then completed by assuming that the observed outcome is the Nash bargaining solution in that simple economic model. The advantage of this approach is that it includes independent realizations from the generalized accelerated failure time model as a special case, and deviations from this special case can be given an economic interpretation. We illustrate the model by studying the joint retirement decisions in married couples using the Health and Retirement Study. We provide a discussion of relevant identifying variation and estimate our model using indirect inference. The main empirical finding is that the simultaneity seems economically important. In our preferred specification the indirect utility associated with being retired increases by approximately 5% if one's spouse is already retired and unobservables exhibit positive correlation. The estimated model also predicts that the indirect effect of a change in husbands' pension plan on wives' retirement dates is about 4% of the direct effect on the husbands.
JEL Codes: J26, C41, C3.

1 Introduction and Related Literature

This paper investigates the determinants of joint couples' retirement decisions. A majority of retirees are married and many studies indicate that a significant proportion of individuals retire within a year of their spouse. Articles documenting the joint retirement of couples (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). Even though this is especially the case for couples closer in age, a spike in the distribution of retirement time differences at zero typically exists for most couples, regardless of the age difference. This is illustrated in Figure 1.

The spike in the distribution of the difference in retirement dates for husbands and wives in Figure 1 suggests that many couples retire simultaneously. There are at least two distinct explanations for such a phenomenon. One is that the husband and wife receive correlated shocks (observable or not), driving them to retirement at similar times. The other is that retirement is jointly decided, reflecting the taste interactions of both members of the couple. In this paper, we primarily focus on the second of these explanations. Our motivation is that 55% of respondents in the Health and Retirement Study expected to retire at the same time as their spouses.[1]

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. In those literatures, the joint determination of a certain outcome of interest $y_i, i = 1, 2$ for two individuals $i = 1, 2$ is represented by the system of equations

$$\begin{aligned} y_1 &= \alpha_1 y_2 + x_1^\top \beta_1 + \epsilon_1 \\ y_2 &= \alpha_2 y_1 + x_2^\top \beta_2 + \epsilon_2 \end{aligned}$$

where x_i and $\epsilon_i, i = 1, 2$ represent observed and unobserved covariates determining y_i. We want to separate the endogenous (direct) effect (the α's) from the correlation in the ϵ's. There, as in this article, discerning these two sources of correlation in outcomes is relevant for analytical

[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.

and policy reasons. For example, when the estimated model does not allow for joint decision making within the couple, the estimate of the effect of a retirement-inducing policy shock can be misleading if the retirement times are indeed chosen jointly. The spillover effects resulting from joint decision making invalidate, for instance, the commonly employed Stable Unit Treatment Value Assumption (SUTVA) used in the treatment effects literature, preventing the clear separation of direct and indirect effects occurring through feedback to the partner's retirement decision [e.g., Burtless (1990)]. Furthermore, the multiplier effect induced by the endogenous direct effect of husband on wife or vice-versa is a potentially important conduit for policy. The quantification of its relative importance is therefore paramount for both methodological and substantive reasons.

Unfortunately, standard econometric duration models are not suitable for analyzing 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. An important contribution of this paper is therefore the specification of an econometric duration model that allows for simultaneity [see also de Paula (2009) and Honoré and de Paula (2010)]. 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.

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 20 years. 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,

[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).

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 Chiappori (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.

In our analysis, we assume that retirement decisions are made through Nash bargaining on the retirement date. This solution concept is attributed to Nash (1950) (see also Zeuthen (1930)). It chooses retirement decisions to maximize the product of differences between spouses utilities and respective outside options (i.e., "threat-points"). 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)].

Our model is a variation of a recently developed model (Honoré and de Paula (2010)) that extends well-known duration models to a (non-cooperative) strategic stopping game, in which endogenous and correlated effects can be disentangled and interpreted (see also de Paula (2009) for a related analysis). As such, it is close to traditional duration models in the statistics and econometrics literature. These are well understood and reasonably straightforward to estimate. Our 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

[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).

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 our model, we resort to indirect inference (Smith (1993); Gourieroux, Monfort, and Renault (1993); and Gallant and Tauchen (1996)), using as auxiliary models standard duration models and ordered discrete choice models, as suggested in Honoré and de Paula (2010) for a similar framework. (For an earlier application of indirect inference in a duration context, see Magnac, Robin, and Visser (1995)).

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

2 Model and Empirical Strategy

2.1 Basic Setup

In this section we formulate 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 model is explicitly designed to have the proportional hazard model as a special case. In this sense, it is a true generalization of a standard econometric model. 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 "family age." In our empirical analysis family age is set to zero when the oldest partner in the couple reaches age 60. Alternatively, we could have worked with the individuals' ages, but that would have been more cumbersome

because the economic motivation is that the person may enjoy utility from being retired at the same time as his or her spouse, and not from being retired at the same age. 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, j, 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 tax or 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. Given these, the discounted utility for individual i, who retires at t_i and whose spouse retires at t_j, is given by

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

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). In the calculations and exposition below, we therefore restrict our attention to the case where δ is greater than or equal to 1. We assume that the function $Z_i(\cdot)$ is strictly 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. Of course, this does not imply 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 this, the functional form for the utility flow 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.

The δ could be made spouse-specific as well, but we focus on homogeneous δ. 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 different δ for husbands and for wives.

This structure is essentially the same as in Honoré and de Paula (2010). There, it is assumed that the observed outcome, (T_1, T_2), is a Nash equilibrium. That assumption is in the spirit of much of the recent work in industrial organization, but it seems inappropriate in the context of retirement decisions. Given a realization (k_1, k_2) for the random vector (K_1, K_2), we therefore assume that retirement timing is obtained as the solution to the Nash bargaining problem (Nash (1950); see also Zeuthen (1930)):

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

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 a more general setting there may be asymmetric bargaining weights that appear as exponents in the objective function. Our analysis could be generalized to include that case, but we ignore this for simplicity and because it is difficult to think about nonparametric features of the data that would allow us to identify such an asymmetry.

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)).

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

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

As a result parameterising either the Nash bargaining approach will impose implicit constraints on c in the corresponding collective model. By the same token, parameterizing the collective model will impose constraints in the corresponding Nash bargaining model. See also 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.

In order to estimate a parameterized version of the Nash bargaining model, we will need to

solve it numerically many times. Note that the first term in (1) can be further simplified to

$$\left(k_1\rho^{-1}\left(1-e^{-\rho t_1}\right)+\varphi_1\widetilde{Z}_1\left(t_1\right)+\varphi_1\left(\delta-1\right)\widetilde{Z}_1\left(\max\left\{t_1,t_2\right\}\right)-A_1\right),$$

where $\widetilde{Z}_i\left(t\right)=\int_t^\infty Z_i(s)e^{-\rho s}ds$ and hence $\widetilde{Z}_i'\left(t\right)=-Z_i\left(t\right)e^{-\rho t}$. An analogous simplification applies to the second term. In other words, the objective function is given by

$$N\left(t_1,t_2\right)=\overbrace{\left(k_1\rho^{-1}\left(1-e^{-\rho t_1}\right)+\varphi_1\widetilde{Z}_1\left(t_1\right)+\varphi_1\left(\delta-1\right)\widetilde{Z}_1\left(\max\left\{t_1,t_2\right\}\right)-A_1\right)}^{\equiv I}\times$$
$$\underbrace{\left(k_2\rho^{-1}\left(1-e^{-\rho t_2}\right)+\varphi_2\widetilde{Z}_2\left(t_2\right)+\varphi_2\left(\delta-1\right)\widetilde{Z}_2\left(\max\left\{t_1,t_2\right\}\right)-A_2\right)}_{\equiv II}.$$

We parameterize $Z_i(t)$ as $Z\left(t;\alpha_i\right)=t^{\alpha_i}$. This will deliver the simple Weibull regression model as a special case of our model, and it yields a concenient expression for $\widetilde{Z}_i$:

$$\widetilde{Z}_i\left(t\right)=\int_t^\infty s^{\alpha_i}e^{-\rho s}ds=\left(\frac{1}{\rho}\right)^{\alpha_i+1}\Gamma\left(\alpha_i+1,\rho t\right),$$

where the upper incomplete gamma function is defined by $\Gamma\left(\alpha_i,x\right)=\int_x^\infty s^{\alpha_i-1}e^{-s}ds$.[4] As will be seen below, our model will have the proportional hazard model as a special case, and in that case $Z_i(t)$ will be the integrated baseline hazard. This is why the functional form assumption $Z_i(t)=Z\left(t;\alpha_i\right)=t^{\alpha_i}$ corresponds to a generalization of the proportional hazards model with a Weibull baseline hazard and why we use this is the parameterization below.

If the two spouses retire sequentially, say, $t_1<t_2$, the first-order condition with respect to t_1 is

$$\left(k_1e^{-\rho t_1}-Z_1(t_1)\varphi_1e^{-\rho t_1}\right)\left(\int_0^{t_2}k_2e^{-\rho s}ds+\int_{t_2}^\infty Z_2(s)\varphi_2\delta(s\geq t_1)e^{-\rho s}ds-A_2\right)=0.$$

[4]This expression can be further manipulated by noting that if the random variable X is Gamma distributed with parameters α and $\beta=1$

$$\overline{F}_{\Gamma(\alpha,1)}\left(x\right)=P\left(X>x\right)=\frac{1}{\Gamma\left(\alpha\right)}\int_x^\infty s^{\alpha-1}e^{-s}ds=\frac{\Gamma\left(\alpha,x\right)}{\Gamma\left(\alpha\right)}.$$

Consequently,

$$\widetilde{Z}\left(t;\alpha\right)=\left(\frac{1}{\rho}\right)^{\alpha+1}\Gamma\left(\alpha+1,\rho t\right)=\left(\frac{1}{\rho}\right)^{\alpha+1}\Gamma\left(\alpha+1\right)\overline{F}_{\Gamma(\alpha+1,1)}\left(\rho t\right)$$

which is useful since both $\Gamma\left(\cdot\right)$ and $\overline{F}_{\Gamma(\cdot,1)}\left(\cdot\right)$ are preprogrammed in many software packages.

This implies that either

$$k_1 = Z_1(t_1)\varphi_1$$

or

$$\int_0^{t_2} k_2 e^{-\rho s} ds + \int_{t_2}^{\infty} Z_2(s)\varphi_2 \delta(s \geq t_1) e^{-\rho s} ds = A_2.$$

The second possibility is ruled out since we specify the threat points so that each person gets a higher utility than his or her threat point at the Nash barganing solution. The first-order condition with respect to t_2 gives

$$Z_2(t_2)e^{-\rho t_2}\varphi_1 (1-\delta) \times (II) + (I) \times \left(k_2 e^{-\rho t_2} - Z_2(t_2)\varphi_2 \delta e^{-\rho t_2}\right) = 0. \tag{2}$$

The t_2 that solves this equation is smaller than the value obtained in Honoré and de Paula (2010): $Z_2^{-1}(k_2/\varphi_2\delta)$.[5] Intuitively, the reason is that with Nash bargaining, the second spouse to retire is willing to forgo some utility if the increase in utility to the other spouse is sufficiently high. Mathematically, we see this by noting that $Z_1(t_2) > 0$, $e^{-\rho t_2} > 0$, $\varphi_1 > 0$, and $1 - \delta < 0$. Moreover, II must be positive in equilibrium. This implies that $Z_1(t_2)e^{-\rho t_2}\varphi_1 (1-\delta) \times (II) \leq 0$ at the solution. So for the first-order condition to be zero, the product $(I) \times \left(k_2 e^{-\rho t_2} - Z_2(t_2)\varphi_2 \delta e^{-\rho t_2}\right)$ should be positive. Since I and $e^{-\rho t_2}$ are both positive, k_2 therefore must be greater than $Z_2(t_2)\varphi_2\delta$. Or equivalently, $t_2 < Z_2^{-1}(k_2/(\varphi_2\delta))$. This implies that

$$\begin{aligned} T_1 &= Z_1^{-1}\left(K_1/\varphi_1\right) \\ T_2 &\leq Z_2^{-1}\left(K_2/\left(\varphi_2\delta\right)\right), \end{aligned}$$

which gives the same timing choice for the first retiree as in Honoré and de Paula (2010) but an earlier one for the spouse. A similar set of calculations is obtained for $T_2 < T_1$.[6]

[5] In Honoré and de Paula (2010), we also require that $Z_1(\cdot) = Z_2(\cdot)$.

[6] For computation purposes we also notice that the objective function is unimodal on t_2. If we start at the critical value, increasing t_2 reduces the function. This is because, for small ρ, $Z_1(t_2)e^{-\rho t_2}\varphi_1 (1-\delta)$ becomes *more* negative and II becomes *more* positive, so the product becomes *more* negative. For the second term, I decreases and $k_2 e^{-\rho t_2} - Z_2(t_2)\varphi_2\delta e^{-\rho t_2}$, *which is positive*, decreases. Their product then decreases. Consequently, the derivative, which is the sum of these two products, becomes negative, and the objective function is decreasing. Analogously, we can also determine that the objective function is increasing for values below the critical value.

A third possibility is for the spouses to retire jointly. In this case,

$$\begin{aligned} T &= \arg\max_t N(t,t) \\ &= \arg\max_t \left(k_1\rho^{-1}\left(1-e^{-\rho t}\right)+\varphi_1\delta\widetilde{Z}_1(t)-A_1\right)\left(k_2\rho^{-1}\left(1-e^{-\rho t}\right)+\varphi_2\delta\widetilde{Z}_2(t)-A_2\right). \end{aligned}$$

The derivative of this with respect to t is

$$\begin{aligned} & e^{-\rho t}\left(K_1-\varphi_1\delta Z_1(t)\right)\left(k_2\rho^{-1}\left(1-e^{-\rho t}\right)+\varphi_2\delta\widetilde{Z}_2(t)-A_2\right) \\ & +e^{-\rho t}\left(k_1\rho^{-1}\left(1-e^{-\rho t}\right)+\varphi_1\delta\widetilde{Z}_1(t)-A_1\right)\left(k_2-\varphi_2\delta Z_2(t)\right). \end{aligned}$$

When $t < Z_1^{-1}(k_1/(\varphi_1\delta))$ and $t < Z_2^{-1}(k_2/(\varphi_2\delta))$ this derivate is positive, and when $t > Z_1^{-1}(k_1/(\varphi_1\delta))$ and $t > Z_2^{-1}(k_2/(\varphi_2\delta))$, it is negative. The optimum is therefore in the interval

$$\min\left\{Z_1^{-1}(k_1/(\varphi_1\delta)), Z_2^{-1}(k_2/(\varphi_2\delta))\right\} \leq t \leq \max\left\{Z_1^{-1}(k_1/(\varphi_1\delta)), Z_2^{-1}(k_2/(\varphi_2\delta))\right\}.$$

This is useful in the numerical solution to the optimal solution condition on joint retirement.

Figure 2 illustrates these cases and plots both T_1 and T_2 as a function of K_2 as K_1 is held fixed. For low values of K_2, $T_1 > T_2$: labor force attachment is higher for spouse 1 than for spouse 2. When K_1 is large, on the other hand, $T_1 < T_2$ and spouse 1 retires sooner. For intermediary values of K_1, $T_1 = T_2$ and the two spouses retire at the same time. This generates probability distributions such as those in Figure 3. Unconditionally, the probability density function for T_1 is smooth. Conditionally on $T_2 = t_2$, though, a point mass at $T_1 = t_2$ arises.

The set of realizations of (K_1, K_2) for which $T = T_1 = T_2$ is an optimum, is larger than the set obtained in the non-cooperative setup from Honoré and de Paula (2010). This is illustrated in Figure 4, where the area between the dotted lines is the joint retirement region in Honoré and de Paula (2010) and the area between solid lines is the joint retirement region in the current paper. Also, in that paper any date within a range $[\underline{T} < \overline{T}]$ (where $\underline{T} = \max\left\{Z_1^{-1}(k_1/(\varphi_1\delta)), Z_2^{-1}(k_2/(\varphi_2\delta))\right\}$) was sustained as an equilibrium for pairs (k_1, k_2) inducing joint retirement. In contrast, the equilibrium joint retirement date for a given realization of (K_1, K_2) is uniquely pinned down in the setup here. Because Nash bargaining implies Pareto

efficiency and because $\underline{T}$ is the Pareto dominant outcome among the possible multiple equilibria in the game analyzed by Honoré and de Paula (2010), it should be the case that joint retirement in the Nash bargaining model occurs on or before $\underline{T}$. In comparison to the non-cooperative paradigm adopted in our previous paper, Nash bargaining allows spouses to "negotiate" an earlier retirement date, which is advantageous to the household.

Finally, we note that 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.

2.2 Parametrization

As discussed above, we assume that $Z_i(t) = t^{\theta_{1i}}$. In order to include the proportional hazard model as the leading special case, we specify the marginal distributions of the utility flows while in the labor force as 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.

2.3 Estimation: Indirect Inference

Because the likelihood function for the model developed in the previous section 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,

$$\widehat{\beta} = \arg\max_{b} \sum_{n=1}^{N} \log \mathcal{L}_a\left(b; z_n\right) \xrightarrow{p} \arg\max_{b} E_{\theta_0}\left[\log \mathcal{L}_a\left(b; z_n\right)\right] \equiv \beta_0\left(\theta_0\right) \tag{3}$$

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 transformation 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 (3) 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 equation (3). As before, we estimate $E_\theta\left[\mathcal{S}_a(\cdot; z_n)\right]$ as a function of θ using

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

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

minimize

$$\left(\frac{1}{R}\sum_{r=1}^{R}\frac{1}{N}\sum_{n=1}^{N}\mathcal{S}_a\left(\widehat{\beta};z_{nr}\left(\theta\right)\right)\right)^{\top}W\left(\frac{1}{R}\sum_{r=1}^{R}\frac{1}{N}\sum_{n=1}^{N}\mathcal{S}_a\left(\widehat{\beta};z_{nr}\left(\theta\right)\right)\right) \tag{4}$$

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.

2.3.1 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 finally the idea that the retirement durations may exhibit correlation (conditionally on the covariates) even when they are not equal. 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). For the forth feature, we exploit the covariance in the residuals in regressions of the two retirement durations on all the covatiates of the model. We present the models in detail below.

2.3.2 Weibull Proportional Hazard Model

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

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

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

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

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:

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

First- and second-order derivatives used in the computation of the MLE for this auxiliary model are presented in the Appendix.

2.3.3 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 $\Lambda\left(\cdot\right)$ is the cummulative distribution function for the logistic distribution.

This allows us to construct the following pseudo-likelihood function:

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

where

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

As before, first- and second-order derivatives are presented in the Appendix.

The explanatory variables in the different 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.

In the data and in the simulations, y is defined using the failure time (i.e., the minimum between censoring and retirement dates). Censored observations do not pose problems when the other person in the household is uncensored and retires earlier, since in that case we can determine that retirement happened sequentially. Whereas we can always mark whether retirement was sequential or simultaneous in the simulations, when censoring happens before the retirement of the uncensored partner or both are censored, we cannot determine in the data whether retirement was sequential or not. Since we use the failure time in both data and simulations, censoring introduces the same degree of "noise" in the definition of y in the data and in the simulations.

2.3.4 Covariance in Failure Times

To allow for correlation in the unobservables variables K_1 and K_2, we use copula functions. We augment our auxiliary models with the covariance in failure times (including censored observations in both data and simulation moments) to perform the estimation. Specifically, we match the covariance between the residuals from a regression of (censored 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. Our procedure is therefore asymptotically equivalent to matching the true errors from those regressions (projections).

In the notation of a an objective function for and auxilyary model, we maximize

$$\mathcal{C}(\rho) = \sum_n (\hat{u}_{1,n}\hat{u}_{2,n} - \rho)^2$$

where $\hat{u}_{i,n} = \ln(t_{i,n}) - x_n^\top (\sum_h x_h x_h^\top)^{-1} (\sum_h x_h \ln(t_{i,h}))$ for $i = 1, 2$ with $t_{i,n}$ representing the failure time (earliest between retirement and censoring time) for partner i in couple n, x_n representing the covariates for couple n and $\hat{u}_{i,nr}$ is defined analogously on the simulated observations.

2.3.5 Overall Auxiliary Model

The overall auxiliary model objective function is then defined by the pseudo-log-likelihood function

$$\log \mathcal{L}_{men}(\alpha_1, \beta_1) + \log \mathcal{L}_{women}(\alpha_2, \beta_2) + \mathcal{Q}(\gamma) + \mathcal{C}(\rho)$$

and the moment conditions used for estimating the parameters of the structural model are the first-order conditions for maximizing this.

As indicated above, we choose as our weighting matrix $W = \widehat{J}_0^{-1}$, where

$$\widehat{J}_0 = \widehat{V} \left[\begin{pmatrix} \frac{\partial \log \mathcal{L}_{mn}}{\partial(\alpha_1, \beta_1)} \\ \frac{\partial \log \mathcal{L}_{wn}}{\partial(\alpha_2, \beta_2)} \\ \frac{\partial \mathcal{Q}_n}{\partial \gamma} \\ \frac{\partial \mathcal{C}_n}{\partial \rho} \end{pmatrix} \right]$$

The (asymptotic) standard errors of the structural estimates are calculated using the formulae in Gourieroux and Monfort (1996).

2.4 Computational Details

Evaluating the objective function (4) for a particular set of parameter values requires us to numerically solve the optimal retirement problem for each individual couple. As pointed out above, the structure makes this fairly straightforward. On the other hand, it is not easy to characterize the derivatives of the objective function with respect to the parameters. We therefore minimize (4) by repeating a multi-step procedure: First, we minimize using the method proposed by Powell (1964). Starting at the optimal parameter value, we then apply MATA's implementation of Nelder and

Mead (1965). If this sequence leads to an improvement of the objective function, we start over. Otherwise, we perturb each parameter by +/- 0.0001 and start over if that leads to an improvement. If it does not, we perform a full line search over each of the parameters. This process is repeated until the whole sequence does not lead to an improvement in the function value. In the estimation we specify $Z_i(t) = t^{\theta_{1i}}, i = 1, 2$. As explained above this simplifies the calculation of $\tilde{Z}_i(t)$.

To maximize the objective function, θ_{1i} is parameterized as $\exp(\alpha_i)$, and δ is parameterized as $1 + \exp(\tau)$. The starting values for the θ_{1i} and the coefficients of the explanatory variables are the maximum likelihood estimates from a Weibull proportional hazards model applied to the retirement times of husbands and wives, respectively, and the starting value for δ is 0.08. The discount factor is fixed and not estimated.

3 Data

We estimate the model using eight waves of the Health and Retirement Study (every two years from 1992 to 2006) and keep households where at least one individual was 60 years old or more. Retirement is observed at a monthly frequency. We use the retirement classification suggested by Rand Corporation, which classifies as retired a respondent who is not working and not looking for work and one for whom there is any mention of retirement through the employment status or the questions that ask the respondent whether he or she considers him- or herself to be retired.[7] To avoid left-censoring, selected households also had both partners in the labor force at the initial period. Right-censoring occurs when someone dies or has his or her last interview before the end of the survey. We excluded individuals who were part of the military. Finally, 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. This leaves us with 1,284 couples. Of those, 407 couples have both the husband's and the wife's uncensored retirement dates. Among the uncensored couples, 31 couples ($\approx$ 7.6%) retire jointly.[8] Figure 5 plots the retirement month of the husbands against the retirement month of the wives for those

[7] Specifically, we use the classification provided by the variable `RwLBRF`.

[8] There are 451 additional couples with only one censored spouse. If those are presumed to have retired sequentially, the proportion of joint retirements among couples with at most one censored spouse is 3.6%. Taking into account the remaining households where both individual retirement dates are censored would place the proportion of simultaneous retirements somewhere between 2.4% (if all additional households are assumed to retire sequentially) and 13.8% (if all additional households are assumed to retire simultaneously).

couples whose retirement month is uncensored for both spouses (January 1931 is month 1). The points along the 45-degree line are the joint retirements.

We condition covariates on the first "household year": when the oldest partner reaches 60 years old.[9] The covariates we use are:

1. the age difference in the couple (husband's age minus wife's age in years);
2. dummies for race (non-Hispanic black, Hispanic and other race with non-Hispanic whites as the omitted category);
3. dummies for education (high school or GED, some college and college or above with less than high school as the omitted category);
4. indicators of region (NE, SO, and WE with MW or other region as omitted category);
5. self-reported health dummies (good health, very good health, with poor health as the omitted category);
6. an indicator for whether the person has health insurance;
7. the total health expenditure per individual in the previous 12 months for the first two waves and the previous 2 years for the subsequent years[10] (inflation adjusted using the CPI to Jan/2000 dollars);
8. indicators for whether the person had a defined contribution (DC) or defined benefit (DB) plan; and
9. financial wealth (inflation adjusted using the CPI to Jan/2000 dollars).[11] This measure includes the value of checking and savings accounts, stocks, mutual funds, investment trusts, CDs, government bonds, Treasury bills and all other savings minus the value of debts such

[9] We take the measurements from the first interview after the oldest spouse turns 60.

[10] We use the transformation $\text{sgn}(\text{total health expenditure}) \times \sqrt{\text{abs}(\text{total health expenditure})}$. This transformation is in the spirit of a logarithmic transformation of positive variables and implies that large quantities have a decreasing effect. In contrast to a log transformation, it allows us to handle zeroes. In the computations, we also divide the transformed variable above by 10^2 to avoid overflow.

[11] For financial wealth we use the transformation $\text{sgn}(\text{financial wealth}) \times \sqrt{\text{abs}(\text{financial wealth})}$. This transformation is in the spirit of a logarithmic transformation of positive variables and implies that large quantities have a decreasing effect. In contrast to a log transformation, it allows us to handle negative numbers. It is concave for positive values and convex for negative ones. In the computations, we also divide the transformed variable above by 10^3 to avoid overflow.

as credit card balances, life insurance policy loans or loans from relatives. It does not include housing wealth or private pension holdings.

In Section 4, we examine whether our focus on covariates recorded at the initial household year might affect results substantially. We verify that couples that retire sequentially did *not* have more volatile realizations of health and financial variables than those couples that retire simultaneously. This might have been expected if shocks displaced couples' plans to simultaneous retirement.

Table 1 presents summary statistics for these variables. Note that we observe potential censoring months even for the observations that are uncensored in the data. This means that even though we assume that the censoring time is independent from retirement dates (conditional on the covariates), we do not need to model the distribution of censoring times to simulate the model. When drawing observations from the model to fit the auxiliary duration models, we are able to censor the simulated observations using the date when respondents were last interviewed or died as censoring dates even for those who retire earlier than that in the data (i.e., the "censoring month" in Table 1).

In Table 2, we present an overview of intra-household differences. Most of the couples marry within their own race but there is substantial variation in term of education. Many couples report different health statuses, and accordingly, there is a substantial difference in health expenditures. There are also differences with respect to insurance and pension ownership. Figure 6 presents the Kaplan-Meier estimates for the retirement behavior in our sample (measured in months since the oldest partner turned 60 years old).

4 Results

We now present our estimation results 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.[12] The number of simulations in each set of estimates is $R = 10$. We assume that $Z_i(\cdot)$ is smooth. As pointed out earlier, it is easy to allow for non-smoothness in $Z_i(\cdot)$. In fact, Figure 6, which displays estimates

[12] In our estimations, we experimented with other multiples of this scaled between 0 and 1 as well. See the discussion at the end of this section.

for marginal cummulative distribution functions of husbands and wives, suggests that a kink might be present at least for men around month twenty-four since the oldest household member turns 60 years old.[13] For simplicity, we abstract from this issue here.

As is often the case in structural estimation, it can be difficult to understand what features of the data identify the parameters of the model. Within the duration literature, this has led to a sizeable number of papers dealing with semiparametric identification. This literature is useful in shedding light on the relation between the parameters of the model and the underlying data. First, note that the functions $Z_i(\cdot)$, $\varphi_i(\cdot)$ and the marginal distribution for K_i are formally identified (up to scale) if covariates have a support large enough so that $\varphi_j(x_j)$ can be made arbitrarily close to zero. For such an individual, it is essentially optimal to have $t_j = \infty$. The other spouse will then optimally retire at T_i such that

$$\log Z_i(T_i) = -\log \varphi_i(x_i) + \log K_i$$

and one can apply the arguments in Ridder (1990) to identify $Z_i(\cdot)$, $\varphi_i(\cdot)$ and the marginal distribution of K_i (up to scale). We note also that this identification argument operates irrespective of the values of A_1 and A_2 (or asymmetries in the bargaining power). Intuitively, this argument would apply if the explanatory variables take values that make one of the spouses strongly attached to the labor force given his or her covariate values. In our data, for example, about 5% of the husbands who do not have a defined benefit pension plan retire after more than 126 months since the oldest member of the household turns 60 years old. Similarly, for the wives, 5% of those without a defined benefit pension plan retire more than 140 months since the oldest member turns 60.

Having identified $Z_i(\cdot), \varphi_i(\cdot)$ and the marginal distribution of K_i, the probability of joint retirement is driven by the interaction parameter δ. When $\delta = 1$ there are no retirement complementarities and joint retirement happens with zero probability. Larger values of δ will induce larger retirement complementarities, which should make joint retirement more likely. Even in the event of sequential retirement, whereas the first person to retire always retires at $Z_i^{-1}(K_i/\varphi_i)$, larger values of δ will lead to earlier retirement of the second retiree, providing additional variation to identify δ (see Appendix for details).

[13] Since the oldest member is usually the husband, this correponds to those turning 62 years old and becoming eligible to early retirement.

To understand why the joint distribution of K_1 and K_2 is identified, consider a point (k_1, k_2). If the joint support of covariates is large enough, then for that point, there is a pair (φ_1, φ_2) that induces sequential retirement in a neighborhood of (k_1, k_2). When there is sequential retirement, the retirement dates t_1 and t_2 are a one-to-one mapping from k_1 and k_2. For example, if $t_1 < t_2$, then t_1 is equal to $Z_1^{-1}(k_1/\varphi_1)$ and t_2 is also uniquely determined (see footnote 6). From the FOC, it is also clear that, given (t_1, t_2) (and $k_1 = Z_1(t_1)\varphi_1$), one can uniquely retrieve the corresponding k_2. Since we have a one-to-one mapping, the joint distribution of (T_1, T_2) is therefore informative about the joint density of (K_1, K_2). A different distribution of (K_1, K_2) in the neighborhood of (k_1, k_2) changes the probability of (T_1, T_2) given the covariates corresponding to the initial choice of (φ_1, φ_2).

Tables 3 and 4 present our estimates. The results are very robust across covariate specifications. 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. Since men are typically older and we count "family age" from the 60th year of the older partner, a larger age difference implies that the wife is younger at time zero and less likely to retire at any "family age" than an older woman (i.e., a similar wife in a household with a lower age difference). Both non-white men and women have a lower retirement hazard than non-Hispanic whites. The hazard of a Hispanic man is about 0.546 ($= \exp(-0.605)$) of a white man's, for example.

More educated women, but especially those with a high school diploma or GED and, in specifications that do not control for financial wealth, with a college degree, seem to retire earlier than those without a high school degree, but the coefficients on those categories are not statistically significant. For men, college-educated husbands and those with some college retire later than the others and the association is statistically significant (even when controlling for financial wealth). Husbands in the Northeast and West tend to retire earlier, whereas those in the South retire later than those in the Midwest. However, only statistically significant coefficients are those associated with the South (and only in some specifications). Geographical region does not seem to play a statistically significant role for women. Furthermore, depending on the covariate specification, women in the Northeast have a lower or higher hazard than those in the Midwest. Western and Southern wives do seem to retire earlier in all covariate specifications, but then again, standard errors are quite imprecise for most of the specifications.

Self-reported health lowers the hazard, with healthier people retiring later than those in poor health. Only the female coefficient on "good health" is significant in some of the specifications. Having health insurance increases the hazard for both husbands and wives, though not in a statistically significant way. Total health expenditures increase the hazard for both genders. Though health expenditures tend to increase as self-reported health decreases, we conjecture that this captures a deterioration in health not completely reflected in self-reported variables. Having a defined benefit pension plan increases the probability of retirement for both genders, but it is numerically and statistically much stronger for men. A defined contribution plan negatively affects the hazard for both, but here female effects are numerically and statistically more pronounced than those for men. Wealthier men and women tend to retire earlier and the effect is particularly more pronounced for women.

The interaction parameter ranges from 1.048 to 1.097 across our various specifications. In terms of our model, this means that the utility flow of retirement increases by more than 5% when one's partner retires. In terms of the effect on the hazard rate of retirement, this corresponds to about 15% of the effect of having a defined benefit plan for men. We also note that the copula parameter hovers around 0.5 in our numerous specifications, yielding a Kendall's rank correlation coefficient of about 0.2. As explained previously, this correlation is potentially due to, e.g., sorting or other commonalities.

To gauge the economic importance of the retirement externality, we simulated an experiment in which all men were assigned to a defined contribution pension plan and compared its outcome to assigning them to a defined benefit plan, holding everything else fixed.[14] This resulted in a 18.2-month change in the median uncensored retirement date for men. The simulated effect on the women was a change in the median uncensored retirement date of 0.75 months. In other words, the indirect effect on the women through the retirement externality is about 4% of the direct effect on the men. Given the large amount of censoring, one might argue that the median uncensored retirement date is not representative of the data we actually use. We therefore also compared the effect on the 25th percentile of uncensored retirement dates. Here, the direct effect on the husbands was 9.1 months, while the effect on the wives corresponded to about 4.3% of that.

We also added spousal variables as covariates to the last specification. Those variables were:

[14] We used 100 simulation draws per individual.

dummies for "very good health" and "good health" and dummies for defined benefit and defined contribution pensions. The coefficient estimates and p-values on all variables are essentially the same as in the tables for both husbands and wives. For males, none of the coefficients on spousal variables is statistically significant. For females, only the coefficient on a defined benefit pension plan for the spouse is statistically significant. The coefficient of the husband having a defined benefit plan on a woman's duration (0.451) is larger than that of the man himself having a defined benefit pension plan on his own duration to retirement, which is 0.315 once we include spousal covariates. (The effect of switching husbands from a defined contribution to a defined benefit plan is still in line with our previous results: median duration until retirement for men increases by 17.153 months whereas median time to retirement for wives increases by 0.787 months, corresponding to 4.3% of the direct effect on husbands duration.) In contrast, the point estimate of the effect of a wife having a defined benefit pension plan on the man's duration is also positive but virtually zero, and statistically insignificant. The absence of an effect of spousal health is in line with previous findings in the literature (e.g., Coile (2004a)).

All of the estimation results presented here can be thought of as GMM results. Rather than working with the same moment conditions for all specifications, we always work with the moment conditions that come from the scores of the Weibull proportional hazard auxiliary model using the same explanatory variables as in the final model, combined with the scores from the pseudo-likelihood function for the ordered logit using only a constant as an explanatory variable. This means that the number of overidentifying restrictions is one for all the specifications. As a specification test, we should therefore compare our minimized objective function to a Chi-squared distribution with one degree of freedom (see Gourieroux, Monfort, and Renault (1993)). The p-values associated with this test of overidentifying restrictions ranges from 20% to 64%. The average p-value across the six specifications is 36%. This suggests that our specification provides a good fit to the moments implicitly used in the estimation and have good predictive power on the retirement behavior of couple.

In Figure 7, we verify the robustness of our estimates to different threat point levels. As mentioned previously, we set $A_i, i = 1, 2$, equal to 0.6 of the utility spouse i would obtain if spouse j never retired. In the graph we plot 95% confidence intervals and point estimates of δ for various proportions of the utility one would get in case the partner were not to retire in the third

specification from Tables 3 and 4.[15] As seen from the point estimates hover around an average of 1.067, which is essentially the estimate presented in our main tables (i.e., 1.070).

Because the differential utility from joint retirement may depend on household characteristics, we also split our estimation into households where husband and wife are within 3 years apart in age and households where their age difference is greater than 3 years. The results are presented in Table 5 for the covariates used in the third specification from Tables 3 and 4. As expected, for households closer in age, Kendall's rank correlation coefficient ($\tau/(2+\tau)$) is higher (0.25 versus 0.21). Interestingly, the interaction coefficient δ appears to be higher for couples farther apart in age, but only slightly so and not in a statistically significant way. Retirement timing for wives in both types of households and for husbands in couples closer in age responds more to health conditions than in the baseline specification in Table 3. More educated wives now seem to systematically retire later than less educated ones in couples farther apart in age. At the same time, the delay in retirement is exacerbated for husbands in those households in comparison to our baseline results whereas it is attenuated for husbands in couples closer in age. The hazard is also comparatively higher for non-Hispanic black wives in couples close in age than in the baseline specification. For households farther apart in age, Hispanic wives retire later than in the baseline specification and college seems to decrease the retirement hazard while it raises it in the baseline specification. The coefficient on `Other Race` is quite different across subsamples. The proportion of individuals of other race is nonetheless small in both: around 3.2% among couples more than 3 years apart in age and only 1.8% among those closer in age.

Finally, to evaluate whether joint retirement is likely to be an outcome from a common shock, as opposed to the interaction between husband and wife, we compare the time variation of our regressors across couples that retire simultaneously and couples that retire sequentially. For the survey waves preceding retirement of any member in household, we look at the average proportional changes in financial assets and health expenditures and average changes in self-reported health status, pension plans (defined benefit and defined contribution) and health insurance. For all of these variables, couples retiring simultaneously displayed at least as much stability (if not more) in the survey waves preceding retirement than those retiring sequentially. For example, financial assets for those who end up retiring simultaneously are much more stable than for couples who

[15] The confidence intervals are asymmetric since lower and upper bounds correspond to symmetric lower and upper bounds on $\ln(\delta - 1)$ (which is the parameterization used in our computations for the estimation of δ).

retire sequentially: the average relative change in financial wealth across survey waves preceding retirement is a factor of 4.847 for those who retire simultaneously versus a factor of 10.850 for those who retire sequentially. Standard deviations were also lower for those couples retiring simultaneously. The same pattern arises even when the factor is deflated by the growth in the S&P500 stock market index. Furthermore, there is no discernible statistical difference between the average change in financial assets from survey wave to survey wave for these two groups. Consequently, it is unlikely that shocks to financial wealth (and, for that matter, that shocks to any of the variables listed above) explain the joint retirement decision in our sample.

5 Concluding Remarks

We have presented a new model that nests the usual generalized accelerated failure time model, but accounts for joint termination of spells and is built upon an economic model of joint decision making. We then applied the model to retirement of husband and wife using data from the Health and Retirement Study. The main empirical finding is that simultaneity seems economically important. In our preferred specification, the indirect utility associated with being retired increases by approximately 5% if one's spouse is already retired. By comparison, a defined benefit pension plan increases indirect utility by 36%. The estimated model also predicts that the indirect effect of a change in the pension plan for husbands on wives' retirement dates is about 4% of the direct effect on the husbands.

The present analysis can be extended in many respects. In future work, we plan to reproduce our analysis for a multi-country comparison of retirement behavior within couples (using, for example, similar data from the Survey of Health, Ageing and Retirement in Europe (SHARE) and the English Longitudinal Study of Ageing (ELSA)).

Appendix for Referees

AUXILIARY MODELS FOR INDIRECT INFERENCE

Log-likelihood Derivatives: Weibull Model

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

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

$$\frac{\partial^2 \log \mathcal{L}_i}{\partial \alpha_i^2} = -\sum_{n=1}^{N}(1-c_{i,n})\frac{1}{\alpha_i^2} - \sum_{n=1}^{N} t_{i,n}^{\alpha_i} \log\left(t_{i,n}\right)^2 \exp\left(x_{i,n}'\beta_i\right)$$

$$\frac{\partial^2 \log \mathcal{L}_i}{\partial \alpha_i \partial \beta_i'} = -\sum_{n=1}^{N} t_{i,n}^{\alpha_i} \log\left(t_{i,n}\right) \exp\left(x_{i,n}'\beta_i\right) x_{i,n}$$

$$\frac{\partial^2 \log \mathcal{L}_i}{\partial \beta_i \partial \beta_i'} = -\sum_{n=1}^{N} t_{i,n}^{\alpha_i} \exp\left(x_{i,n}'\beta_i\right) x_{i,n} x_{i,n}'$$

To impose $\alpha_i > 0$ in our computations we parameterize $\alpha_i = \exp(\theta)$. Then,

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

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

$$\begin{aligned}
\frac{\partial^2 \log \mathcal{L}_i}{\partial \theta^2} &= \frac{\partial}{\partial \theta}\left(\frac{\partial \log \mathcal{L}_i}{\partial \alpha_i}\frac{\partial \alpha_i}{\partial \theta}\right) \\
&= \frac{\partial^2 \log \mathcal{L}_i}{\partial \alpha_i^2}\left(\frac{\partial \alpha_i}{\partial \theta}\right)^2 + \frac{\partial \log \mathcal{L}}{\partial \alpha_i}\frac{\partial^2 \alpha_i}{\partial \theta^2} \\
&= \left(-\sum_{n=1}^{N}(1-c_{i,n})\frac{1}{\alpha_i^2} - \sum_{n=1}^{N} t_{i,n}^{\alpha_i} \log\left(t_{i,n}\right)^2 \exp\left(x_{i,n}'\beta_i\right)\right)\alpha_i^2 \\
&\quad - \left(\sum_{n=1}^{N}(1-c_{i,n})\left(\frac{1}{\alpha_i} + \log\left(t_{i,n}\right)\right) - \sum_{n=1}^{N} t_{i,n}^{\alpha_i} \log\left(t_{i,n}\right) \exp\left(x_{i,n}'\beta_i\right)\right)\alpha_i \\
\frac{\partial^2 \log \mathcal{L}_i}{\partial \theta \partial \beta_i'} &= \frac{\partial^2 \log \mathcal{L}_i}{\partial \alpha_i \partial \beta_i'}\frac{\partial \alpha_i}{\partial \theta} = \left(-\sum_{n=1}^{N} t_{i,n}^{\alpha_i} \log\left(t_{i,n}\right) \exp\left(x_{i,n}'\beta_i\right) x_{i,n}\right)\alpha_i \\
\frac{\partial^2 \log \mathcal{L}_i}{\partial \beta_i \partial \beta_i'} &= -\sum_{n=1}^{N} t_{i,n}^{\alpha_i} \exp\left(x_{i,n}'\beta_i\right) x_{i,n} x_{i,n}'
\end{aligned}$$

Pseudo-likelihood Derivatives: Ordered Model

$$\frac{\partial \mathcal{Q}}{\partial \gamma} = \sum_n \left[\left(1\{y_n \neq 0\} - \Lambda\left(x_{0n}^\top \gamma\right)\right) x_{0n} + \left(1\{y_n = 2\} - \Lambda\left(x_{1n}^\top \gamma\right)\right) x_{1n}\right]$$

$$\frac{\partial^2 \mathcal{Q}}{\partial\gamma\partial\gamma^\top} = -\sum_n \left[\left(\left(1-\Lambda\left(x_{0n}^\top\gamma\right)\right)\Lambda\left(x_{0n}^\top\gamma\right)\right)x_{0n}x_{0n}^\top + \left(\left(1-\Lambda\left(x_{1n}^\top\gamma\right)\right)\Lambda\left(x_{1n}^\top\gamma\right)\right)x_{1n}x_{1n}^\top\right]$$

ADDITIONAL DERIVATIONS FOR IDENTIFICATION DISCUSSION

Here we provide details for the effect of δ on the retirement date of the second spouse to retire when there is sequential retirement. First, note that when $t_1 \approx 0$, applying the Implicit Function Theorem to the FOC for t_2 (see equation (2)) gives

$$\frac{dt_2}{d\delta} = -\left[\frac{\frac{\partial^2 I}{\partial t_2 \partial\delta}\times(II) + \frac{\partial II}{\partial\delta}\times\frac{\partial I}{\partial t_2} + \frac{\partial^2 II}{\partial t_2\partial\delta}\times(I) + \frac{\partial I}{\partial\delta}\times\frac{\partial II}{\partial t_2}}{\frac{\partial^2 I}{\partial^2 t_2}\times(II) + \frac{\partial II}{\partial t_2}\times\frac{\partial I}{\partial t_2} + \frac{\partial^2 II}{\partial^2 t_2}\times(I) + \frac{\partial I}{\partial t_2}\times\frac{\partial II}{\partial t_2}}\right], \tag{5}$$

where (I) and (II) are defined as in equation (2). The various terms can be signed as shown below:

$$\begin{aligned}
&\tfrac{\partial I}{\partial\delta} = \varphi_1\tilde{Z}_1(t_2) > 0 && \tfrac{\partial II}{\partial\delta} = \varphi_2\tilde{Z}_2(t_2) > 0\\
&\tfrac{\partial I}{\partial t_2} = Z_1(t_2)e^{-\rho t_2}\varphi_1(1-\delta) < 0 && \tfrac{\partial II}{\partial t_2} = k_2e^{-\rho t_2} - Z_2(t_2)\varphi_2\delta e^{-\rho t_2} > 0\\
&\tfrac{\partial^2 I}{\partial t_2\partial\delta} = -Z_1(t_2)e^{-\rho t_2}\varphi_1 < 0 && \tfrac{\partial^2 II}{\partial t_2\partial\delta} = k_2e^{-\rho t_2} - Z_2(t_2)\varphi_2\delta e^{-\rho t_2} > 0\\
&\tfrac{\partial^2 I}{\partial t_2^2} = Z_1'(t_2)e^{-\rho t_2}\varphi_1(1-\delta) < 0 && \tfrac{\partial^2 II}{\partial t_2^2} = -\rho e^{-\rho t_2}(k_2 - Z_2(t_2)\varphi_2\delta) - Z_2'(t_2)e^{-\rho t_2} < 0.
\end{aligned}$$

These and the fact that $(I) \geq 0$ and $(II) \geq 0$ imply that the denominator in expression (5) is *strictly* negative. To see that the numerator is also negative notice that

$$\lim_{\delta\to 1}\left[\frac{\partial II}{\partial\delta}\times\frac{\partial I}{\partial t_2} + \frac{\partial I}{\partial\delta}\times\frac{\partial II}{\partial t_2}\right] = \varphi_1\tilde{Z}_1(t_2)[k_2 - Z_2(t_2)\varphi_2] = 0$$

where the last equality follows because $k_2 = Z_2(t_2)\varphi_2$ at the optimally chosen t_2 when $\delta = 1$. Since

$$\frac{\partial}{\partial\delta}\left[\frac{\partial II}{\partial\delta}\times\frac{\partial I}{\partial t_2} + \frac{\partial I}{\partial\delta}\times\frac{\partial II}{\partial t_2}\right] = -\varphi_1\varphi_2\left(Z_1(t_2)\tilde{Z}_2(t_2) + Z_2(t_2)\tilde{Z}_1(t_2)\right)e^{-\rho t_2} < 0,$$

it follows that

$$\frac{\partial II}{\partial\delta}\times\frac{\partial I}{\partial t_2} + \frac{\partial I}{\partial\delta}\times\frac{\partial II}{\partial t_2} < 0.$$

The other two remaining terms in the numerator are negative, which then implies that the numerator is negative. Consequently, (5) is negative: larger values of δ lead to earlier retirement by the second agent (i.e., lower t_2). Because $t_1 \approx 0$, I (and, consequently, t_2) will not depend on k_1. Having identified $Z_i(\cdot)$, $\varphi_i(\cdot)$ and the marginal distribution of K_2, this allows one to identify δ.

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.

Burtless, G. (1990): "Comment on M. Hurd's '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.

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).

de Paula, A. (2009): "Inference in a Synchronization Game with Social Interactions," *Journal of Econometrics*, 148(1), 56–71.

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.

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.

Magnac, T., J.-M. Robin, and M. Visser (1995): "Analysing Incomplete Individual Employment Histories Using Indirect Inference," *Journal of Applied Econometrics*, 10, S153–S169.

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.

Nelder, J. A., and R. Mead (1965): "A Simplex Method for Function Minimization," *The Computer Journal*, 7(4), 308–313.

Powell, M. J. D. (1964): "An Efficient Method for Finding the Minimum of a Function of Several Variables without Calculating Derivatives," *The Computer Journal*, 7(2), 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

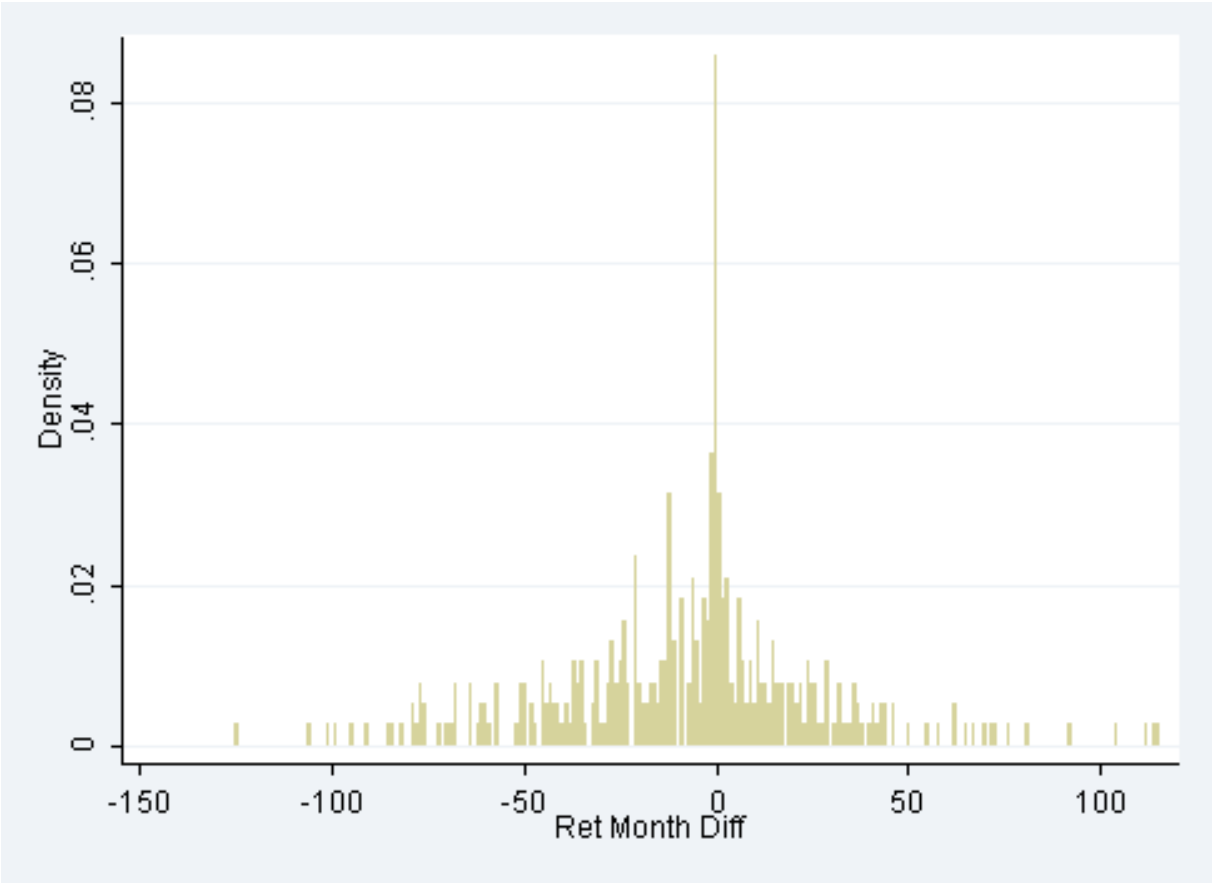

Figure 1: Difference in Retirement Months (Husband-Wife)

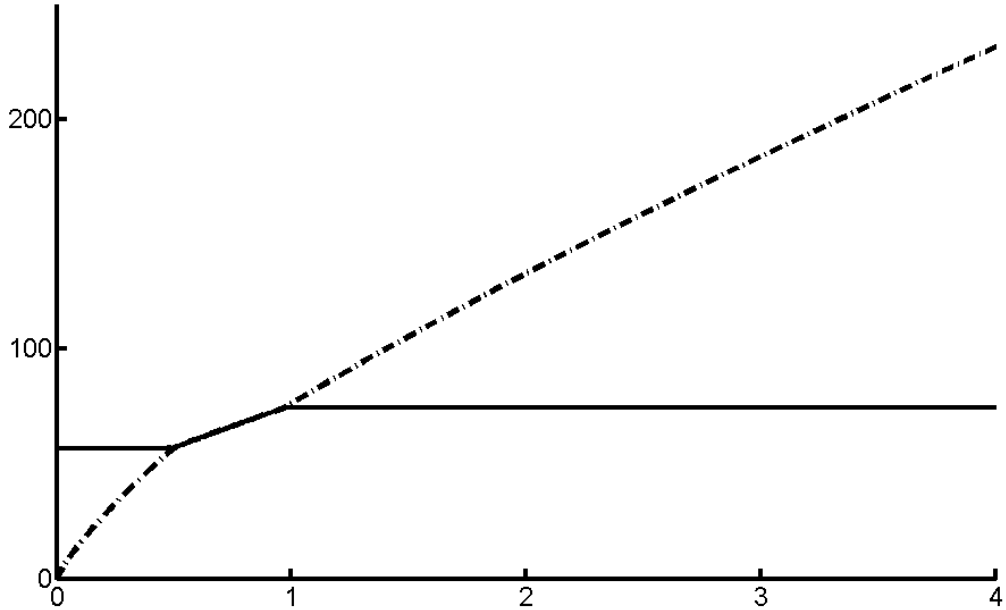

Figure 2: T_1 and T_2 as Functions of K_2 (For K_1 Fixed)

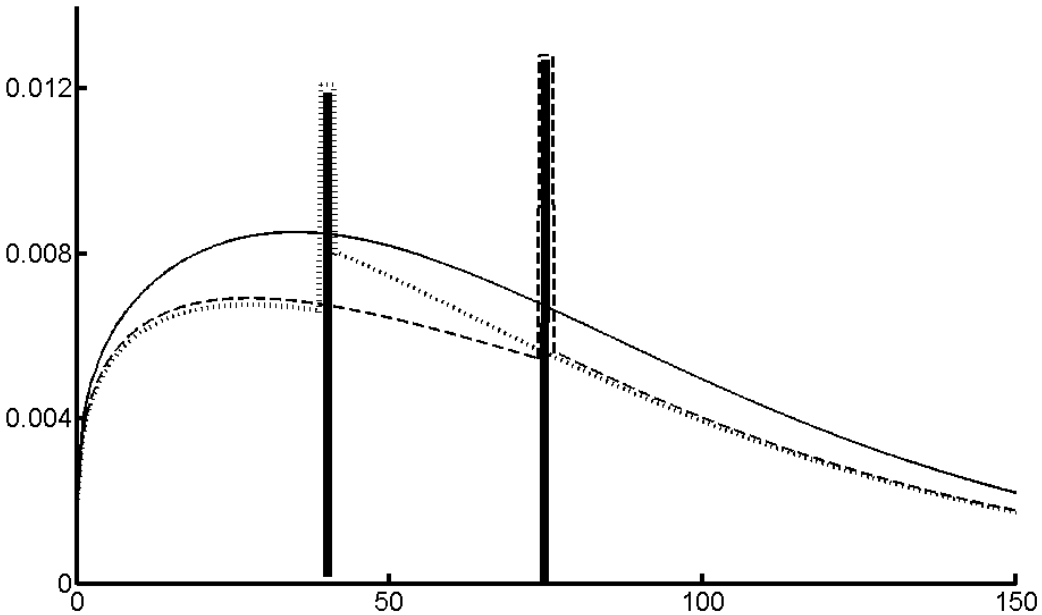

Figure 3: Marginal Density for T_1 and Conditional Given $T_2 = 45$ and $T_2 = 75$, Respectively.

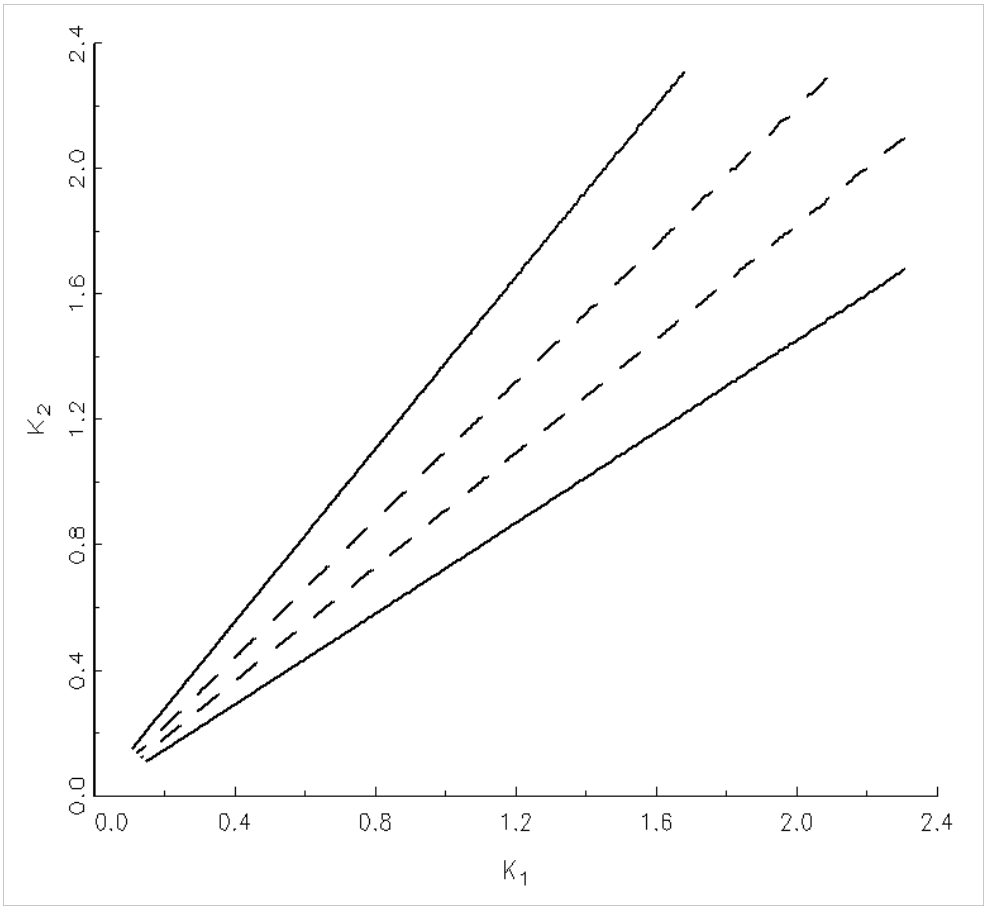

Figure 4: Joint Retirement Regions

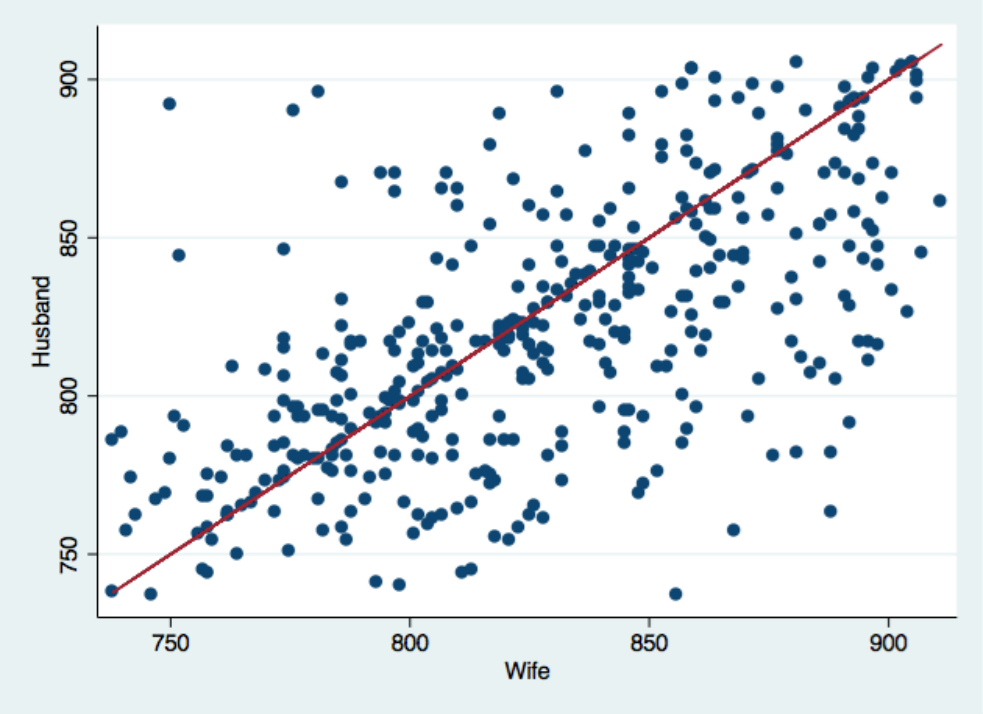

Figure 5: Retirement Months: Husband vs Wife

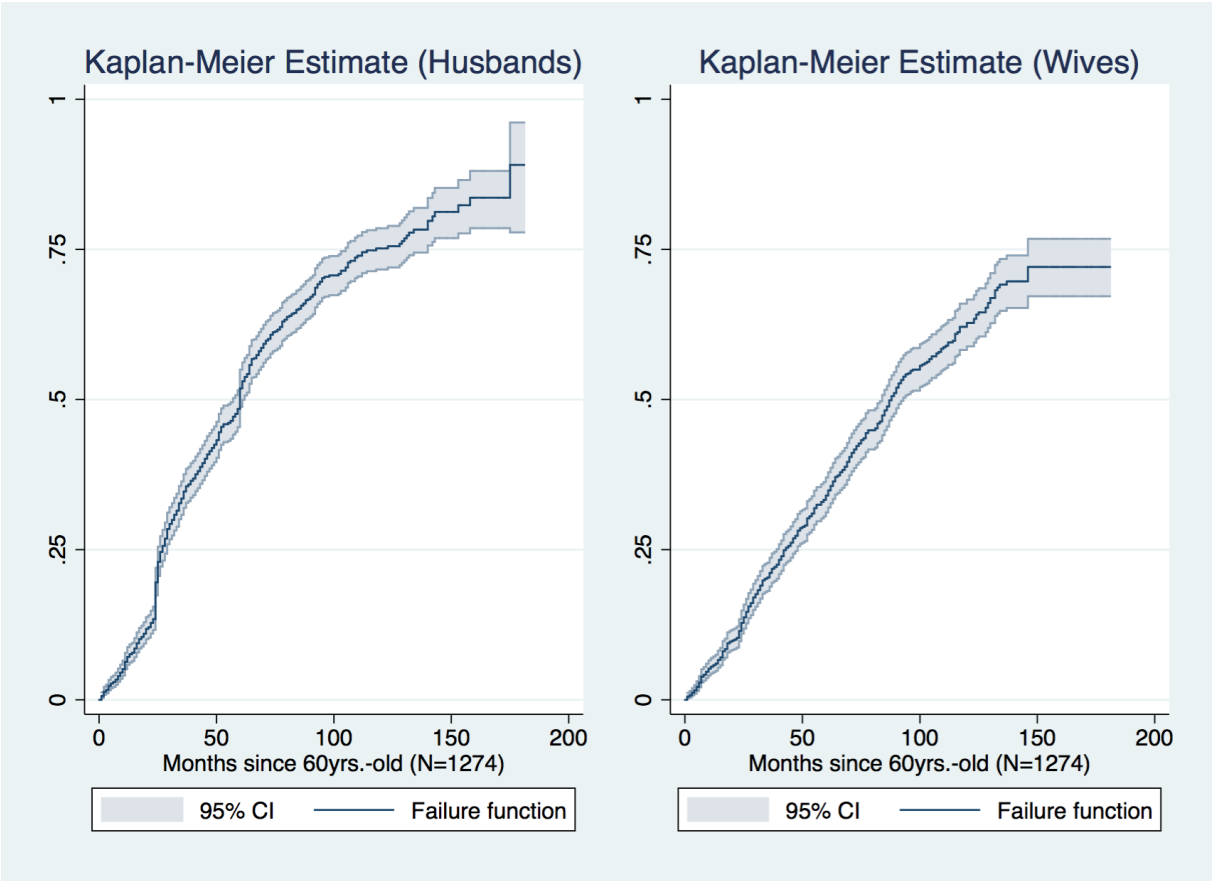

Figure 6: Kaplan-Meier Estimates: Husband and Wife

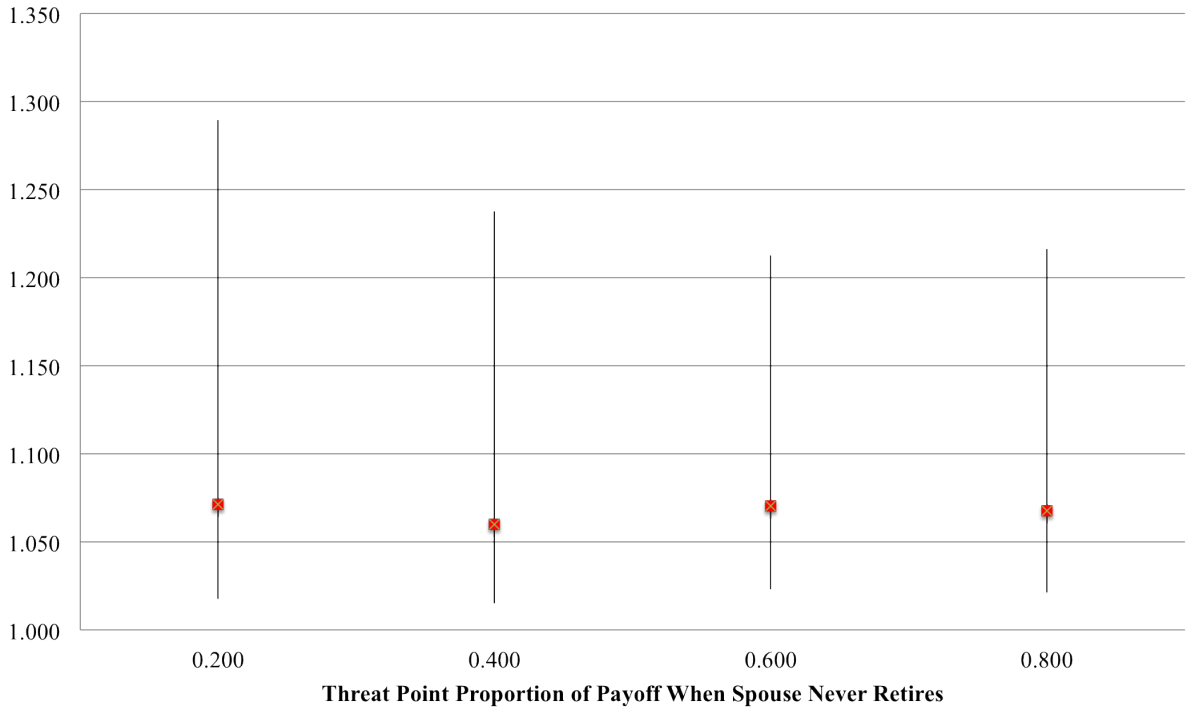

Figure 7: Robustness of $\hat{\delta}$ to Different Threat Point Specifications

Table 1: Summary statistics

	All Observations		Uncensored		Censored	
Variable	**Mean**	**N**	**Mean**	**N**	**Mean**	**N**
Gender	0.50	2568	0.56	1265	0.44	1303
Min(Ret. Month, Cens. Month)	52.67	2568	44.43	1265	60.66	1303
Censored	0.51	2568	0.00	1265	1.00	1303
Censoring Month[a]	85.04	2568	110.15	1265	60.66	1303
Age Diff.	4.11	2454	3.80	1234	4.43	1220
Non-Hisp. White	0.79	2568	0.82	1265	5.07	1303
Non-Hisp. Black	0.11	2568	0.09	1265	0.12	1303
Other Race	0.02	2568	0.02	1265	0.03	1303
Hispanic	0.08	2568	0.06	1265	0.10	1303
< High School	0.16	2568	0.18	1265	0.15	1303
HS or GED	0.36	2568	0.39	1265	0.34	1303
Some College	0.24	2568	0.22	1265	0.26	1303
College or Above	0.23	2568	0.21	1265	0.25	1303
NE	0.18	2568	0.18	1265	0.04	1303
MW	0.26	2568	0.26	1265	0.25	1303
SO	0.40	2568	0.37	1265	0.43	1303
WE	0.16	2568	0.17	1265	0.15	1303
Health Insurance	0.88	2554	0.89	1257	0.86	1297
Very Good Health	0.56	2568	0.57	1265	0.85	1296
Good Health	0.30	2568	0.30	1265	0.30	1303
Poor Health	0.14	2568	0.13	1265	0.14	1303
Pension (DB)	0.26	2568	0.30	1265	0.22	1303
Pension (DC)	0.24	2568	0.21	1265	0.27	1303
Tot. Health Expen.[b]	8.22	2180	9.48	1220	6.61	960
Financial Wealth[c]	81.22	2568	88.76	1265	73.90	1303

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.
b. Inflation-adjusted using the CPI to thousands of 2000 US dollars.

Table 2: Intra-Household Differences

	Prop. or Diff.	N of Couples
Same Race (proportion)	0.9502	1284
Same Education (proportion)	0.4618	1284
Same Self-Reported Health (proportion)	0.4938	1284
Health Insurance (both) (proportion)	0.8386	1270
Health Insurance (neither) (proportion)	0.0882	1270
DB Pension (both) (proportion)	0.0833	1284
DB Pension (neither) (proportion)	0.5600	1284
DC Pension (both) (proportion)	0.0685	1447
DC Pension (neither) (proportion)	0.5927	1284
Health Exp. (difference) (US$1,000)	3.084	1090

Only couples with no missing variables. Inflation-adjusted health expenditures in Jan/2000 USD.

Table 3: WIVES' Simultaneous Duration

Variable	**Coef.** (Std. Err.)	**Coef.** (Std. Err.)	**Coef.** (Std. Err.)	**Coef.** (Std. Err.)	**Coef.** (Std. Err.)	**Coef.** (Std. Err.)
δ	1.097	1.058	1.070	1.048	1.055	1.048
	(0.040)	(0.045)	(0.040)	(0.036)	(0.029)	(0.033)
θ_1	1.262	1.282	1.292	1.299	1.294	1.302
	(0.043)	(0.055)	(0.058)	(0.068)	(0.069)	(0.055)
Constant	-5.855 **	-5.937 **	-5.761 **	-5.950 **	-5.960 **	-6.088 **
	(0.184)	(0.241)	(0.291)	(0.401)	(0.355)	(0.388)
Age Diff.	-0.074 **	-0.074 **	-0.078 **	-0.079 **	-0.083 **	-0.077 **
	(0.015)	(0.014)	(0.013)	(0.016)	(0.018)	(0.017)
Nonhisp. Black		-0.192	-0.203	-0.117	-0.089	-0.036
		(0.148)	(0.151)	(0.180)	(0.163)	(0.182)
Other race		-0.845 **	-0.856 *	-0.580 *	-0.579 *	-0.531 †
		(0.311)	(0.381)	(0.258)	(0.243)	(0.283)
Hispanic		-0.461 †	-0.521 **	-0.372 *	-0.361 †	-0.279
		(0.280)	(0.183)	(0.189)	(0.205)	(0.217)
High school or GED		0.085	0.098	0.099	0.130	0.085
		(0.153)	(0.154)	(0.167)	(0.155)	(0.182)
Some college		-0.074	-0.063	-0.126	-0.098	-0.177
		(0.170)	(0.162)	(0.191)	(0.169)	(0.193)
College or above		0.033	0.064	0.081	0.056	-0.065
		(0.191)	(0.189)	(0.206)	(0.181)	(0.211)
NE		0.040	0.017	-0.046	-0.045	-0.092
		(0.151)	(0.149)	(0.156)	(0.164)	(0.175)
SO		0.123	0.108	0.143	0.150	0.149
		(0.114)	(0.115)	(0.126)	(0.119)	(0.125)
WE		0.208	0.205	0.256 †	0.265 †	0.266
		(0.138)	(0.134)	(0.152)	(0.148)	(0.163)
V Good Health			-0.232	-0.166	-0.129	-0.146
			(0.151)	(0.184)	(0.169)	(0.170)
Good Health			-0.279 †	-0.250	-0.228	-0.217
			(0.156)	(0.193)	(0.174)	(0.176)
Health Insurance				0.095	0.130	0.118
				(0.136)	(0.147)	(0.145)
Tot. Health Exp.				0.094	0.091	0.108
				(0.088)	(0.080)	(0.087)
Pension (DC)					-0.294 *	-0.238 †
					(0.131)	(0.130)
Pension (DB)					0.033	0.022
					(0.109)	(0.113)
Fin. Wealth						0.727 **
						(0.218)
τ	0.256	0.555	0.477	0.531	0.495	0.586
	(0.313)	(0.343)	(0.339)	(0.285)	(0.248)	(0.513)
Function Value	0.214	1.551	1.657	1.166	1.475	0.280
Number of Obs.	1227	1227	1227	1037	1037	1037

Significance levels : † : 10% $*$: 5% $**$: 1%. Significance levels are not displayed for θ_1 or δ. Ommitted categories are `Non-Hisp. White`, `Less than high school`, `Midwest or Other Region`, and `Poor Health` The threat point scale factor is 0.6, $\rho = 0.004$ and $R = 10$.

Table 4: HUSBANDS' Simultaneous Duration

Variable	**Coef.** (Std. Err.)	**Coef.** (Std. Err.)	**Coef.** (Std. Err.)	**Coef.** (Std. Err.)	**Coef.** (Std. Err.)	**Coef.** (Std. Err.)
δ	1.097	1.058	1.070	1.048	1.055	1.048
	(0.040)	(0.045)	(0.040)	(0.036)	(0.029)	(0.033)
θ_1	1.163	1.224	1.224	1.224	1.238	1.230
	(0.041)	(0.061)	(0.044)	(0.039)	(0.039)	(0.041)
Constant	-5.337 **	-5.291 **	-5.205 **	-5.426 **	-5.467 **	-5.479 **
	(0.175)	(0.249)	(0.254)	(0.292)	(0.299)	(0.279)
Age Diff.	0.023 **	0.027 **	0.028 **	0.029 **	0.027 **	0.028 **
	(0.008)	(0.007)	(0.008)	(0.009)	(0.008)	(0.008)
Nonhisp. Black		-0.207	-0.200	-0.166	-0.201	-0.179
		(0.163)	(0.157)	(0.180)	(0.154)	(0.147)
Other race		-0.154	-0.192	-0.124	-0.167	-0.161
		(0.265)	(0.253)	(0.377)	(0.212)	(0.259)
Hispanic		-0.640 **	-0.642 **	-0.615 **	-0.591 **	-0.605 **
		(0.167)	(0.199)	(0.198)	(0.229)	(0.204)
High school or GED		-0.064	-0.055	-0.118	-0.113	-0.129
		(0.113)	(0.119)	(0.130)	(0.130)	(0.130)
Some college		-0.278 *	-0.245 †	-0.324 *	-0.322 *	-0.349 *
		(0.134)	(0.137)	(0.134)	(0.144)	(0.138)
College or above		-0.466 **	-0.409 **	-0.485 **	-0.496 **	-0.536 **
		(0.123)	(0.124)	(0.131)	(0.136)	(0.144)
NE		0.099	0.090	0.106	0.127	0.117
		(0.120)	(0.125)	(0.136)	(0.132)	(0.133)
SO		-0.183 †	-0.177	-0.128	-0.131	-0.112
		(0.103)	(0.111)	(0.111)	(0.109)	(0.110)
WE		0.114	0.117	0.136	0.138	0.162
		(0.120)	(0.122)	(0.123)	(0.125)	(0.124)
V Good Health			-0.162	-0.075	-0.098	-0.090
			(0.151)	(0.152)	(0.164)	(0.152)
Good Health			-0.094	-0.051	-0.062	-0.077
			(0.154)	(0.156)	(0.163)	(0.151)
Health Insurance				0.157	0.104	0.115
				(0.151)	(0.165)	(0.153)
Tot. Health Exp.				0.139 *	0.138 *	0.137 *
				(0.057)	(0.062)	(0.059)
Pension (DC)					-0.095	-0.089
					(0.108)	(0.111)
Pension (DB)					0.307 **	0.309 **
					(0.107)	(0.110)
Fin. Wealth						0.233
						(0.199)
τ	0.256	0.555	0.477	0.531	0.495	0.513
	(0.313)	(0.343)	(0.339)	(0.285)	(0.248)	(0.280)
Function Value	0.214	1.551	1.657	1.166	1.475	0.280
Number of Obs.	1227	1227	1227	1037	1037	1037

Significance levels : † : 10% $*$: 5% $**$: 1%. Significance levels are not displayed for θ_1 or δ. Ommitted categories are `Non-Hisp. White`, `Less than high school`, `Midwest or Other Region`, and `Poor Health` The threat point scale factor is 0.6, $\rho = 0.004$ and $R = 10$.

Table 5: Simultaneous Duration by Age Diff.

	$\leq$ **3 yrs.**		$>$ **3 yrs.**	
	Wife	**Husband**	**Wife**	**Husband**
Variable	**Coef.** (Std. Err.)	**Coef.** (Std. Err.)	**Coef.** (Std. Err.)	**Coef.** (Std. Err.)
δ	1.058 (0.063)		1.077 (0.051)	
θ_1	1.340 (0.060)	1.268 (0.053)	1.332 (0.062)	1.235 (0.052)
Constant	-5.740 ** (0.450)	-5.164 ** (0.415)	-6.071 ** (0.407)	-5.363 ** (0.328)
Age Diff.	-0.054 ** (0.037)	0.060 † (0.033)	-0.063 ** (0.013)	0.036 ** (0.010)
Nonhisp. Black	0.240 (0.215)	0.190 (0.185)	-0.522 † (0.267)	-0.466 * (0.203)
Other race	-1.153 * (0.512)	-0.696 (0.626)	-0.253 (0.464)	0.030 (0.253)
Hispanic	-0.458 (0.383)	-0.794 * (0.334)	-0.759 * (0.308)	-0.580 * (0.276)
High school or GED	0.231 (0.278)	-0.012 (0.210)	-0.267 (0.263)	-0.131 (0.164)
Some college	-0.072 (0.304)	-0.175 (0.219)	-0.232 (0.260)	-0.435 * (0.185)
College or above	0.359 (0.297)	-0.329 (0.201)	-0.414 (0.317)	-0.618 ** (0.172)
NE	0.041 (0.188)	0.298 (0.186)	0.070 (0.250)	-0.004 (0.207)
SO	-0.084 (0.160)	-0.335 † (0.193)	0.414 † (0.212)	-0.013 (0.165)
WE	0.164 (0.186)	0.152 (0.191)	0.534 † (0.286)	0.087 (0.188)
V Good Health	-0.293 (0.267)	-0.387 (0.243)	-0.382 † (0.193)	-0.087 (0.167)
Good Health	-0.456 (0.299)	-0.424 (0.263)	-0.346 † (0.209)	0.026 (0.176)
τ	0.671 (0.522)		0.530 (0.505)	
Function Value	0.687		0.550	

Significance levels :† : 10% $*$: 5% $**$: 1%. Significance levels are not displayed for θ_1 or δ. Omitted categories are `Non-Hisp.` `White`, `Less than high school`, `Midwest or Other Region`, and `Poor Health`. The threat point scale factor is 0.6, $\rho = 0.004$ and $R = 5$.