Happy Together or Home Alone: A Structural Model of The Role of Health Insurance in Household Joint Retirement

Job Market Paper

Dina Guo*

University of Virginia

October 2016

*PhD in economics, University of Virginia. Email: dg2df@virginia.edu. I would like to thank my advisors Steven Stern, Leora Friedberg, and John Pepper for their time and guidance. I would also like to thank the University of Virginia's Bankard Fund and the Dissertation Fellowship from Center for Retirement Research at Boston College for financial support. Any errors are my own. The most recent version of my paper can be found at http://dinaguo.weebly.com

Abstract

The baby boomers are approaching retirement, and the majority of them are married. Simultaneously, employers are less likely to provide retiree or spousal health insurance, making it important to understand how health insurance affects couples' joint retirement decisions. I develop a dynamic programming model of household retirement in which married couples jointly decide when to retire, how to use available insurance, and how much to save. Insurance plans vary by plan characteristics (including premium, deductible, and coinsurance rates). The members of each couple coordinate their retirement decisions in response to the following motivations: (1) they share economic resources through the household budget constraint; (2) they care about spending leisure time with each other; and (3) their health insurance coverage choices are interdependent. I also model two channels through which people value health insurance: (1) insurance smooths consumption by reducing the mean and volatility of medical expenses; and (2) insurance can improve health and thus decrease individuals' value of leisure relative to work.

I estimate my model with Maximum Simulated Likelihood estimation using data from the Health and Retirement Study and the Medical Expenditure Panel Survey. I find that employer-provided health insurance (EPHI) plays an important role in retirement decisions. For workers with tied health insurance, who lose employer-provided coverage if they retire, gaining employer-provided retiree coverage would decrease the average retirement age by 1.1 and 0.5 years for husbands and wives, respectively. Similarly, raising the Medicare eligibility age is predicted to delay retirement (by 0.7 and 0.4 years), while the Affordable Care Act (ACA), which makes health insurance independent of employment status, is predicted to accelerate it (by 0.4 and 0.3 years). The effects of Medicare are bigger than the effects of the ACA but smaller than the effects of EPHI due to the differences in plan quality, which has been overlooked in the previous literature. In addition, in decomposing the employment response to EPHI coverage, I find that over 80% of the response reflects the valuation of the consumption smoothing effects of health insurance, and less than 20% reflects the valuation of the health improvement effects. Furthermore, I find that spousal coverage motivates simultaneous retirement by delaying husbands' retirement and accelerating wives' retirement, and it explains about one-fourth of the simultaneous retirement observed in the data. Lastly, I find that husbands and wives enjoy spending leisure time together, which explains nearly one-third of the observed simultaneous retirement.

JEL Classification: I13, J26, H31, C51

Keywords: household retirement, health insurance, Medicare, saving, spousal insurance coverage, health insurance plan characteristics

1 Introduction

The baby boomer generation is approaching retirement, and most of them are married. At the same time, the health insurance system has undergone massive changes. Understanding the effects of health insurance on couples' joint retirement decisions is important for evaluating health insurance policy. In the United States, the absence of retiree health insurance coverage might be one of the biggest obstacles to retiring before age 65. People with health insurance are not only better protected financially in case of serious accident or sickness; they also tend to be healthier because they are more likely to get preventative care. Health insurance coverage is particularly important for people facing retirement, who tend to have more health problems and usually are less able to financially recover from costly medical expenses. Many American workers are reluctant to retire before becoming eligible for Medicare at age 65 because they fear losing their employer-provided health insurance and because health insurance from the private market is expensive.[1] The link between health insurance and retirement is even more complicated in a household in which two spouses face different employment conditions but coordinate their retirement decisions. This paper studies how health insurance affects retirement decisions at the family level.

I develop a dynamic structural model of household retirement that accounts for the coordination between spouses and includes heterogeneity in the availability of health insurance plans and in health insurance plan characteristics. My model captures unobserved heterogeneity in household-level tastes for two household decisions–retirement decisions and health insurance plan choices. I estimate the unobserved heterogeneity parameters jointly with the preference parameters.

This paper makes three key contributions to the literature. First, it allows health insurance to affect retirement decisions through two channels: (1) insurance smooths consumption by reducing the mean and volatility of medical expenses (Rust and Phelan (1997)); and (2) insurance can improve health and thus decrease individuals' value of leisure relative to work (Currie and Gruber (1996), Levy and Meltzer (2008), and Gustman and Steinmeier (2014)). Second, my model includes the interdependence of two spouses' health insurance coverage as an additional factor that motivates coordinated retirement. Papers in the previous literature that either overlook or simplify this interdependence would risk underestimating the effect of health insurance on couples' retirement. Third, this paper differentiates each spouse's employer-provided plan by plan characteristics (premium, deductible, and coinsurance rate). Including the heterogeneity in plan characteristics enables me to model the endogenous choice of EPHI coverage. In other words, this paper evaluates two dimensions of health insurance: coverage and plan characteristics.The latter dimension has been

[1] According to a Kaiser Family Foundation survey of employers, only 28 percent of large firms with 200 or more workers offered retiree health insurance in 2010, down from 66 percent in 1988. Only three percent of small firms that have between three and 199 workers offer health plans for retirees.

largely overlooked in the literature. [2]

I estimate my model with Maximum Simulated Likelihood (MSL) estimation (Keane and Wolpin (1997), Rust and Phelan (1997), and Brien et al. (2006), and Blau and Gilleskie (2008)) using three different data sets: the Health and Retirement Study (HRS); the Medical Expenditure Panel Survey (MEPS); and the Panel Study of Income Dynamics (PSID). Using the parameter estimates, I run several counterfactual simulations to measure the causal effects of health insurance on spouses' retirement behavior. I find that employer-provided health insurance has an important influence on retirement. For workers with tied health insurance, who lose employer-provided coverage if they retire, gaining employer-provided retiree coverage would decrease the average retirement age by 1.1 and 0.5 years for husbands and wives, respectively. In decomposing the employment response to EPHI coverage, I find that over 80% of the response reflects the valuation of the consumption smoothing effects of health insurance, and less than 20% reflects the valuation of the health improvement effects. Furthermore, I find that spousal coverage motivates simultaneous retirement by delaying husbands' retirement and accelerating wives' retirement, and it explains about 24% of the simultaneous retirement observed in the data. Lastly, I find that husbands and wives enjoy spending leisure time together, which explains 34% of the observed simultaneous retirement. I also conduct several policy simulations to predict labor supply responses to the implementation of the ACA, which makes health insurance independent of employment status, and responses to changes in the Medicare and Social Security retirement rules.

The rest of the paper proceeds as follows. In Section 2, I discuss how my work relates to the literature on household retirement behavior and the relationship between health insurance and retirement decisions. I present my dynamic programming model of household joint retirement in Section 3, and, in Section 4, I describe the different data sets used to estimate the model. In Section 5, I explain the estimation strategy. Section 6 presents the estimates of my structural parameters and tests how well the model performs in various aspects. In Section 7, I run several counterfactual simulations to examine the causal effects of health insurance on retirement and to predict the labor supply responses to relevant policies. Section 8 concludes.

2 Literature Review

My paper draws from two important branches of the literature on retirement. The first branch considers the relationship between health insurance and an individual's retirement decisions. The second branch examines married couples' coordinated retirement decisions.

[2] Whereas the effect of health insurance coverage has been intensively studied (Madrian et al. (1994), Lumsdaine et al. (1994), Gustman and Steinmeier (1994), Gruber and Madrian (1995), Rust and Phelan (1997), and French and Jones (2011)), very little attention has been paid to how health insurance plan characteristics affect retirement (Gustman et al. (1994), and Fields and Mitchell (1984)).

Papers in the second branch identify two main factors that lead spouses to coordinate their retirement decisions: preference for shared leisure, and shared economic resources. However, there are very few studies that combine both branches to examine the role of health insurance in a couple's joint retirement decisions. My paper contributes to this limited literature by considering the interdependence of husbands' and wives' health insurance coverage, which is also a crucial factor in spousal retirement coordination. In addition, I model both health and medical expenses as channels linking health insurance and retirement behavior, whereas previous studies consider medical expenses as the only channel. I also add to the existing literature by modeling the heterogeneity in health insurance plan characteristics. In the subsequent sections, I comprehensively review the current literature and identify my contribution to ongoing research in these areas.

2.1 Health Insurance and an Individual's Retirement

Researchers have used different models to examine the effect of health insurance on individuals (usually elderly men's) retirement behavior, but their findings do not align. Madrian et al. (1994) and Gruber and Madrian (1995) regress the individuals' retirement decision on health insurance coverage. Both papers find that having health insurance can significantly increase the propensity to retire. Lumsdaine et al. (1994) and Gustman and Steinmeier (1994) develop structural models assuming that individuals value health insurance at the cost paid by employers. Both papers conclude that health insurance coverage has only a moderate effect on early retirement. Rust and Phelan (1997) and Blau and Gilleskie (2008) construct dynamic programming models that account for uncertain medical expenses as well as risk aversion. Both find larger effects of health insurance on older men's retirement behavior. However, these two studies overlook the possibility that older people can "self-insure" against out-of-pocket medical expenses through savings. French and Jones (2011) address this shortcoming by estimating a dynamic model that includes saving and retirement decisions. Their findings provide empirical evidence that models which exclude saving decisions overstate the effect of health insurance. Following French and Jones (2011), my paper develops a dynamic structural model of saving and retirement decisions that accounts for health insurance and uncertain medical expenses.

In contrast to my study, which examines the effects of health insurance on family retirement decisions, most papers in the literature on health insurance and individual retirement ignore spousal retirement coordination, and they exclude the dynamic relationship between one spouse's retirement and the other spouse's health insurance coverage. Ignoring this dynamic relationship may cause biased estimates of the effects of health insurance.

2.2 Married Couples' Retirement

With the increase in women's labor force participation since the 1960s, a growing number of scholars study the retirement of married couples. Regardless of the data set used, researchers find evidence of coordinated retirement among couples (Hurd (1990), Gustman and Steinmeier (2000), Blau (1998), and Coile (2004)).

Economists have identified two motivations that can lead spouses to coordinate their retirement: (1) the complementarity in leisure (Gustman and Steinmeier (2000, 2004), Banks et al. (2010)); and (2) shared economic resources affecting both spouse (Blau (1998), Michaud (2003), Coile (2004), and Kapur and Rogowski (2007)). The main limitation of these papers is that they do not consider uncertainty in future environments because they employ static models. Van der Klaauw and Wolpin (2008), Gustman and Steinmeier (2009), and Casanova (2010) expand on this literature by developing dynamic structural models that account for the uncertainty of future income, health costs, and survival upon retirement. They also contribute to this literature by jointly analyzing retirement and saving decisions. Following the literature, my dynamic model incorporates complementarity in leisure, correlation in economic environments, and uncertainty in future environments.[3]

Most papers in the literature on married couples' retirement do not pay enough attention to health insurance coverage. However, health insurance is an important factor that can link spouses' retirement behaviors for two reasons. First, health insurance can affect the household out-of-pocket medical expenditures, which in turn can affect the household budget constraints. Second, one spouse's retirement decision might affect the other spouse's health insurance coverage. In the next subsection, I discuss the few studies that focus on the effects of health insurance on married couples' retirement.

2.3 Health Insurance and Married Couples' Retirement

There is very little research examining the effects of health insurance on married couples' retirement (Blau and Gilleskie (2006) and Kapur and Rogowski (2007)), even though most people coordinate their retirement decisions with their spouses. My paper contributes to the existing literature in three ways. First, I carefully model the interdependence of the two spouses' health insurance coverage. Analysts who overlook this interdependence risk underestimating the effect of health insurance on couples' retirement. Using the information on the two spouses' employer-provided health insurance eligibility, my model constructs the available insurance plans for each spouse under each possible household retirement de-

[3] In my model, household utility is defined as a weighted sum of each spouse's utility (Van der Klaauw and Wolpin (2008) and Casanova (2010)), and the weight is imputed using spouses' discrete responses to a question in the HRS about who has the final say when major household decisions are made (Friedberg and Webb (2006)). The weight varies across households but is constant over time. The marginal utility of leisure for one spouse is modeled as a function of age, health, and the other spouse's retirement status (Gustman and Steinmeier (2000, 2004, 2009).

cision. Each spouse's health insurance coverage is modeled to depend on both spouses' retirement decisions.

Second, my paper includes heterogeneity in health insurance plan characteristics. For spouses who are eligible for both the husband and wife's employer-provided health insurance plans, my model allows them to choose one plan after comparing the plan characteristics (including paid premium, co-insurance rate, and deductible). By including this heterogeneity, I can evaluate the effects on retirement choices of policies affecting health insurance plan characteristics.[4]

Last, my paper models health as an additional channel that links health insurance and retirement. Previous papers studying health insurance and retirement assume that people value health insurance because they are risk averse and face uncertain future medical expenses (Rust and Phelan (1997), Blau and Gilleskie (2006, 2008), and French and Jones (2011)). Besides the medical expense channel, my paper models health status as an additional channel that connects health insurance coverage to the retirement decision (Currie and Gruber (1996), Levy and Meltzer (2008), and Gustman and Steinmeier (2014)).

3 Theoretical Model

In this section, I develop a finite-horizon, discrete-time, dynamic model to show how households make decisions about retirement, employer-provided health insurance (EPHI) plans, and household consumption. Each household consists of two spouses ("husband" and "wife") who are both initially working, each with his or her own preferences.[5]

At the beginning of the first period, a household knows both spouses' labor incomes and pension benefits, both spouses' EPHI eligibilities, and the characteristics (premium, coinsurance rate, and deductible) of the insurance plans provided by both spouses' employers for all periods.[6,7] In addition, at the beginning of each period, the household observes household assets and the two spouses' health and survival statuses for the period. Although the household does not observe the two spouses' medical expenditures until the end of the period, it does know the joint distribution of the two spouses' medical expenditures.

With the above information in hand, a household makes three decisions at the beginning of each period: whether each spouse should retire if he or she is working; how to

[4] Apply insurance plans with different characteristics to the same amount of total medical expenses can result out-of-pocket medical expenses with different volatility and mean. In other words, insurance plans with different characteristics can generate different financial incentives for spouses making retirement decisions.

[5] In the rest of this paper, I use "he" as a generic pronoun.

[6] Employer-provided health insurance eligibility is defined in section 3.1.2.

[7] I assume that EPHI eligibility does not depend on job choice. If individuals could become eligible for employer-provided retiree health insurance by changing jobs, then the effect of health insurance on their employment decisions might be less than what the model finds because the incentive to stay employed, and thus keep health insurance coverage, is weaker (Madrian et al. (1994)).

use available insurance; and how much the household should consume this period.[8] These three decisions remain fixed for the duration of the period.

During each two-year period, from 1992 ($t = 1$) to a known terminal period ($t = T$), each household makes the three decisions with the goal of maximizing the expected discounted value of remaining lifetime utility, subject to budget constraints. The terminal period T is a period by which both spouses in a household have retired permanently.[9] Let T^* be the period by which both spouses in a household have died.[10] For $T \leq t < T^*$, a household makes only the decision about consumption. The main features of the model are described in the following subsections.

3.1 Discrete Choice Set

At each discrete period t, other than a continuous household consumption choice, a household makes two discrete choices: the household retirement status and the EPHI plan choice.

3.1.1 Retirement Status

I define retirement as a state where an individual works less than full-time.[11] Let

$$L_{it} = \begin{cases} 1 & if\, spouse\ i\ is\ in\ retirement\, in\, period\, t \\ 0 & if\, spouse\ i\ is\ working\, in\, period\, t \end{cases} \qquad i = \{m, f\},$$

where i is m for the husband and f for the wife. Retirement is assumed to be an absorbing state: once retired, one cannot return to work full-time in any future period.[12] Let $\mathscr{L}_{it}$ denote the set of retirement statuses available to spouse i in period t,

$$\mathscr{L}_{it} = \{L_{i,t-1} \cup 1\}.$$

The household retirement status consists of the two spouses' retirement statuses, $L_t = (L_{mt}, L_{ft})$. The choice set of household retirement status in period t is

$$\mathscr{L}_t = \mathscr{L}_{mt} \times \mathscr{L}_{ft}.$$

[8]I define retirement as a state where an individual works less than full-time. Thus, I treat part-time work as retirement. Details about retirement decision are discussed in subsection 3.1.

[9]I assume that everyone has retired at age 72 (Blau and Gilleskie (2006)). T is the period in which the younger spouse in a household turns 72.

[10]I assume that everyone has died at age 100. T^* is the period in which the younger spouse in a household turns 100.

[11]Recall that all spouses are initially working full-time. I define both full retirement and part-time jobs as retirement. Blau and Gilleskie (2006) also define part-time jobs as retirement.

[12]Very few people (less than 1%) in my sample go back to full-time work from either part-time work or full retirement. Therefore, I assume in this paper that retirement is an absorbing state, and thus I do not model the transition from retirement to work.

3.1.2 Employer-Provided Health Insurance Plan Choice

A household also chooses the EPHI plan for each spouse. Let j_t denote the household EPHI plan choice, which consists of the two spouses' EPHI plan choices, $j_t = (j_{mt}, j_{ft})$. For each spouse i, j_{it} is defined as

$$j_{it} = \begin{cases} 0 & if\, spouse\; i\; has\, no\, EPHI\, coverage \\ m & if\, spouse\; i\; is\, covered\, by\; the\, husband's\; EPHI\, plan \\ f & if\, spouse\; i\; is\, covered\, by\; the\, wife's\; EPHI\, plan. \end{cases}$$

If no EPHI plan is available to spouse i, then he has no EPHI coverage. If he is eligible for just one employer-provided plan–either his own or that of his spouse–then he automatically chooses to be covered by the plan. If he is eligible for insurance plans provided by both spouses' employers, then he picks one of the two plans.[13]

The choice set of household EPHI plans in period t, $\mathscr{J}_t$, depends on the household retirement status, L_t, and the household EPHI eligibility, $e_t = (e_{mt}, e_{ft})$. Each spouse's EPHI eligibility, e_{it}, is expressed using a vector of four dummy variables, $e_{it} = (e_{it}^{w1}, e_{it}^{w2}, e_{it}^{r1}, e_{it}^{r2})$. Table 1 defines these four variables. See Guo (2016) for details about how the household EPHI eligibility, e_t, and the household retirement status, L_t, affect the choice set of household EPHI plan, $\mathscr{J}_t$.[14]

Table 1: Spouse i's EPHI Eligibility, $e_{it} = (e_{it}^{w1}, e_{it}^{w2}, e_{it}^{r1}, e_{it}^{r2})$

e_{it}^{w1}=1	if spouse i's employer provides health insurance while he is working;
e_{it}^{w1}=0	otherwise.
e_{it}^{w2}=1	if the working insurance provided by i's employer can cover the other spouse;
e_{it}^{w2}=0	otherwise.
e_{it}^{r1}=1	if spouse i's employer provides health insurance while he is retired;
e_{it}^{r1}=0	otherwise.
e_{it}^{r2}=1	if the retiree insurance provided by i's employer can cover the other spouse;
e_{it}^{r2}=0	otherwise.

3.2 Preferences

The household utility flow in period t is defined as the weighted sum of each spouse's utility. For each spouse i, his individual utility flow is a function of his retirement status, L_{it},

[13] I assume that people always choose to be covered by some EPHI plan if they are eligible for one because that is what I observe in the data.

[14] I assume that retirees switch from EPHI (if available) to Medicare at 65, and non-retirees are covered by both EPHI (if available) and Medicare until they retire. I make this assumption because the HRS data provides information only about whether an employer provides health insurance for retirees until they turn 65, and less than 6% of spouses in my sample did not enroll in Medicare when they turned 65.

the household consumption, C_t, and unobserved preferences, $\varpi_{it}(d_t, \iota_{dt})$, for the household discrete choice, $d_t = (L_t,\, j_t)$. Let $u_i(L_{it}, C_t, \varpi_{it}(d_t, \iota_{dt}))$ denote spouse i's utility flow and

$$U(d_t, C_t, \iota_{dt}) = \gamma u_m(L_{mt}, C_t, \varpi_{mt}(d_t, \iota_{dt})) + (1-\gamma) u_f(L_{ft}, C_t, \varpi_{ft}(d_t, \iota_{dt})) \tag{3.1}$$

be the household utility function, where γ is the husband's bargaining power (or some household sharing rule), which I assume is constant over time but differs across households.[15]

The utility flow for each spouse, u_i, is assumed to be non-decreasing and twice differentiable in household consumption, C_t. The function u_i is assumed to take the form

$$u_i(L_{it}, C_t, \varpi_{it}(d_t, \iota_{dt})) = \frac{C_t^{1-\alpha}}{1-\alpha} + exp\{\beta^i X_t^i\} L_{it} + \varpi_{it}(d_t, \iota_{dt}), \tag{3.2}$$

where α measures the degree of risk aversion over consumption, and $exp\{\beta^i X_t^i\}$ determines the value of leisure to spouse i. Spousal retirement status, $L_{-i,t}$, is a variable included in the vector of X_t^i.[16] The parameter associated with $L_{-i,t}$ measures spouse i's preference for retiring at the same time as his spouse (simultaneous retirement) (Gustman and Steinmeier (2000, 2004, 2009)).[17]

The unobserved preference for the household discrete choice, $\varpi_{it}(d_t, \iota_{dt})$, is modeled as the sum of three unobserved variables:

$$\varpi_{it}(d_t, \iota_{dt}) = \eta^i_{L_{it}} + \tau^i_{j_{it}} + \iota_{dt}, \tag{3.3}$$

where the first two variables, $\eta^i_{L_{it}}$ and $\tau^i_{j_{it}}$, represent spouse i's time-invariant unobserved preferences for his retirement status, L_{it}, and EPHI plan choice, j_{it}, respectively. Including these sources of unobserved heterogeneity can help explain why some spouses always choose to work and why some always choose to be covered by their own EPHI plan, even when their spousal EPHI plan has better characteristics. The last variable, ι_{dt}, is an idiosyncratic shock to the individual utility flow that spouse i receives at time t resulting from the household choice d.[18] For computational purposes, I assume that the two spouses receive the same idiosyncratic shock resulting from the household discrete choice. Including this assumption makes the household's idiosyncratic shock equal to each spouse's idiosyncratic

[15] I use a unique survey question in the HRS to impute the value of γ for each household. Details are discussed in data section 4.8.

[16] The subscript $-i$ denotes the spouse of i.

[17] Note that the two spouses' preferences for simultaneous retirement cannot be separately identified. I assume that the husband's and wife's preferences for simultaneous retirement are equal. Details are discussed in the estimation section.

[18] In the literature, it is standard to include in the individual utility an Extreme Value (EV) distributed idiosyncratic shock resulting from an individual choice. This paper allows widows or widowers to make choices to maximize their individual choice-specific value function. I assume that each spouse's idiosyncratic shock is distributed EV, so the value function of surviving spouses has a closed-form expression. Details are discussed in the estimation section.

shock. In this way, the household's shock is also distributed Extreme Value (EV).[19] In addition, I assume that the two spouses' time-persistent choice-specific individual preferences, $(\eta^i_{L_{it}}, \tau^i_{j_{it}})$, are known to the household at the beginning of the first period and that the idiosyncratic shock, ι_{dt}, is known to the household at the beginning of period t.[20]

3.3 Budget Constraints

In each period t, a household has household assets saved from the last period, A_t. In addition, the household receives household income during this period. The household has several sources of income: household asset income, rA_t, where r is the constant asset return rate;[21] the two spouses' labor incomes, $\sum_i w_{it}(1-L_{it})$, where w_{it} denotes wage; pension benefits, $\sum_i b_{it}$; Social Security benefits, $\sum_i s_{it}$; and household government transfers, g_t. The household post-tax income, y_t, is a function of taxable income (including rA_t, $\sum_i w_{it}(1-L_{it})$, and $\sum_i b_{it}$), which captures the tax structure.[22] I describe the computation of y_t in Guo (2016).

The household spends money on consumption, C_t, and on household out-of-pocket medical expenditures, $\sum\limits_{i=m,f} o_{it}$, this period. The rest of the household's money is the household assets saved for the next period, A_{t+1}. Thus, the asset accumulation equation is

$$A_{t+1} = A_t + y_t + \sum_{i=m,f} s_{it} + g_t - \sum_{i=m,f} o_{it} - C_t. \tag{3.4}$$

Households cannot borrow against future Social Security benefits, and it is very difficult to borrow against future pension benefits. Thus, I assume that the household faces the borrowing constraint,

$$C_t \leqslant A_t + y_t + \sum_{i=m,f} s_{it} + g_t. \tag{3.5}$$

Following Hubbard et al. (1995), government transfers are modeled as

$$g_t = max\{0, C_{min} - (A_t + y_t + \sum_{i=m,f} s_{it})\}. \tag{3.6}$$

[19]Note that the idiosyncratic shock to the household is a weighted sum of the shock to each spouse. If the shocks to the two spouses are different, the shock to the household is not distributed EV, even if both spouses' individual shock are distributed EV.

[20]I discuss the empirical specification for the distributions of these unobserved preferences in the estimation section.

[21]I assume that the asset return rate is constant over time. With this assumption, my model ignores the stochastic return to assets. Gustman and Steinmeier (2000) find that the stochastic return to assets has very limited effects on retirement. This is because, for average households on the cusp of retirement, stock market holdings account for a small portion of their total assets.

[22]Some people who get Social Security must pay federal income taxes on their benefits if their modified adjusted gross income (AGI) exceeds a certain amount. I assume that Social Security benefits are not taxable because the HRS data has limited information to compute modified AGI.

The parameter C_{min}, which I will estimate, is the consumption floor: the minimum amount, or sustenance level of consumption, that a household needs in every period.[23]

In the next two subsections, I describe how I model three important parts of household budget constraints: each spouse's out-of-pocket medical expenditures, o_{it}; Social Security benefits, s_{it}; and pension benefits, b_{it}.

3.4 Medical Expenditure

I model medical expenditure as a channel through which health insurance affects retirement decisions. Health insurance can affect both total medical expenditures and out-of-pocket medical expenditures. In this subsection, I first describe how I calculate each spouse's out-of-pocket medical expenditures, given the characteristics of health insurance plans and the total medical expenditures. Then, I discuss the processes that generate the two spouses' total medical expenditures.

3.4.1 Out-of-Pocket Medical Expenditure

Spouse i's out-of-pocket medical expenditures, o_{it}, are computed by applying the insurance plan characteristics (including paid premium, Γ_{it}, co-insurance rate, Λ_{it}, and deductible, Ξ_{it}) to the total medical expenditures, m_{it},[24]

$$o_{it} = m_{it} - (1 - \Lambda_{it}) max\{0, (m_{it} - \Xi_{it})\} + \Gamma_{it}. \tag{3.7}$$

3.4.2 Total Medical Expenditure

I assume that each spouse's total medical expenditures are generated by two separate processes (e.g., Pohlmeier and Ulrich (1995), Buntin and Zaslavsky (2004), and Frees et al. (2011)): (1) whether each spouse has zero or positive total medical expenditures; and (2) the amount of total medical expenditures conditional on having positive total medical expenditures. Within each process, I assume that a spouse's total medical expenditures depend on four components: (i) his health insurance coverage, I_{it}; (ii) his health status, H_{it}; (iii) his retirement status, L_{it}; and (iv) his demographic and socioeconomic factors (including age, race, education level, and household wealth), X_{it}.

To model the first process, I use a probit framework. Let an indicator variable, P_{it}, denote whether spouse i has positive total medical expenditures in period t. The latent

[23] Several papers in the literature on health insurance (e.g., Casanova (2010) and French and Jones (2011)) model government transfers in the same way.

[24] I assume that plan characteristics are exogenous. Workers do not make job choices based on the plan characteristics. Employers also do not choose plan characteristics based on their employees' health, nor do they structure plans to influence the distribution of workers with respect to health.

variable, P^*_{it}, is modeled as a function of the four components and an error term, ϑ_{it},

$$P^*_{it} = P(I_{it}, H_{it}, L_{it}, X_{it}) + \vartheta_{it}, \tag{3.8}$$

$$\begin{pmatrix} \vartheta_{mt} \\ \vartheta_{ft} \end{pmatrix} \sim N\left[\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & \rho_\vartheta \\ \rho_\vartheta & 1 \end{pmatrix}\right].$$

To model the second process, I assume that positive total medical expenditures, m^P_{it}, are a function of the same four components and an error term, u_{it}. Medical expenditures for spouses facing retirement have two main features: (1) the distribution of medical expenditures has a long right tail (Casanova (2010)); and (2) medical expenditures are volatile, and the variance of medical expenditures varies across people (De Nardi et al. (2010)).[25] To capture these two features, I model the log of positive total medical expenditures as

$$ln(m^P_{it}) = \mu(I_{it}, H_{it}, L_{it}, X_{it}) + \sigma(I_{it}, H_{it}, L_{it}, X_{it})u_{it}, \tag{3.9}$$

$$\begin{pmatrix} u_{mt} \\ u_{ft} \end{pmatrix} \sim N\left[\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & \rho_u \\ \rho_u & 1 \end{pmatrix}\right].$$

Combining the two processes described above, spouse i's total medical expenditures, m_{it}, are

$$m_{it} = P_{it} m^P_{it}. \tag{3.10}$$

3.5 Social Security and Pension Benefits

Two sources of income, Social Security and pension benefits, can generate retirement incentives. I model these two programs in detail.

3.5.1 Social Security Benefits

Social Security benefits are determined by the age at which a worker claims the benefits and by his average indexed monthly earnings (AIME), which roughly equal his average labor income during his 35 highest earnings years, adjusted for inflation. Precisely calculating AIME requires keeping track of a worker's entire earnings history, which is computationally burdensome. To simplify the computation, I follow French and Jones (2011) in assuming that spouse i's annualized AIME in the next year is a function of his annualized AIME, wage, retirement status, and age this year (described in data section 4.7).

The Social Security system has four features. First, for individuals who work less than 35 years, working more years automatically increases their AIME; and for individuals who have already worked more than 35 years, working more years increases their benefits

[25] De Nardi et al. (2010) find that the variance of medical expenditures rises with age, bad health, and income.

only if labor income earned later is higher than income earned in some previous years. Because Social Security benefits increase in AIME, this decreases work incentives after 35 years of work. Second, the age at which the individual claims Social Security affects the level of benefits. Third, each spouse is eligible for spousal benefits when he is at least 62 and his spouse is receiving Social Security benefits. Fourth, widows or widowers are entitled to survivor benefits.

In summary, the Social Security benefits that one receives depend on his own benefits, s_{it}^{O}, the survival status of his spouse, $S_{-i,t}$, his spousal benefits, s_{it}^{S}, and his survivor benefits, s_{it}^{W}. Thus, spouse i's Social Security benefits in period t, s_{it}, are

$$s_{it} = s(s_{it}^{O}, S_{-i,t}, s_{it}^{S}, s_{it}^{W}). \tag{3.11}$$

I describe details of the computation of s_{it} in Guo (2016).[26]

3.5.2 Pension Benefits

There are two main types of pension plans: defined-benefit (DB) plans and defined-contribution (DC) plans. Different types of pension plans generate different incentives. DB plans give strong incentives for retirement at specific ages.[27] On the other hand, DC plan does not encourage or discourage retirement at specific ages (Lumsdaine et al. (1996)).[28]

I model the pension benefits for DC and DB plans differently. If spouse i has a DC plan, I assume that he withdraws the pension wealth when he retires, or at the early withdrawal age if he retires earlier.[29] His pension wealth in period t, b_{it}^{DC}, depends on pension wealth in the last period, and the pension accrual this period. If spouse i has a DB plan, following French and Jones (2011), I model spouse i's DB pension benefits, b_{it}^{DB}, as a function of his age this period, and his age, PIA, and EPHI eligibility type at the time of his retirement.

In summary, spouse i's pension benefits, b_{it}, are

$$b_{it} = b_{it}^{DC} 1_{[DC]} + b_{it}^{DB} 1_{[DB]}, \tag{3.12}$$

where $1_{[DC]}$ is a variable indicating whether spouse i has a DB plan, and $1_{[DB]}$ is a variables

[26] I assume that a spouse starts claiming benefits when he retires, or at age 62 if he retires earlier than 62. In my sample, 26% of spouses claim benefits either before or after they retire. This assumption simplifies the dynamic problem because it treats Social Security benefits as a variable determined by the retirement decision and previous work and earnings history.

[27] The pension accrual rate is greatly reduced after a certain number of years of service in a firm, or past the early or normal retirement age.

[28] Benefits from DC plans are determined by a worker's salary, his and his employer's contribution rates, and the asset return rate. Most DC pension plans, such as 401(K) plans or IRAs, specify an early withdrawal age of age 59 1/2. Withdrawing benefits before this age is strongly penalized. This may encourage spouses with low household wealth to continue working at least until the early withdrawal age.

[29] Most DC plan recipients withdraw their pension wealth when they retire, and treat their pension wealth as household wealth.

indicating whether he has a DC plan. I describe details of the computation of b_{it} in Guo (2016).

3.6 Remaining Parts of The Model

In every period t, the household is uncertain about future health and survival status of the two spouses. The household adjusts its behavior in every period based on its subjective belief about future health and survival status. I allow the two spouses' health transitions to be interdependent because they experience similar events that can affect health. For similar reasons, I allow the two spouses' survival rates to be interdependent. Details of how I model household health transitions and survival rates are discussed in Guo (2016).

If one spouse dies, the surviving spouse continues to choose retirement status and consumption in order to maximize the household's remaining lifetime utility.[30] The household utility flow now is the surviving spouse's individual utility flow times his original bargaining power. If both spouses die, following De Nardi (2004), the household utility flow is modeled as a function of assets, A_t, bequeathed to survivors in the family.[31]

3.7 Value Function

I assume households are forward-looking. Each period, a household makes choices in order to maximize its present discounted value of expected lifetime utility subject to the budget constraints. Let $z_t = (A_t, H_t, S_t, \iota_{dt})$ be the state variables in period t that are either endogenous or stochastic.[32] Note that $H_t = (H_{mt}, H_{ft})$ denotes household health status, and $S_t = (S_{mt}, S_{ft})$ denotes household survival status. The optimization problem can be represented in terms of choice-specific value functions which give the lifetime discounted value of a vector of household choices, (d_t, C_t), for a given set of state variables, z_t.

The value function for a household in period t is

$$V(z_t) = \underset{(d_t, C_t)}{Max}\{v(d_t, C_t; z_t)\} \tag{3.13}$$

subject to

$$A_{t+1} = A_t + y_t + \sum_{i=m,f} s_{it} + g_t - \sum_{i=m,f} o_{it} - C_t,$$

$$C_t \leqslant A_t + y_t + \sum_{i=m,f} s_{it} + g_t,$$

[30] The surviving spouse is usually no longer eligible for the deceased spouse's EPHI. I assume that a surviving spouse is covered by his own EPHI (if available).

[31] This paper ignores divorce because less than 1% of couples in the sample divorce. Thus, I do not have enough observations to estimate any parameters associated with divorce. For a similar reason, I ignore remarriage of a widow or widower.

[32] Recall that I assume a household observes the two spouses' wage and EPHI eligibility for all period. Thus, wage and EPHI eligibility are neither endogenous nor stochastic state variables.

$$g_t = max\{0, C_{min} - (A_t + y_t + \sum_{i=m,f} s_{it})\}.$$

The term $v(d_t, C_t; z_t)$ in equation (3.13) is the choice-specific value function

$$\begin{aligned} v(d_t, C_t; z_t) &= E_m\{v(d_t, C_t; z_t | m_t)\} \\ &= U(d_t, C_t, \iota_{dt}) + \beta E_m\{E[V(z_{t+1}) \,|\, A_t, H_t, S_t, d_t, C_t, m_t]\}, \end{aligned} \tag{3.14}$$

where β is the time discount factor. $E[V(z_{t+1}) \,|\, A_t, H_t, S_t, d_t, C_t, m_t]$ is the expected value function in period $t+1$, conditional on three state variables, (A_t, H_t, S_t), household choices (d_t, C_t), and the two spouses' total medical expenses, $m_t = (m_{mt}, m_{ft})$, in period t. Given the household consumption choice in period t, the household assets at the beginning of the next period, A_{t+1}, depends on the two spouses' total medical expenses in current period, m_t. Different realizations of m_t generate different values of A_{t+1}, and then generate different values of $E[V(z_{t+1}) \,|\, A_t, H_t, S_t, d_t, C_t, m_t]$. Recall that the household does not know the realization of m_t at the beginning of period t, and thus, when the household makes choices, it knows only the expected value of $E[V(z_{t+1}) \,|\, A_t, H_t, S_t, d_t, C_t, m_t]$ with respect to the joint distribution of m_t, $F(\cdot)$,

$$E_m\{E[V(z_{t+1}) \,|\, A_t, H_t, S_t, d_t, C_t, m_t]\} = \int_{m_t} E[V(z_{t+1}) \,|\, A_t, H_t, S_t, d_t, C_t, m_t] dF(m_t). \tag{3.15}$$

Expectations in the term $E[V(z_{t+1}) \,|\, A_t, H_t, S_t, d_t, C_t, m_t]$ are taken over the idiosyncratic utility shock, $\iota_{d,t+1}$, and future household health and survival status, (H_{t+1}, S_{t+1}). At period t, a household does not know the two spouses' health and survival statuses in the next period, but it has subjective beliefs about each spouse's probability of being in good health and probability of being alive in the next period. Let $\pi(H_{t+1}, S_{t+1} | A_t, H_t, S_t, d_t,)$ denote the household's subjective beliefs about household health and survival status in the next period, (H_{t+1}, S_{t+1}), conditional on three state variables, (A_t, H_t, S_t), and household discrete choices in the current period, d_t. Then, the expected value function in period $t+1$ can be expressed as

$$E[V(z_{t+1}) | A_t, H_t, S_t, d_t, C_t, m_t] = \sum_{(H_{t+1}, S_{t+1})} EV(z_{t+1} | C_t, m_t) \pi(H_{t+1}, S_{t+1} | A_t, H_t, S_t, d_t) \tag{3.16}$$

where

$$EV(z_{t+1} | C_t, m_t) = \underset{(d_{t+1}, C_{t+1})}{Emax} \{v(d_{t+1}, C_{t+1}; z_{t+1} | C_t, m_t)\}. \tag{3.17}$$

The household uses the value functions to determine the optimal choices each period.

4 Data

4.1 Data Sources

I use data from three sources: the Health and Retirement Study (HRS), the Medical Expenditure Panel Survey (MEPS), and the Panel Study of Income and Dynamics (PSID). The primary source of data is the HRS, which is a detailed panel survey of individuals over age 50 and their spouses. It collects extensive information about household characteristics, labor force participation, health insurance coverage, health transitions, income, assets, pension plans, and health care expenditures. The HRS began in 1992 and, since then, it has re-interviewed the same households every two years.[33] I use data from the first nine waves, which cover 1992 to 2008.

Although the HRS provides most of the information I need to estimate my model, it is not an ideal data set for two types of information: total medical expenditures and health insurance plan characteristics.[34] First, the HRS data have no information on individuals' total medical expenses, which I need to predict out-of-pocket expenses on alternative plans.[35] To overcome this problem, I impute total medical expenses using MEPS, which is "a set of large-scale surveys of families and individuals, their medical providers, and their employers."[36] The MEPS began in 1996 and provides precise information about household total medical expenses. I use households in the MEPS to estimate the part of the model that determines total medical expenses (equations (3.8)-(3.9)), and then use the estimates to impute total medical expenses for each spouse each period.

Second, the HRS data have no information on health insurance plan characteristics. To impute plan characteristics for each spouse, I use MEPS, which provides the average plan characteristics (including average paid premium, co-insurance rate, and deductible) of employer-provided plans by industry type and firm size in the private sector and for government institutions in the public sector. I assign the averages of insurance plan characteristics to each spouse in my sample according to his employment industry type, firm size, and sector.[37]

Another problem with the HRS is that it collects individuals' labor income histories in a restricted file, and I cannot use this file because it can be used only on a computer that is not connected to other computers.[38] Yet, I need a worker's earnings history to calculate his

[33] New cohorts are added every three waves.

[34] Health insurance plan characteristics include premiums and deductibles (for both family and single plans), as well as the coinsurance rate.

[35] The HRS only collects data on out-of-pocket medical expenses.

[36] See the Medical Expenditure Panel Survey homepage. http://meps.ahrq.gov/mepsweb/

[37] The MEPS divides firms into two sectors: private and public. Firms in the private sector are divided into 10 industry types using the two-digit SIC (Standard Industrial Classification) codes (e.g., construction, wholesale, and retail, etc.). Details about averages of plan characteristics are discussed in subsection 4.5.

[38] I cannot use this restricted file because, to run parallel computing for my code, I need to use the High-Performance Computing system at University of Virginia, and this violates the restrictive of data rules.

average indexed monthly earnings (AIME), Social Security benefits, and pension benefits. To solve this data problem, I assume that an individual's labor income follows a modified AR(1) process (described in subsection 4.7) (Guvenen (2009)). With the estimates of this modified AR(1) process, I construct each individual's earnings history backward from the labor income recorded in the 1992 HRS. A useful data set for estimating the modified AR(1) process is the PSID, which is a longitudinal household survey that has been operating since 1968. The PSID collects data on employment and income for individuals in a nationally representative sample of households. In subsection 4.7, I describe how I use the PSID and the HRS together to impute an individual's earnings history before 1992.

4.2 Sample

I restrict my sample to couples in long-term marriages[39] in which both spouses were full-time workers in the first wave.[40] Table 2 describes the criteria that I use and the number of couples I delete due to each criterion.

Table 2: Sample Selection Criteria for the HRS Sample

Criteria	Observation Deleted	Observation Remaining
Couples with only one spouse interviewed		4,746
Changed spouses after age 35	1,310	3,436
Not both full-time workers in 1992	2,513	**923**

The resulting estimation sample is 923 couples and 6461 household-period observations. My sample does not include cohabitants because the HRS does not survey them about the length of their current relationship, and my sample requires couples in long-term relationships.

4.3 Health Insurance Eligibility

My structural model requires information about employer-provided health insurance (EPHI) eligibility for both spouses in a household.[41] However, in the HRS, the only people who report their EPHI eligibility are those covered by their own EPHI plan. Consequently, there is no information about the EPH eligibility of those who are covered by their spouse's employer.[42]

[39] Long-term marriages are defined as beginning before both spouses were aged 35. As Gustman and Steinmeier (2002) point out, in the cases of those who change spouses after age 35, it is necessary to consider how much wealth each spouse brings into the marriage and how they split obligations to children–factors that are not recorded in the HRS data.

[40] Full-time work is defined as 30+ hours per week and 36+ weeks per year.

[41] The details of EPHI eligibility are discussed in the model section.

[42] Figure A1 in Guo (2016) shows how the HRS surveys households about their EPHI coverage.

To solve the missing data problem, I need a model that accounts for two elements: (1) the relationship between the observed EPHI eligibility and explanatory variables for those who report their EPHI eligibility; and (2) the selection of reporting EPHI eligibility. It would be ideal to use structural equations to model the observed EPHI eligibility and the selection of reporting EPHI eligibility. However, modeling the EPHI eligibility in a structural way requires modeling how workers choose a job based on employer-provided plan characteristics. Job choice-related information is not recorded in the HRS data. In addition, although my structural model describes the spouses' health insurance coverage choice, using my structural model to explain the selection of reporting EPHI eligibility makes EPHI eligibility a vector of four state variables. To simplify the dynamic computation, rather than treating EPHI eligibility as a vector of four state variables, I instead assume that, at the beginning of the first period, a worker knows his EPHI eligibility for all periods. I use a multivariate probit framework to model the EPHI eligibility and the selection of reporting EPHI eligibility, and I use the estimates to impute EPHI eligibility for spouses who do not report their EPHI eligibility.[43]

First, I model a spouse's EPHI eligibility. Recall that I use a vector of four dummy variables, $e_{it} = (e_{it}^{w1}, e_{it}^{w2}, e_{it}^{r1}, e_{it}^{r2})$, to represent spouse i's EPHI eligibility.[44] Note that $e_{it}^{w2} = 0$ if $e_{it}^{w1} = 0$, $e_{it}^{r1} = 0$ if $e_{it}^{w1} = 0$, and $e_{it}^{r2} = 0$ if $e_{it}^{r1} = 0$ or $e_{it}^{w2} = 0$. Spouses who have more education are more likely to have a job that provides health insurance, and employers who provide pension benefits are more likely to provide health insurance (especially retiree health insurance). Thus, the latent variable of each of these dummy variables is modeled as a function of one's demographic characteristics (including race, gender, age, health, and education), X_{it}^{d}, his employment characteristics (including firm size, working hours, hourly wage, tenure, and pension availability), X_{it}^{E}, and an error term. These four latent variables can be expressed as

$$e_{it}^{w1*} = e(X_{it}^{d}, X_{it}^{E}; \varsigma_{w1}) + \omega_{it}^{w1}, \tag{4.1}$$

$$e_{it}^{w2*}|_{[e_{it}^{w1}=1]} = e(X_{it}^{d}, X_{it}^{E}; \varsigma_{w2}) + \omega_{it}^{w2}, \tag{4.2}$$

$$e_{it}^{r1*}|_{[e_{it}^{w1}=1]} = e(X_{it}^{d}, X_{it}^{E}; \varsigma_{r1}) + \omega_{it}^{r1}, \tag{4.3}$$

$$e_{it}^{r2*}|_{[e_{it}^{r1}=1; e_{it}^{w2}=1]} = e(X_{it}^{d}, X_{it}^{E}; \varsigma_{r2}) + \omega_{it}^{r2}. \tag{4.4}$$

Different combinations of these four binary variables represent different types of EPHI eligibility. I assume that there are three mutually exclusive categories of EPHI eligibility. The first category is None, which consists of workers whose employers provide no health insurance at all. The second category is Tied, where workers receive their employer-provided insurance as long as they continue to work. The third category is Retiree, where workers

[43]Because the multivariate probit model is used to impute EPHI eligibility, I hereafter describe it as the EPHI eligibility imputation model.

[44]The definitions of these dummy variables are discussed in the model section.

keep their health insurance even after leaving their jobs. Each category consists of one or more types of EPHI eligibility. Table 3 defines the three categories and the six types of EPHI eligibility using the four binary variables.

Table 3: EPHI Eligibility

Eligibility Category	Eligibility Type	e^{w1}	e^{w2}	e^{r1}	e^{r2}
None	N1N2	0	0	0	0
Tied	T1N2	1	0	0	0
	T1T2	1	1	0	0
Retiree	R1N2	1	0	1	0
	R1T2	1	1	1	0
	R1R2	1	1	1	1

Note: In the eligibility type column, N, T, and R are short for None, Tied, and Retiree, and numbers 1 and 2 represent self and spouse, respectively. For example, T1T2 means that a sposue is eligible for his EPHI only when he is working, and this working health insurance can cover his spouse.

Next, I model the selection of reporting EPHI eligibility. Let an indicator variable, e^s_{it}, denote whether spouse i reports his EPHI eligibility. Spouses who report EPHI eligibility are those who are covered by their own EPHI plan, either because they are not eligible for their spouse's EPHI, or because their employer provides better plans than their spouse's employer. I assume that one's own EPHI plan characteristics are correlated with employment characteristics. Spouses who have similar employment characteristics might have different coverage. For example, some are covered by their own EPHI plan because their spouse has a part-time job that does not provide insurance; while others are covered by their spousal EPHI plan because their spouse works for a big firm that provides a good spousal plan. Thus, whether a spouse is covered by his own EPHI is correlated with his own and his spouse's employment characteristics. The latent variable, e^{s*}_{it}, is modeled as

$$e^{s*}_{it} = e^s(X^d_{it}, X^E_{it}, X^E_{-i,t}; \varsigma_s) + \omega^s_{it} \tag{4.5}$$

where X^d_{it}, X^E_{it}, and $X^E_{-i,t}$ represent spouse i's demographic characteristics, his employment characteristics, and his spouse's employment characteristics, respectively; and ω^s_{it} is an error term. Let $\omega_{it} = (\omega^{w1}_{it}, \omega^{w2}_{it}, \omega^{r1}_{it}, \omega^{r2}_{it}, \omega^s_{it})'$, and assume that $\omega_{it} \sim N(0, \Sigma_\omega)$. The EPHI eligibility imputation model consists of equations (4.1)-(4.5). To estimate this model, I use married people who are working in the HRS. The estimates are used to predict the probability of the occurrence of each type of EPHI eligibility for spouses in my sample who do not report their EPHI eligibility.[45]

Tables 4 and 5 list the sample statistics for EPHI eligibility for husbands and wives,

[45] The details of the estimation of the EPHI eligibility imputation model are discussed in Guo (2016).

respectively, in the first wave.[46] In both tables, the third column lists the statistics for observed (binary) EPHI eligibility for those who report their EPHI eligibility. The fourth column lists the statistics for imputed (continuous) EPHI eligibility for those who do not report their EPHI eligibility. The last columns in tables 4 and 5 list statistics for EPHI eligibility for all husbands and all wives, respectively.

Table 4: EPHI Eligibility Distribution in 1992: Husband

		Observed		Imputed		All Husbands	
Category	Type	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
None	N1N2	0.144	(0.35)	0.196	(0.40)	0.156	(0.35)
Tied	T1N2	0.034	(0.18)	0.107	(0.31)	0.051	(0.17)
	T1T2	0.192	(0.39)	0.149	(0.33)	0.182	(0.44)
Retiree	R1N2	0.005	(0.07)	0.009	(0.01)	0.006	(0.06)
	R1T2	0.024	(0.15)	0.054	(0.15)	0.031	(0.16)
	R1R2	0.601	(0.49)	0.485	(0.33)	0.574	(0.44)

Note: See the note in Table 3 for an explanation of the type column.

Table 5: EPHI Eligibility Distribution in 1992: Wives

		Observed		Imputed		All Wives	
Category	Type	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
None	N1N2	0.142	(0.35)	0.146	(0.33)	0.144	(0.33)
Tied	T1N2	0.042	(0.20)	0.102	(0.14)	0.070	(0.16)
	T1T2	0.243	(0.43)	0.228	(0.34)	0.236	(0.31)
Retiree	R1N2	0.012	(0.11)	0.008	(0.02)	0.010	(0.09)
	R1T2	0.030	(0.17)	0.028	(0.21)	0.029	(0.15)
	R1R2	0.531	(0.49)	0.488	(0.32)	0.511	(0.44)

Note: See the note in Table 3 for an explanation of the type column.

The last column in Table 4 shows that about 63% of husbands in my sample are in the retiree category, and 23.3% of them are in the tied category. The last column in Table 5 shows that about 55% of wives in my sample are in the retiree category, and 30.6% of them are in the tied category.[47]

4.4 Employment

Figures 4.1-4.2 present some of the labor market behavior that my model should explain. Figure 4.1 shows the HRS full-time job participation rates by health insurance category. Regardless of the health insurance category, the full-time job participation rate declines

[46] In the data, the observed EPHI eligibility type rarely changes over time, and thus, I assume that each spouse's EPHI eligibility type remains the same.

[47] In contrast, the Department of Labor estimates that in 1993, 44% of full-time employees in medium and large private establishments offered retiree health insurance (Clark (1999). My sample might overpredict the percentage of people in the retiree category and underpredict the percentage of people in the tied category.

with age, and the decline is especially sharp between the ages of 62 and 65. Those in the tied category have the highest participation rates, while those in the retiree category have the lowest participation rates. Note that the participation rates for workers in the none category are higher than those for workers in the retiree category. This may be because, in my sample, most workers in the none category are self-employed. On average, self-employed workers have a strong labor market attachment.[48]

Figure 4.1: Full-Time Job Participation Rates: Husband

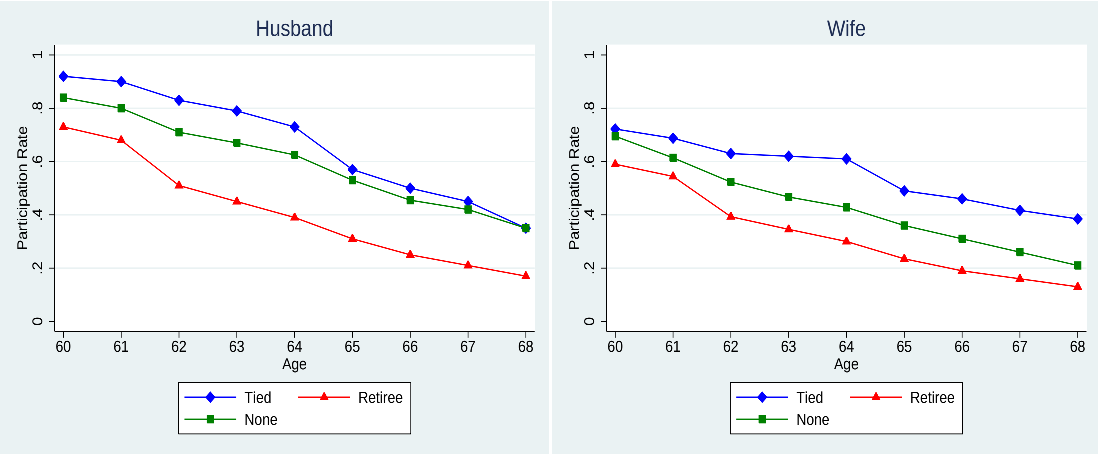

Figure 4.2 shows the HRS full-time job exit rates by health insurance category. For both husbands and wives, the job exit rates jump at age 62 and at age 65. At age 62, the jump in the job exit rates is large for workers in the retiree category. This pattern provides evidence that being eligible for SSB gives workers in the retiree category incentives to retire. At age 65, the jump in the job exit rates is large for workers in the tied category. This pattern provides evidence that workers in the tied category tend to work until age 65, when they become eligible for Medicare. French and Jones (2011) also find these patterns for males in different insurance categories. Casanova (2010) finds jumps in the job exit rates at age 62 and at age 65 for both husbands and wives, but she does not distinguish between spouses in different insurance categories.

At age 65, the job exit rate is similar for those in the tied and the retiree categories. At almost every age other than 65, workers in the tied category have lower job exit rates than those in the retiree or the none category. For example, at age 62, the job exit rates for husbands in the tied, retiree, and none categories are 7.8%, 25%, and 11.3%, respectively. In Rust and Phelan (1997), at around age 62, the job exit rates for males in the tied and the non-tied (including the retiree and the none) categories are 7.5% and 25.9%, respectively. Compared to Rust and Phelan (1997), I find a smaller job exit rate for males in the non-tied category. This might be because: (1) the males in my sample have higher household in-

[48]Several papers (e.g., Zissimopoulos et al. (2007) and Hochguertel (2015)) find that self-employed workers aged 50 and over are less likely to retire than wage and salary workers. One potential reason for this phenomenon is that self-employed workers have low access to pension benefits.

Figure 4.2: Full-Time Job Exit Rates: Husband

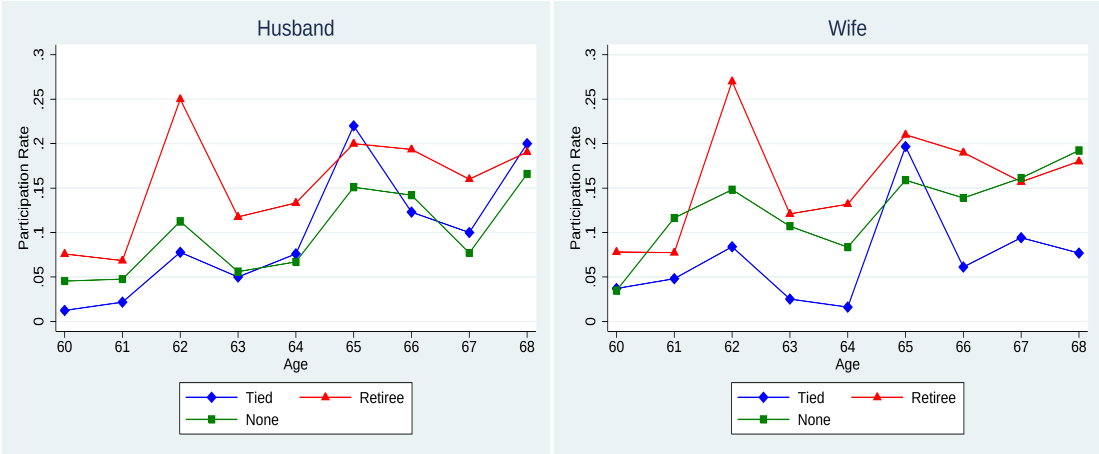

comes, and thus, the social security benefits account for a smaller portion of their retirement income and generate smaller incentives for them to retire;[49] and (2) most of the males in the none category in my sample are self-employed workers, and they, on average, have a strong labor market attachment.

Furthermore, at almost every age, the job exit rates for workers in the tied category are lower than those for workers who are eligible for retiree insurance. This suggests that factors other than differences in health insurance eligibility cause differences in job exit rates.[50]

4.5 Employer-Provided Health Insurance Plan Characteristics

To compute household budget constraints, this paper requires the following information about EPHI plan characteristics: the co-insurance rate, the paid premium and deductible for family coverage, and the paid premium and deductible for single coverage. Unfortunately, the HRS data have no information on plan characteristics. To solve this data problem, I use information provided in MEPS to impute plan characteristics for each spouse in my sample. The MEPS Insurance Component surveys private- and public-sector employers to collect data on their private health insurance plan characteristics. Firms in the private sector are divided into cells by industry type and firm size. Specifically, there are ten industry types and two firm size categories (more or less than 100 employees), which leads to total of 20 cells. All firms in the public sector are treated as one cell. The data are used to generate the average of each EPHI plan characteristic for these 21 cells. For each spouse in my sample, I define his EPHI plan characteristics as the averages of the cell to which his firm belongs.

Using the cell averages in the MEPS to impute the plan characteristics of each

[49] Rust and Phelan (1997) include low- to middle-income males who initially aged 58-63 from 1969-1979.

[50] As French and Jones (2011) point out, those who are eligible for retiree health insurance usually have higher pension wealth than those who are eligible for tied health insurance. Thus, pension wealth could be another factor that causes workers who have retiree health insurance to retire earlier.

spouse in my sample poses two problems: (1) the MEPS does not contain data before 1996 (the "censored data problem"); and (2) the cell averages of plan characteristics can differ significantly from the true individual plan characteristics (the "measurement error problem"). In the rest of this subsection, I describe and address the censored data problem. I discuss how the measurement error problem is likely to affect the model estimates in Guo (2016).

The MEPS collecting information about paid premiums and deductibles in 1996, and it began collecting information about co-insurance rates in 2002. Thus, I cannot observe paid premiums and deductibles for periods before 1996, and I cannot observe co-insurance rates for periods before 2002. To impute missing information about paid premiums and deductibles, I assume that the real paid premiums and deductibles in 1992 and 1994 are the same as those in 1996. This assumption is reasonable because the paid premiums increased only modestly from the mid-1990s to the late 1990s (Cutler (2003)).[51] I calculate the nominal values of the paid premiums and deductibles in 1992 and 1994 using the Medical Care Price Index (MCPI) from the St. Louis Federal Reserve Bank as a deflator.[52]

To impute missing information about co-insurance rates, I assume that the co-insurance rates in 1992, 1994, 1996, 1998, and 2000 equal those in 2002. This assumption is reasonable because the co-insurance rates remained approximately the same from the mid-1990s to the early 2000s (Eibner and Marquis 2008).[53]

4.6 Wage

Annual wage is defined as the product of hourly wage, working hours per week, and working weeks per year. However, in the HRS, information about annual wage can be missing for two reasons. First, some individuals did not report one or more of these variables. Second, as in the well-known wage selection problem, econometricians cannot observe what annual wage retired people would have earned.

To solve the missing data problem, I use a wage equation to impute annual wage for spouses whose wage cannot be observed in the HRS. I model the log real annual wage in period t as

$$ln(w_{it}) = \beta^w X_{it}^w + u_{it} \tag{4.6}$$

where X_{it}^w is a vector of explanatory variables (including age, gender, education, and annual working hours), and $u_{it} \sim N(0, \sigma_u^2)$ is an error term. To reduce computational burden, I estimate the wage equation separately from my structural model. Details about the estimation of the wage equation are discussed in Guo (2016).

[51]Cutler (2003) finds that the nominal paid premiums increased by less than 6% between 1993 and 1999.

[52]The Medical Care Price index (MCPI) was 190.058, 211.025, and 228.267 in 1992, 1994 and 1996, respectively.

[53]Elbner and Marquis (2008) find that the co-insurance rates either declined modestly or remained about the same from 1995 to 2003.

4.7 Average Indexed Monthly Earnings (AIME)

As described in model section 3.5, a spouse's AIME is a key factor that determines Social Security and pension benefits. Precisely calculating AIME requires keeping track of a worker's entire earnings history, which is computationally burdensome. Following French and Jones (2011), I assume that spouse i's annualized AIME in the next year is a function of his annualized AIME, Δ_{it}, labor income, $w_{it}L_{it}$, and age, a_{it}, in the current year.[54] I refer to this function as the AIME updating function hereafter.

Using the AIME updating function and the annualized AIME in 1992, I compute the annualized AIME under each possible retirement choice for every year after 1992. To calculate the AIME in 1992, I need a worker's earnings history up to 1992. Due to the HRS data usage restriction of individuals' earnings history,[55] I use an AR(1) process and the PSID data to impute the earnings history prior to 1992 for each spouse in my sample. Details are discussed in Guo (2016).

4.8 Bargaining Power

In the first wave, the HRS asks each spouse a question about decision-making power: "When it comes to making major family decisions, who has the final say-you or your (husband/wife/partner)?"[56] Friedberg and Webb (2006) treat the answers as noisy discrete measures of the true continuous bargaining power, and they develop a bivariate ordered probit model to examine the effects of explanatory variables on bargaining power.[57] Using the estimates and the explanatory variables included in their paper, I impute the continuous bargaining power for each husband in my sample.

4.9 Summary Statistics

Table 6 shows the summary statistics for the first period for the variables used in my model.[58] Two variables that I haven't previously discussed are household race and health status. Household race is defined as white if both spouses are white, black if both spouses are black, and other if otherwise. Most (84%) households in my sample are white, and very few (5%) households have mixed races. To measure the health status of the HRS respondents, I examine their answer to the question: "Would you say your health is excellent, very good, good, fair or poor?"[59] Health status is defined as good if the answer is excellent, very

[54]The details of the AIME updating function are described in Guo (2016).

[55]For an explanation of the data usage restriction, see data section 4.1.

[56]See the Health and Retirement Study Questionnaire, Question E135.

[57]Details of the model developed by Friedberg and Webb (2006) are described in Guo (2016).

[58]Because the summary statistics for EPHI eligibility are listed in Tables 4 and 5, I exclude them from Table 6.

[59]See the Health and Retirement Study Questionnaire, Question B1.

good, or good, and bad if the answer is fair or poor.[60]

Table 6: Summary Statistics For The Initial Distribution

Variables	Mean	Std Dev
Bargaining Power	0.51	0.10
Household Assets [a]	197	229
AIME_h [a]	24.6	4.34
AIME_w [a]	13.8	2.71
Annual Income_h [a]	35.7	21.5
Annual Income_w [a]	21.9	13.8
Health_h [b]	0.89	0.32
Health_w [b]	0.93	0.25
Age_h	54.5	3.02
Age_w	51.6	3.11
EPHI Plan Characteristics		
co-insurance rate_h	0.18	0.02
paid premium-Family_h [c]	1220	198
deductible-Family_h [c]	1026	269
paid premium-Single_h [c]	295	62
deductible-Single_h [c]	464	114
co-insurance rate_w	0.18	0.01
paid premium-Family_w [c]	1333	164
deductible-Family_w [c]	1045	238
paid premium-Single_w [c]	303	58
deductible-Single_w [c]	456	94
Household Race		
white	0.84	0.36
black	0.11	0.31
other	0.05	0.21
Education_h		
no degree	0.16	0.37
high school	0.60	0.49
college+	0.24	0.43
Education_w		
no degree	0.12	0.32
high school	0.70	0.46
college+	0.18	0.39

Note: a: measured in thousands of dollars; b: fraction in good health;
c: measured in dollars.

[60] Although several researchers raise concerns about the reliability of self-reported health status, I assume, following Benitez-Silva et al. (2004), that it is reported accurately.

5 Estimation

In this section, I first define the structure of errors in my theoretical model (see subsection 5.1). Secondly, I use the error structure to define the likelihood function and describe the estimation strategy (see subsection 5.2). I conclude by discussing how the parameters in the model are identified (see subsection 5.3).

5.1 Error Structure

The household utility function includes two types of errors: an idiosyncratic error, ι_{dt}, and unobserved heterogeneity terms, $(\eta^i_{L_i}, \tau^i_{j_i})$ (equation (3.3)). As described in the model section, the term ι_{dt} accounts for the household utility shock that varies across discrete choices and time. The term $\eta^i_{L_i}$ is spouse i's time-invariant preference for his retirement status L_i, and the term $\tau^i_{j_i}$ is spouse i's time-invariant preference for his EPHI plan choice j_i. Both types of errors are assumed to be known to the household but not the econometrician.

I assume that $\iota_{dt} \sim iid\,EV(0, \sigma_\iota, 0)$. The parameter σ_ι is the scale of the distribution of the idiosyncratic error. The vector of unobserved heterogeneity, $u^h = (\eta, \tau)$, is assumed to take two values, u^h_1, u^h_2, with probabilities p^u, $1 - p^u$, and I restrict $u^h_1 = 0$. (Heckman and Singer (1985), Keane and Wolpin (1997), Cameron and Heckman (1998), Mroz (1999)).

5.2 Likelihood Function

As described in the model section, before both spouses in a household retire, the household makes three decisions at the beginning of each period to maximize the expected present discounted value (EPDV) of its remaining lifetime utility. The three decisions include two discrete choices (household retirement status and EPHI plan choice, $d_t = (L_t,\, j_t)$) and one continuous choice (household consumption, C_t). After both spouses in a household retire, the household chooses only the household consumption, C_t, at the beginning of each period to maximize the EPDV of its remaining lifetime utility.

Including the density of the observed C_t in the likelihood function not only increases the computational burden; it also requires a data set with quality information about household assets and tracks each household's asset transitions over time. However, the HRS fails to keep track of asset transitions.[61] To avoid the computational burden and the data problem, I assume that, in each period, households optimally choose C_t, conditional on the discrete decision d_t (Rust (1988) and Casanova (2010)). This means that the Bellman equation (equation (3.13)) can be rewritten as

$$V(z_t) = \max_{d_t}\{\max_{C_t}\{v(d, C_t; z_t) | d_t = d\}\} \tag{5.1}$$

[61] For example, for several households in my sample, the difference between the household assets observed in two adjacent waves is much bigger than the observed household income gained during the period.

where $z_t = (A_t, H_t, S_t, \iota_{dt})$ is a vector of state variables in period t (defined in model section 3.8) and $v(d_t, C_t; z_t)$ is the choice-specific value function (defined in equation (3.14)).

In the rest of this subsection, I first describe how to compute optimal consumption, conditional on a discrete decision d_t (see 5.2.1). Next, I discuss how to compute choice probabilities (see 5.2.2). Then, I derive the likelihood function (see 5.2.3). Finally, I discuss the estimation strategy (see 5.2.4).

5.2.1 Optimal Consumption

As defined in the model section (equations (3.14)), $v(d_t, C_t; z_t)$ is

$$v(d_t, C_t; z_t) = U(d_t, C_t; z_t, \iota_{dt}) + \beta E_m\{E[V(z_{t+1})|z_t, d_t, C_t, m_t]\} \tag{5.2}$$

where $E_m\{E[V(z_{t+1})|z_t, d_t, C_t, m_t]\}$ is defined in equation (3.15).[62] Given any d_t, optimal household consumption, denoted as $C^*(d_t)$, satisfies the First Order Condition (FOC) of equation (5.2),

$$\frac{\partial U(d_t, C_t; z_t)}{\partial C_t} + \beta \frac{\partial E_m\{E[V(z_{t+1})|z_t, d_t, C_t, m_t]\}}{\partial A_{t+1}} \frac{\partial A_{t+1}(C_t)}{\partial C_t} = 0. \tag{5.3}$$

Computing the term $\frac{\partial E_m\{E[V(z_{t+1})|z_t, d_t, C_t, m_t]\}}{\partial A_{t+1}}$ analytically is computationally burdensome. My approach is to use four linear splines to approximate $E_m\{E[V(z_{t+1})|z_t, d_t, C_t, m_t]\}$. Let $C(d_t, q)$ be the the optimal household consumption for the q^{th} spline.[63] The optimal household consumption, conditional on d_t, is

$$C^*(d_t) = \underset{C_t \in \{C(d_t, q), q=1,\ldots,4\}}{argmax} \{v(d_t, C_t; z_t)\}. \tag{5.4}$$

5.2.2 Choice Probabilities

Before the two spouses in a household retire, the household's optimization problem in each period is to choose the d_t that maximizes $v(d_t, C^*(d_t); z_t)$. Because $v(d_t, C^*(d_t); z_t)$ is a function of d_t only, I simplify the notation by rewriting $v(d_t, C^*(d_t); z_t)$ as $v^*(d_t, ; z_t)$. I decompose the term $v^*(d_t; z_t)$ as the sum of three parts:

$$v^*(d_t; z_t) = \hat{v}(d_t; z_t) + (\eta_L + \tau_j) + \iota_{dt} \tag{5.5}$$

where $\eta_L = \gamma \eta^m_{L_m} + (1-\gamma)\eta^f_{L_f}$ and $\tau_j = \gamma \tau^m_{j_m} + (1-\gamma)\tau^f_{j_f}$ are the household's time-invariant preferences for the household retirement status, L, and the household EPHI plan choice, j,

[62] The term $v(d_t, C_t; z_t)$ is strictly concave in C_t because: (1) the household utility flow (equations (3.1)-(3.2)), $U(\cdot)$, is strictly concave in C_t; (2) the distribution function of m_t (equations (3.8)-(3.10)) are independent of C_t; and (3) the transition functions in $E[V(z_{t+1})|z_t, d_t, C_t, m_t]$ (equation (3.16)) are independent of C_t. Thus, given any d_t, there exists a unique household consumption decision that maximizes $v(d_t, C_t; z_t)$.

[63] Details about the computation of $C(d_t, q)$ are discussed in Guo (2016).

respectively.[64] In any period before both spouses retire, the probability of interest for a household is the probability that the observed discrete decision $d^* = (L^*, j^*)$ generates a higher value of $v^*(d_t; z_t)$ than any other possible discrete choice $d = (L, j)$. Because the idiosyncratic error, ι_{dt}, is assumed to be distributed Extreme Value, the choice probability, conditional on the household EPHI eligibility type and the unobserved heterogeneity, is

$$P_t(d_t = d^*; z_t, \theta | e, \eta, \tau) = \frac{exp\{[\hat{v}(d^*; z_t|e) + (\eta_{L^*} + \tau_{j^*})]/\sigma_\iota\}}{\sum_{d \in \mathscr{L}_t \times \mathscr{J}_t} exp\{[\hat{v}(d; z_t|e) + (\eta_L + \tau_j)]/\sigma_\iota\}}. \tag{5.6}$$

where θ is the full set of parameters to be estimated and e is the household EPHI eligibility type.[65]

5.2.3 Likelihood Function

Let T_n^R be the period by which both spouses in household n have retired, and let T_n^L be the last period by which the household has stayed in the HRS. For household n, the number of periods included in its likelihood contribution is

$$T_n = min\{T_n^R, T_n^L\}.$$

Given panel data $\{z_t^n; d_t^n\}$ $(t = 1, ..., T_n; n = 1, ..., N)$ on the observed states and discrete decisions of N households, the likelihood contribution for household n, conditional on the household EPHI eligibility type and the unobserved heterogeneity, (e^n, η, τ), is

$$L^n(\theta | e^n, \eta, \tau) = \prod_{t=1}^{T_n} \left\{ P_t(d_t^n; z_t^n, \theta | e^n, \eta, \tau) \pi_t(z_t^n, \theta_\pi | z_{t-1}^n, d_{t-1}^n) \right\} \tag{5.7}$$

where $\pi_t(\cdot|\cdot)$ represents the household's health transition probabilities and survival rates. The vector of parameters $\theta = (\theta_P, \theta_m, \theta_\pi)$ includes three parts: (1) the vector of parameters that affect the household utility function and the consumption floor, θ_P; (2) the vector of parameters that affect household total medical expenses, θ_m; and (3) the vector of parameters that affect the the household's subjective belief about future events, θ_π.[66]

Recall that $u^h = (\eta, \tau)$ is assumed to take two values, $u_1^h = (\eta^1, \tau^1)$ and $u_2^h = (\eta^2, \tau^2)$, with probabilities p^u and $(1 - p^u)$. Thus, the likelihood contribution for household n, conditional on the household EPHI eligibility type, e^n, is

$$L^n(\theta | e^n) = \sum_{k=1}^{2} \left\{ p^k L^n(\theta | e^n, \eta^k, \tau^k) \right\} \tag{5.8}$$

where $p^1 = p^u$ and $p^1 = 1 - p^u$.

[64] As introduced in model section 3.2, the parameter γ is the husband's bargaining power.

[65] As I described in data section 4.3, I assume that EPHI eligibility type does not change over time.

[66] This decomposition is useful later in the estimation strategy (see section 5.2.4.).

As described in data section 4.3, for spouses who are covered by their spouse's EPHI plan, their EPHI eligibility is not reported in the HRS. For these spouses, I use the EPHI eligibility imputation model (equations (4.1)-(4.5)) to impute the probability for each of the six possible types of EPHI eligibility.[67] This means that, for some households, I know only the distribution of the household EPHI eligibility over the 36 possible types.[68] Let $P_n^e(e^n = e)$ represent the probability of having a particular EPHI eligibility type for household n. Then, a household likelihood contribution over the distribution of the household EPHI eligibility type is

$$L^n(\theta) = \sum_{e=1}^{36} P_n^e(e^n = e) L^n(\theta|e). \tag{5.9}$$

The likelihood function is

$$L(\theta) = \prod_{n=1}^{N} L^n(\theta). \tag{5.10}$$

A key part of computing the likelihood function, $L(\theta)$, is to compute the term $\hat{v}(d_t; z_t|e)$, which is

$$[U(d_t, C^*(d_t); z_t, \iota_{dt}|e) - (\eta_L + \tau_j) - \iota_{dt}] + \beta E_m\{E[V(z_{t+1})|e, z_t, d_t, m_t]\}.$$

Thus, computing $\hat{v}(d_t; z_t|e)$ requires an evaluation of integrals over the joint distribution of the two spouses' total medical expenses, $m_t = (m_{mt}, m_{ft})$. Recall that, when a household makes decisions, it does not observe the two spouses' medical expense-related shocks, $(\vartheta_{mt}, \vartheta_{ft}, u_{mt}, u_{ft})$, which later determine the two spouses' total medical expenses, m_t.[69] This means that computing $\hat{v}(d_t; z_t|e)$ requires an evaluation of the four-dimensional integrals over the joint distribution of $(\vartheta_{mt}, \vartheta_{ft}, u_{mt}, u_{ft})$.

Evaluation of the four-dimensional integrals in $\hat{v}(d_t; z_t|e)$ is not possible analytically, so I calculate the value of $\hat{v}(d_t; z_t|e)$ numerically using a simulation method. I make R draws of $(\vartheta_{mt}, \vartheta_{ft}, u_{mt}, u_{ft})$ from their joint distribution defined in model section 3.4.2, and I compute the term $\hat{v}(d_t; z_t|e, \vartheta_{mt}^r, \vartheta_{ft}^r, u_{mt}^r, u_{ft}^r)$ for each random draw. I define the simulated analog to $\hat{v}(d_t; z_t|e)$ as

$$\hat{v}^R(d_t; z_t|e) = \frac{1}{R}\sum_{r=1}^{R} \hat{v}(d_t; z_t|e, \vartheta_{mt}^r, \vartheta_{ft}^r, u_{mt}^r, u_{ft}^r). \tag{5.11}$$

The simulated choice probability, $P_t^R(d_t; z_t, \theta|e, \eta, \tau)$, is computed by replacing all $\hat{v}(d_t; z_t|e)$ terms in equation (5.6) with their simulated analogs, $\hat{v}^R(d_t; z_t|e)$. The simulated conditional likelihood contribution for household n, $L^{nR}(\theta|e^n, \eta, \tau)$, is computed by replac-

[67] Different types of individual EPHI eligibility are defined in data section 4.3.

[68] There are six possible types of EPHI eligibility for a spouse, and thus, there are 36 possible types of household EPHI eligibility for a household.

[69] As described in model section 3.4.2, each spouse has two medical expense-related shocks, (ϑ_{it}, u_{it}). The first determines whether a spouse has positive medical expenses, and the second determines the total medical expenses conditional on having positive medical expenses.

ing $P_t(d_t; z_t, \theta|e, \eta, \tau)$ in equation (5.7) with their simulated analogs, $P_t^R(d_t; z_t, \theta|e, \eta, \tau)$. The simulated unconditional likelihood contribution is

$$L^{nR}(\theta) = \sum_{e=1}^{36} P_n^e(e^n = e) \left[\sum_{k=1}^{2} p^k L^{nR}(\theta|e, \eta^k, \tau^k) \right], \tag{5.12}$$

and the simulated likelihood is

$$L^R(\theta) = \prod_{n=1}^{N} L^{nR}(\theta). \tag{5.13}$$

The estimate $\hat{\theta}$ is the vector of parameter values that maximizes the simulated likelihood function $L^R(\theta)$.[70]

5.2.4 Estimation Strategy

While a relatively small number (16) of parameters are used to specify a household's preferences, a large number (98) of parameters are needed to specify its health transition function and survival function and to specify the distribution of the two spouses' medical expenses (I refer to it as computational burden problem). In addition, the HRS data have no information on the two spouses' total medical expenses (I refer to it as omitted variables problem).

Due to the computational burden and the omitted variables problem, I use a two-stage estimation procedure to estimate the model (Rust (1987, 1988), Rust and Phelan (1997), and French and Jones (2011)). In the first stage, I use the HRS demographic data and Maximum Likelihood (ML) Estimation to estimate the parameters (θ_π) that determine household health transitions and survival rates. I also use the MEPS data and ML Estimation to estimate the parameters (θ_m) that determine total medical expenditures (equations (3.8)-(3.10)). In the second stage, I use my HRS sample and MSL Estimation (Keane and Wolpin (1997), Rust and Phelan (1997), Brien et al. (2006), and Blau and Gilleskie (2008)) to estimate the parameters (θ_P) that determine utility function (equations (3.1)-(3.3)) and the consumption floor (equation (3.6)).[71]

5.3 Identification

In this subsection, I discuss identification of preference parameters (θ_P). In Guo (2016), I discuss identification of other parameters (θ_π and θ_m).

Recall that household utility is defined as the weighted sum of each spouse's utility,

[70] Although the Maximum Simulated Likelihood (MSL) estimator, $\hat{\theta}$, is inconsistent, the magnitude of the inconsistency is frequently small, and with a modest R, MSL estimation can construct an practically consistent estimator (Börsch-Supan and Hajivassiliou (1993), and Hajivassiliou (2000).

[71] As discussed in Rust and Phelan (1997), the two-stage estimation procedure is not as efficient as the full likelihood estimation using the full likelihood function (equation (5.13)).

u_i. Spouse i's utility depends on household consumption, C_t, his retirement status, L_{it}, and his unobserved preferences, $\varpi_{it}(d_t, \iota_{dt})$, for household discrete choices, d_t.[72] As defined in model section 3.2, spouse i's utility function is

$$u_i(L_{it}, C_t; \varpi_{it}(d_t, \iota_{dt})) = \frac{C_t^{1-\alpha}}{1-\alpha} + exp\{\beta^i X_t^i\} L_{it} + \varpi_{it}(d_t, \iota_{dt}) \quad i \in \{m, f\}$$

where

$$\beta^i X_t^i = \beta_0^i + \beta_1^i a_{it} + \beta_2^i H_{it} + \beta_3^i L_{-i,t}. \tag{5.14}$$

The risk aversion parameter α is identified by the co-variation in household savings and future uncertainty. For example, suppose that two risk averse households are similar in everything except that one has family health insurance coverage and one does not. The household with health insurance coverage will save less than the household without insurance coverage for two reasons. First, the one with insurance coverage has less uncertainty about future medical expenses (the "risk reducing" aspect of insurance). Second, the health insurance company helps to pay part of the medical expenses for the household with insurance coverage (the "gift" aspect of insurance). For healthy people, the effect of the gift aspect of insurance on saving decisions is limited because their expected total medical expenses are small. Therefore, the degree of risk aversion is identified by the extent to which healthy, insurance-eligible households save less than healthy, ineligible households.

For spouse i, β^i is a vector of parameters that affect his preference for leisure. The parameters associated with spouse i's age and health, (β_1^i, β_2^i), are identified by the co-variation in spouse i's retirement status, L_{it}, and his age and health status, (a_{it}, H_{it}), conditional on other observable characteristics. The last element of β^i, β_3^i, represents spouse i's preference for spending leisure time with the other spouse (simultaneous retirement). Both β_3^m and β_3^f are identified by the co-variation in the two spouses' retirement choices: retiring together or not. Thus, these two parameters, (β_3^m, β_3^f), cannot be identified separately, and I assume that $\beta_3^m = \beta_3^f$.[73]

Recall that u^h represents the vector of the two spouses' unobserved time-invariant preferences for different discrete household choices.[74] The vector u^h differs across households, but its distribution remains the same. I assume that u^h takes two values, u_1^h, u_2^h, with probabilities p^u, $1 - p^u$, and I restrict $u_1^h = 0$. The parameters u_2^h and p^u are identified by the variance of the household time-specific residuals and the variance of the household-specific residuals.

An identification problem I face is the difficulty of separately identifying the effects of unobserved heterogeneity from duration dependence. This is because both factors

[72]As I defined in model section 3.2, $L_{-i,t}$ denotes the other spouse's leisure (or retirement status).

[73]Gustman and Steinmeier (2000, 2004, 2009) allow the husband and the wife to have asymmetric preferences for spending leisure with the other spouse. Their ability to separately identify these two parameters relies on an important assumption: the decision-making process in a family is a non-cooperative bargaining process.

[74]Details of u^h are explained in section 5.1.1.

can account for the observed decreasing hazard out of employment over time. Elbers and Ridder (1982) prove that variations in observed explanatory variables that are included in hazard functions provide enough restrictions on the observed hazard to separately identify the effects of unobserved heterogeneity and duration dependence. Because the choice probability functions (or hazard functions) in my model include explanatory variables that vary across people and time, the effects of unobserved heterogeneity can be separately identified from duration dependence.

Additionally, there are three parameters that affect household choices but do not appear in the household utility function directly. These parameters are: (1) the time discount factor, β; (2) the consumption floor, C_{min}; and (3) the scale of the distribution of idiosyncratic utility shocks, σ_ι. The time discount factor is identified by the co-variation in household asset quantile and retirement decision. The consumption floor is identified by the same co-variation in households in the bottom asset quantile. If the consumption floor is sufficiently low, the risk of a catastrophic medical expense shock will encourage work (delaying retirement) among the poor. Conversely, a high consumption floor discourages work among the poor (Hubbard et al. (1995)). The scale parameter can be identified separately from other parameters because the value function is nonlinear in other parameters (equation (5.6)).[75]

6 Results

In this section, I first discuss the parameter estimates. Then, I assess the model's performance using three different specification tests.

6.1 Model Parameter Estimates

In this subsection, I present the preference parameter estimates that affect household utility. I discuss the parameter estimates that determine the distribution of total medical expenditures, health transitions, and survival rates in Guo (2016).

Recall that a household's utility is the weighted average of each spouse's utility, weighted by each spouse's bargaining power (equation (3.1)). A spouse's utility function depends on his leisure, his time-invariant preferences for discrete choices (unobserved heterogeneity), household consumption, and an idiosyncratic shock (equation (3.2)). His preference for leisure depends on his age, his health, and the other spouse's leisure.

Table 7 presents estimates of preference parameters that affect household utility. The first and the second panels list the estimates of parameters that determine husbands' and wives' preferences for leisure, respectively. The sign of these parameter estimates are

[75] Rust and Phelan (1997) and Blau and Gilleskie (2006) include EV distributed idiosyncratic shocks in their models, and they normalize the scale parameter to 1.

as expected and are significant at 1% significance level. The estimates show that, for both husbands and wives, their preferences for leisure increase with age. Husbands and wives value leisure more if they are in bad health or their spouse is retired. The estimated coefficients of age and health are similar to those found in Gustman and Steinmeier (2000, 2004). Unlike these two papers, I assume that husbands and wives value spousal leisure (shared retirement) equally. I make this assumption because the two parameters representing the two spouses' preferences for shared leisure cannot be separately identified (explained in estimation section 5.3.1). For Gustman and Steinmeier (2000, 2004), their ability to separately identify these two parameters relies on an important assumption: the decision-making process in a family is a non-cooperative bargaining process. My estimated coefficient of spousal leisure falls between the estimates of the husbands' and wives' preferences for shared retirement found in these two papers.[76] To put the size of the estimates into perspective, I compute the average marginal effect (AME) for each of these variables. The AMEs show that increasing the age by one year increases the propensity to retire by 0.9 and 1.2 percentage points for husbands and wives, respectively. Being in good health decreases the propensity to retire by 3.3 and 2.9 percentage points for husbands and wives, respectively. Spending leisure with the other spouse increases the propensity to retire by 2.4 and 3.2 percentage points for husbands and wives, respectively.

The third panel presents the parameter estimates associated with the distribution of unobserved heterogeneity. The first estimate in this panel, 0.72, is the estimated probability of being Type I unobserved heterogeneity. Recall that Type I unobserved heterogeneity is assumed to be a vector of zeros. The rest of the estimates in this panel are the estimated Type II unobserved heterogeneity. These estimates show that there exists, at least, one factor that makes husbands prefer working over retirement and prefer obtaining health insurance from their own employer rather than from their wife's employer. In addition, this factor makes wives prefer retirement over working and prefer obtaining insurance from their husband's employer rather than from their own employer. This factor could be the unobserved ability/willingness to do house work and care for family members. The estimate of the wife's preference for her own employer-provided health insurance (EPHI) relative to spousal EPHI has a much larger scale than other estimates, which indicates that my model does not explain the observed wives' EPHI coverage choices very well. This might be because my model does not include the notion that many wives prefer to let husbands provide family health insurance coverage, which allows wives to be more flexible and to take care of family members.

The last panel shows other parameters that affect household utility. My estimated relative risk-aversion coefficient is 1.67, which lies within the range of $[1,5]$ that many

[76]Gustman and Steinmeier (2000, 2004) found that the wife's preference for retirement is small and statistically insignificant and the husband's preference for shared retirement is large and statistically significant.

Table 7: Estimates of Preference Parameters

	Estimate	Std. Err.	AME
Husband's Preference for Leisure			
H's age	0.589**	0.143	0.009
H has good health	-0.573**	0.147	-0.033
W's leisure	0.792**	0.129	0.024
Constant	-1.583**	0.353	
Wife's Preference for Leisure			
W's age	0.503**	0.035	0.012
W has good health	-1.262**	0.378	-0.029
H's leisure	0.792**	0.129	0.032
Constant	-1.523**	0.269	
Unobserved Heterogeneity			
Probability of type I	0.724**	0.024	
H's taste for working relative to retirement	0.555**	0.144	
H's taste for own EPHI relative to spousal EPHI	0.425**	0.109	
W's taste for working relative to retirement	-0.819**	0.247	
W's taste for own EPHI relative to spousal EPHI	-4.561**	0.488	
Relative risk aversion	1.674**	0.019	
Time discount factor	0.942**	0.163	
Consumption floor	7,734**	1,578	
The Scale of idiosyncratic errors (σ_ι)	0.816**	0.204	

Note: 1) * is 5 percent significance level, and ** is 1 percent significance level;
2) variable age here is (age-55);
3) the column AME lists average marginal effects of variables on the propensity to retire.

economists believe contains this parameter (Chetty (2006)).[77] This implies that my estimate of the coefficient of relative risk aversion is reasonable. Another important estimate is the time discount factor. Each period contains two years, and thus, the estimated time discount factor 0.94 is a biannual discount factor. The yearly time discount factor is 0.97, which matches the empirical time discount factor widely assumed in the literature. For example, Blau and Gilleskie (2006) fix the annual time discount factor at 0.96; Rust and Phelan (1997) fix it at 0.98. My estimated household consumption floor is $\$7,734$ per year, which is similar to the consumption floor assumed or estimated in the literature. For example, Casanova (2010) assumes that the household consumption floor is $\$7,596$, and French and Jones (2011) estimate that the individual consumption floor is $\$4,380$.

6.2 Specification Tests

I conduct three different specification tests to assess my model's performance. First, I test how well the model fits the data by comparing the predicted and observed labor market be-

[77] The estimated values of the relative risk-aversion coefficient reported by Rust and Phelan (1997), Blau and Gilleskie (2006, 2008), and French and Jones (2011) are 1.1, 1.8, 1.0 and 5.0, respectively.

havior by age and health insurance types (see 6.2.1). Then, I use a Chi-Square Goodness-of-Fit test to assess whether the distribution of predicted probabilities matches the distribution of observed frequencies of household retirement choices (see 6.2.2). Last, I use three Lagrange Multiplier (LM) tests to check whether the model is properly specified (see 6.2.3).

6.2.1 Model Fit

To examine how well the structural model fits the data, I simulate full-time work participation rates and job exit rates for each spouse in my sample, using the model and the estimated parameters. I compare these simulated labor market behaviors with the observed labor market choices in the HRS.

Participation Rates

Figures 6.1-6.3 compare the simulated and observed full-time work participation rates by age and health insurance categories. The focus is on employment patterns for spouses who are in three different health insurance categories: (1) health insurance is available only while employed ("Tied HI"); (2) health insurance is available while employed and retired ("Retiree HI"); and (3) no insurance is available, whether employed or retired ("No HI"). The area between the two grey, smooth lines represents the 95 percent confidence interval for the predicted participation rates. As these figures show, although the simulated data differ from the actual data, the observations fall within the 95 percent confidence interval for all points.

My model is able to replicate three key features of how full-time work participation rates vary with age and health insurance categories. First, my model is able to capture the overall pattern that full-time participation rates decline with age. Second, my model is able to capture several sharp declines around age 62 and 65. Specifically, the model captures (a) the sharp declines in participation rates at age 65 for both husbands and wives who have tied health insurance; (b) the sharp declines at age 62 and 65 for husbands who have retiree insurance; and (c) the sharp decline at age 62 for wives who have retiree insurance. Third, my model is able to capture the large differences in full-time work participation rates across health insurance categories.

Job Exit Rates

Figures 6.4-6.6 compare the simulated and observed full-time job exit rates by age and health insurance categories for husbands and wives, respectively. As these figures show, most of the observations fall within the 95 percent confidence interval for the predicted job exit rates. This indicates that my model fits the overall patterns of job exit rates quite well. To be more specific, the model captures the spike in job exit rates at age 65 for both husbands and wives who have tied health insurance. It also captures the spike in job exit rates

Figure 6.1: Participation Rates: Tied EPHI

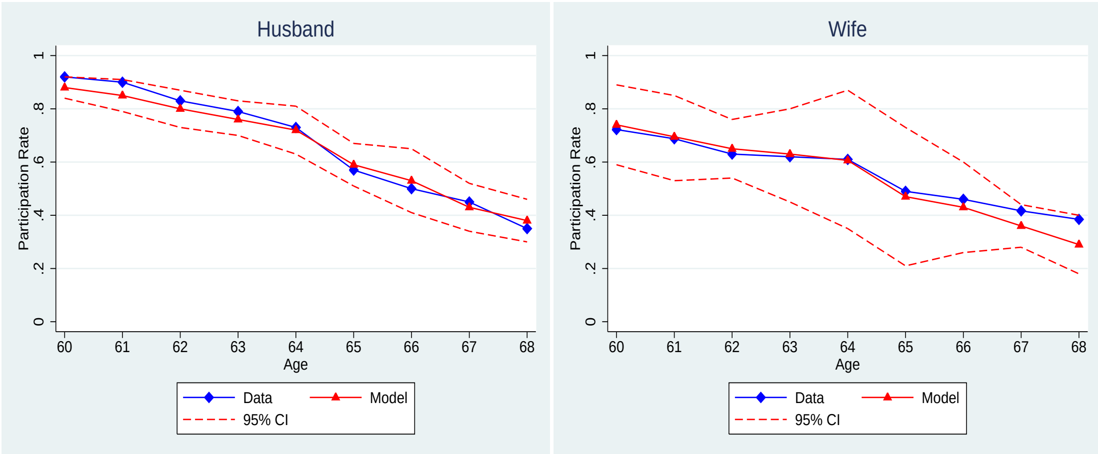

Figure 6.2: Participation Rates: Retiree EPHI

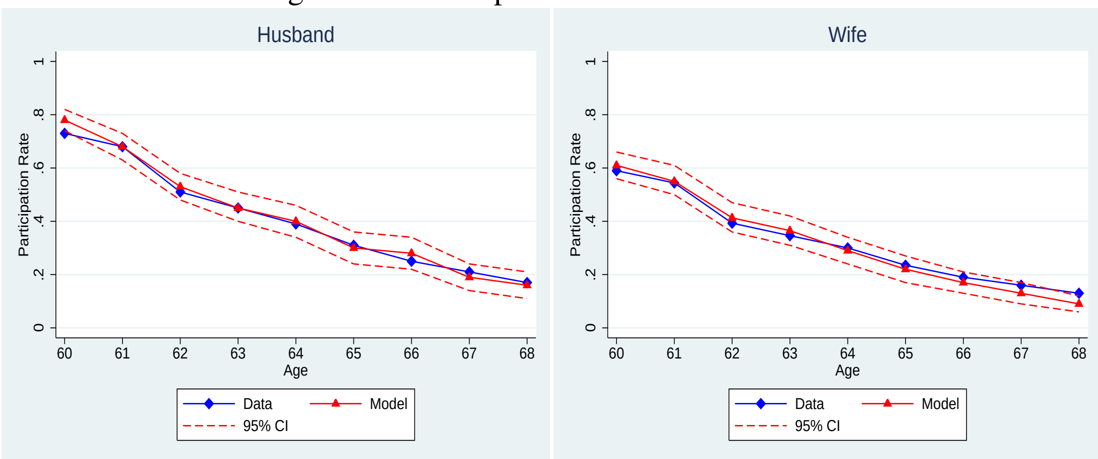

Figure 6.3: Participation Rates: No EPHI

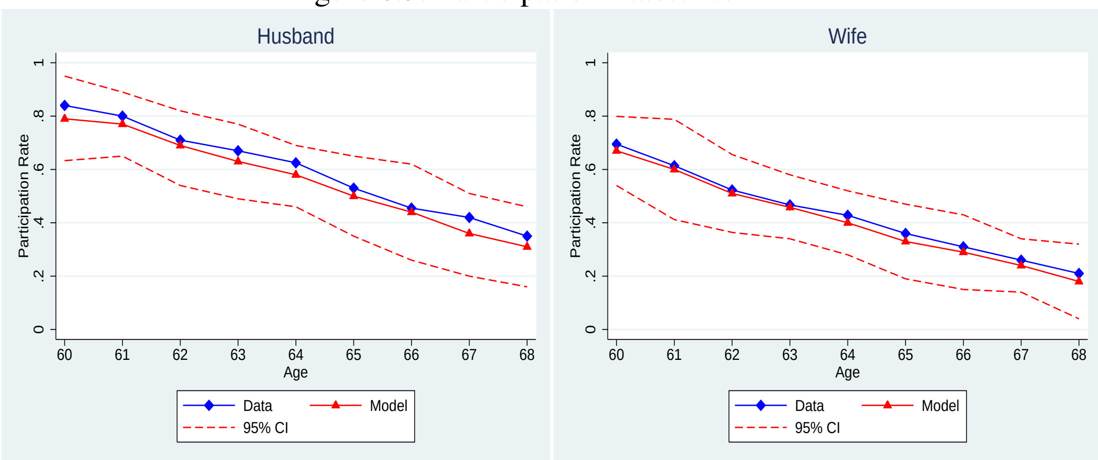

at age 62 and 65 for both husbands and wives who have retiree health insurance. However, some differences exist between the simulated data and the actual data. For example, for husbands with tied insurance, the model underpredicts the age-65 job exit rate. For husbands with retiree insurance, the model underpredicts the age-62 job exit rate and overpredicts the age-65 job exit rate. For wives with tied insurance, the model overpredicts the age-65 job exit rate, and for wives with retiree insurance, the model underpredicts the age-62 job exit rate. For both husbands and wives with no insurance, the model overpredicts the job exit rates at almost every age. The mismatch for spouses with no insurance may be due to the following reasons: (1) the sample size of households that have no health insurance is very small; and (2) most spouses who have no EPHI in my sample are self-employed, and self-employed people may have a high preference for working due to some unobserved factors, such as workaholism, that are not included in my model.

The ability of this paper to capture the spikes at ages 62 and 65 relies mostly on two aspects of the model: (1) carefully model the interaction between two spouses' financial benefits (including Social Security and pension benefits); and (2) include the interdependence of two spouses' health insurance. Papers ignore the interdependence of two spouses' health insurance may underpredict the high job exit rate at ages 65. This is because some spouses with retiree insurance may delay their retirement to age 65, in order to provide health insurance to their spouses.[78]

Figure 6.4: Job Exit Rates: Tied EPHI

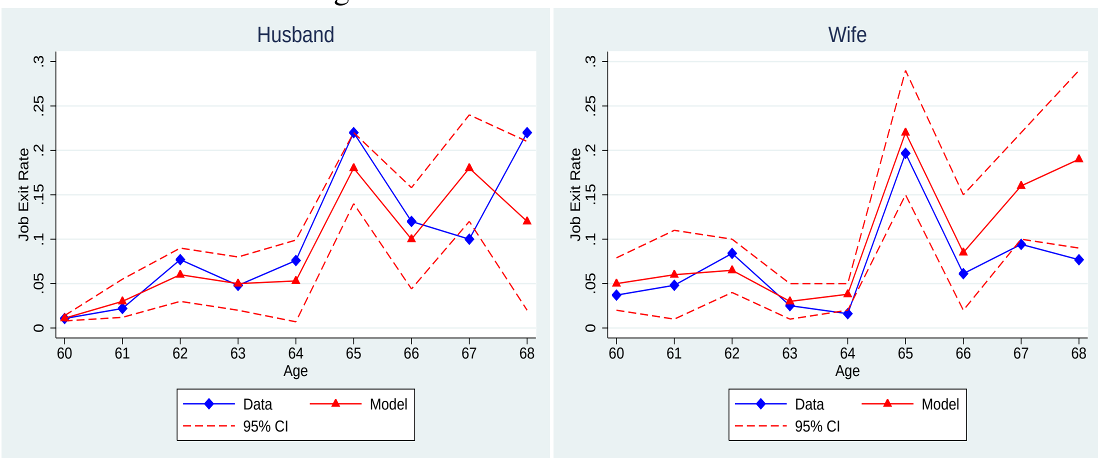

6.2.2 Chi-Square Goodness-of-Fit Test

In my model, there are four possible household labor supply choices: (1) both spouses work full-time; (2) only husband works full-time; (3) only wife works full-time; or (4) both spouses are retired. For each household in my sample, the model predicts the probability of these four household choices. For each household labor supply choice, I divide households into five strata by the percentile of their predicted probabilities and then compare the average predicted probabilities to the observed frequencies for each stratum. The null hypothesis I test is that the distributions of observed frequencies and predicted probabilities

[78]For example, Blau and Gilleskie (2006) do not pick up the especially high exit rates at ages 62-65, and they underpredict the difference in job exit rates between workers with and without retiree insurance. French and Jones (2011) fail to capture the spike at 65 for those who have retiree insurance.

Figure 6.5: Job Exit Rates: Retiree EPHI

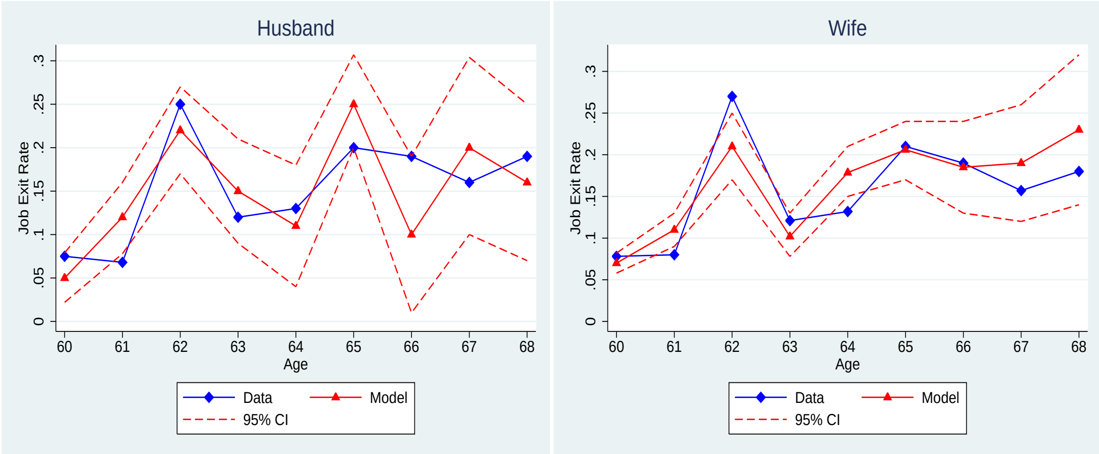

Figure 6.6: Job Exit Rates: No EPHI

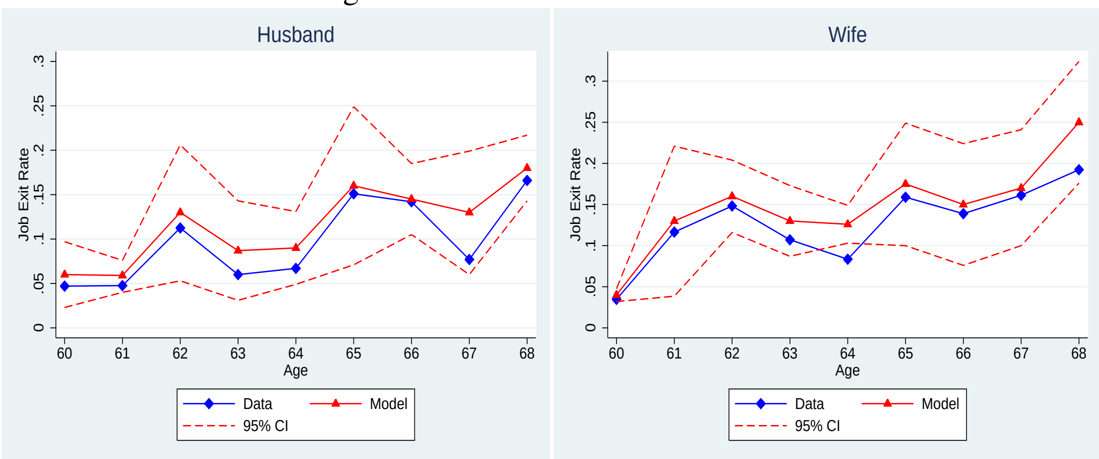

are identical. The Pearson Chi-square test statistic is

$$\sum_{c=1}^{4}\sum_{g=1}^{5}\frac{(O_{c,g}\times N_g - P_{c,g}\times N_g)^2}{P_{c,g}\times N_g}$$

where $O_{c,g}$ and $P_{c,g}$ represent the average observed frequency and the average predicted probability for choice c in stratum g, and N_g is the number of observations in stratum g. This Pearson Chi-square test statistic has an χ^2 distribution with $19(= cg - 1)$ degrees of freedom. For 19 degrees of freedom, the 5% critical value of the χ^2 distribution is 30.14 and the 1% critical value is 36.19. The chi-square statistic is 21.8465, which is well below these critical values. This indicates that there is little evidence that the model does not fit the data well.

6.2.3 Lagrange Multiplier Test

I use three Lagrange Multiplier (LM) tests to check whether the utility function is properly specified. First, I test the specification of individuals' preference for leisure. In my model, one spouse's preference for leisure is assumed to be an exponential function of age, health status, and the other spouse's leisure. However, other variables, such as the other spouse's health status, could also affect the preference for leisure due to caregiving. I assume spouse i's true preference for leisure in period t is

$$exp\{\beta^i X_t^i + Z_t^i \gamma^i\}$$

where $Z_t^i = H_{-i,t}$ is the other spouse's health, which is omitted in the model. I run an LM test to check whether or not parameters in $\gamma = (\gamma^m, \gamma^f)$ are jointly significantly different from zero. The null hypothesis is that $\gamma = 0$. The LM test statistic is $\chi^2(2) = 14.17$, and the P-value is 0.001. Thus, the null hypothesis can be rejected at the 1% significance level. This indicates that the omitted variable, spousal health, has a statistically significant effect on individuals' preference for leisure. This suggests the existence of another channel through which spousal retirement may be correlated, and it should be the subject of future research.

Second, I test whether one spouse cares about the other spouse's utility (also called caring preference). In the model, I assume that each spouse has egoistic preferences and does not value the other spouse's utility. To test caring preferences, I add the product of the two spouses' self-regarding utility functions to the original household utility function. The new household utility function is

$$U_{new} = U_{old} + \gamma_A (u_h \times u_w).$$

The lowercase u_h and u_w present the two spouses' self-regarding utility functions, as defined in the model section. The null hypothesis is $\gamma_A = 0$. The LM test statistic is $\chi^2(1) = 8.91$ and the P-value is 0.003. Thus, the null hypothesis can be rejected at the 1% significance level. This indicates that caring preferences might be more appropriate than egoistic preferences to describe married people's utility. This test result is consistent with the evidence found in Friedberg and Stern (2014).

Last, I test for heteroskedasticity of idiosyncratic errors. In the model, I assume that idiosyncratic utility shocks, ι_{dt}, are independent across households, time, and discrete choices. Specifically, I assume that idiosyncratic utility shocks are iid EV. In other words, the variance of the idiosyncratic shocks is assumed to be the same across different groups of households. However, this may not always be the case. For example, in healthy households, some spouses may retire from full-time work because they want to travel, while others may continue working full-time because they want to bring in more money. Thus, households in good health may have a greater dispersion in labor supply decisions than households in bad health. To test for heteroskedasticity, I add a vector of observed explanatory variables to the original scale parameter. The new scale parameter is

$$\sigma_{new} = \sigma_\iota + X_\sigma \gamma_\sigma$$

where σ_ι is the scale parameter under the homoskedasticity assumption. The variable X_σ is a vector of variables that might affect the variance of idiosyncratic shocks, which includes the two spouses' education level and health status. Then the null hypothesis is $\gamma_\sigma = 0$, which means that, among people who have different education levels and health statuses, there is no heteroskedasticity of idiosyncratic errors. The LM test statistic is $\chi^2(4) = 2.95$, and the P-value is 0.23. Therefore, since the null hypothesis cannot be rejected significantly, it is reasonable to assume the homoskedasticity of idiosyncratic errors.

7 Counterfactual Simulations

In this section, I use the model estimates to run several counterfactual simulation experiments for two purposes. First, I run several simulation experiments to identify the following causal inferences: (1) the causal effects of employer-provided retiree health insurance (EPRHI) on spouses' retirement (see 7.1.1); (2) the importance of including the health channel in evaluating the effects of health insurance on retirement (see 7.1.2); (3) the causal effects of spousal coverage on spouses' coordinated retirement (see 7.1.3); and (4) the causal effects of enjoying shared leisure time on household simultaneous retirement (see 7.1.4). Second, I conduct several policy simulations to predict the consequences of the following potential policy reforms: (1) the implementation of the Affordable Care Act (ACA) (see 7.2.1); (2) raising the Medicare eligibility age to 67 (see 7.2.2); and (3) raising the normal retirement age to 67 (see 7.2.3).

7.1 Causal Inferences

In the simulation experiments, I either pretend people are slightly different in their pension benefits or spousal coverage, or set some parameter estimates to zero, to see how spouses' retirement behaviors change.

7.1.1 The Effects of Employer-Provided Retiree Health Insurance

As shown in Figures 4.1 and 4.2 in data section 4.4, spouses who have retiree health insurance have lower full-time work participation rates, at every age, than those who have tied health insurance. This may be so for two reasons: (1) risk-averse workers value the effect of health insurance on reducing both the amount and the volatility of future out-of-pocket medical expenses; and (2) workers who have retiree insurance may have greater pension wealth than those with tied insurance. To isolate the effects of EPHI on labor supply from the effects of pension wealth, I assign the average pension accrual rates of workers who have tied health insurance to every spouse in my sample, so that pension incentives are the same across health insurance categories. Next, I conduct two simulation experiments. In the first experiment, I assume that all spouses have tied insurance; and in the second experiment, I assume that all spouses have retiree insurance. The difference in these two simulation experiments is a measure of the effects of EPHI on retirement, after controlling for pensions.

Figure 7.1 compares the job participation rates in the two simulation experiments. The two solid lines represent the model-predicted participation rates for spouses with tied insurance and those with retiree insurance, respectively. The two dotted lines represent the simulated participation rates for spouses in these two groups. At almost every age, spouses with tied insurance have higher simulated participation rates than those with retiree insurance. Even after the Medicare eligibility age (65), the simulated participation rates of husbands with tied insurance are still higher than the simulated participation rates of husbands with retiree insurance. This might be because some husbands with tied insurance delay their retirement to provide insurance coverage to their wife until their wife is eligible for Medicare. By contrast, the simulated participation rates of wives with tied insurance are close to the simulated participation rates of wives with retiree insurance even before 65. This might be because some wives are covered by spousal coverage, and thus, they would retire before 65 without losing insurance coverage.

Figure 7.1: Effects of EPRHI on Job Participation Rates

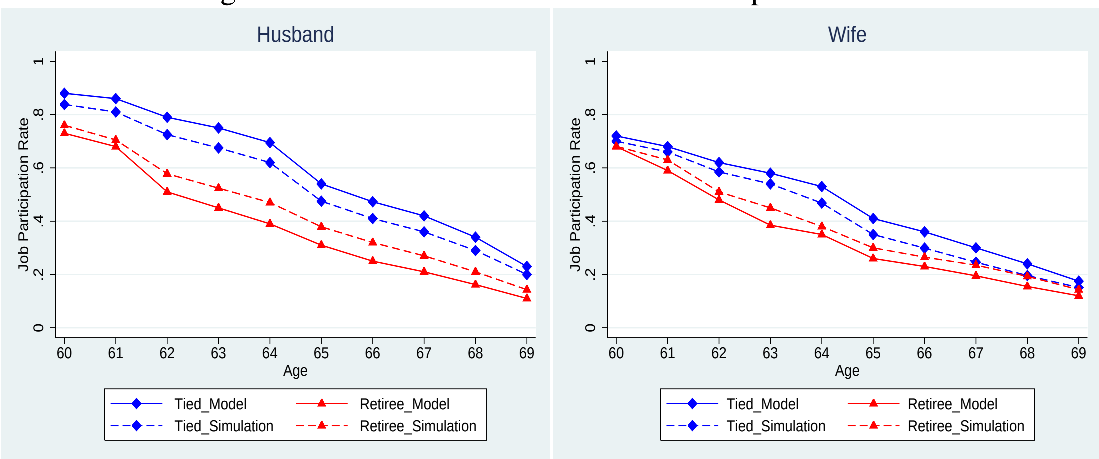

One measurement of the effects of EPRHI is the difference in the average retirement age, which is the sum of the differences in the two dotted lines for all ages between 60 and 69 (French and Jones (2011)). The results show that, giving employer-provided retiree insurance to spouses with tied insurance, the average retirement age decreases by 1.1 and 0.5 years for husbands and wives, respectively. Comparatively, the sum of the differences in the model-predicted participation rates in two insurance categories (retiree insurance and tied insurance) is 2.3 and 1.2 years for husbands and wives, respectively. Thus, the low job participation rates of workers who have retiree coverage are due partly to more generous pensions. However, even after controlling for pensions, health insurance remains an important factor that explains more than 40% of the difference in the job participation rates for spouses in the two insurance categories.

In the literature, another measurement of the effects of EPHI used is the difference in the age-62 job exit rates (Rust and Phelan (1997)). Figure 7.2 compares the job exit rates in the two simulation experiments. The difference at age 62 means that, giving retiree coverage to spouses with tied insurance increases the age-62 job exit rate by 9 and 8 percentage points for husbands and wives, respectively.[79]

Figure 7.2: Effects of EPRHI on Job Exit Rates

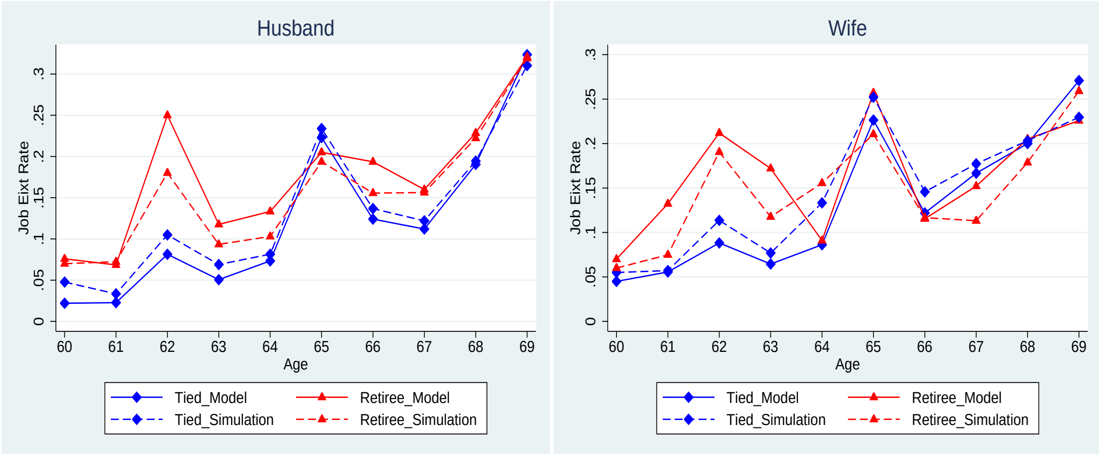

This paper finds larger effects of health insurance on retirement than previous structural studies that include risks related to medical expenses.[80] The reason why I find larger effects of health insurance on retirement may be because my model accounts for the interdependence of both spouses' insurance coverage, which has been overlooked in the previous studies. My estimates of the effects of health insurance lie within the range of estimates that have been provided by previous reduced-form studies.[81]

7.1.2 The Importance of Including the Health Channel

Recall that this paper models two channels through which health insurance affects household retirement decisions: medical expense and health. The medical expense channel is the only one that has been considered in the literature on health insurance and retirement. To evaluate the importance of including the health channel in my model, I run several simulation experiments to answer two research questions. First, without including the health channel, how well does my model explain the observed labor market behaviors? Second,

[79]If all spouses are eligible for retiree insurance rather than tied insurance, the husbands' age-62 job exit rate would increase from 9 to 18 percentage points, and the wives' age-62 job exit rate would increase from 8 to 16 percentage points.

[80]For example, Rust and Phelan (1997) find that being eligible for employer-provided retiree insurance increases men's age-62 job exit rate by about 8 percentage points. French and Jones (2011) find that retiree coverage reduces men's average retirement age by 0.34 year. Blau and Gilleskie (2006, 2008) find that retiree coverage reduces married men's average participation rates by 1.7 percentage points. By contrast, I find that retiree coverage increases husbands' age-62 exit rate by 9 percentage points, accelerates husbands' retirement by 1.1 years, and reduces husbands' average participation rates by 11 percentage points.

[81]For example, Madrian et al. (1994) find that retiree coverage reduces men's retirement age by 0.4 – 1.2 years.

does EPRHI raise job exit rates primarily through the medical expense channel or the health channel?

To answer the first question, I simulate household labor supply by taking out the health channel, which appears in health transition functions. I set the parameters of health status in health transition functions to zero. Figure 7.3 compares observed job participation rates in the data, model-predicted participation rates with both channels, and simulated participation rates without the health channel. Figure 7.4 compares observed, model-predicted, and simulated job exit rates. For both husbands and wives, without the health channel, the model would overpredict the job participation rates and underpredict the job exit rates, at almost every age. This is because, without the health channel, spouses would expect worse health in the model. This would affect their retirement decisions in two directions: (1) spouses with worse health have higher preference for leisure relative to work, and thus, they would retire earlier; and (2) spouses with worse health would expect higher and more volatile medical expenses, and thus, they would delay their retirement to have better financial protection. The simulation results show that the second direction outweighs the first one.

Figure 7.3: Job Participation Rates without the Health Channel

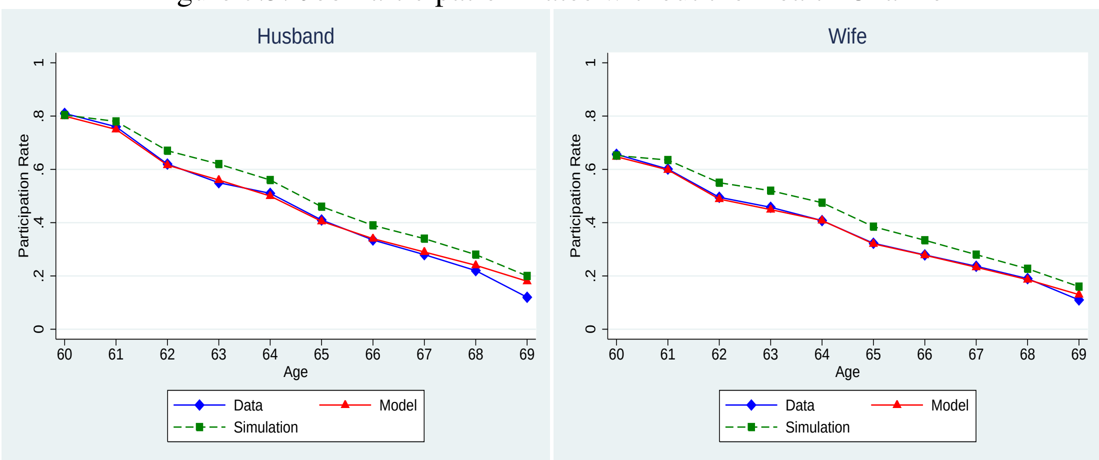

To answer the second question (does EPRHI raise job exit rates primarily through the medical expense channel or the health channel?), I conduct two simulation experiments to separate the effects of health insurance through the medical expense channel from the effects through the health channel. In this subsection (7.1.2), I rerun the two simulations from subsection 7.1.1 but take out the health channel. The difference in these two simulation experiments is a measure of the effects of EPRHI on retirement without the health channel. Figure 7.5 compares effects of EPRHI on job participation rates with and without the health channel. Without the health channel, although the effects of EPHI on participation rates becomes smaller at almost every age, the model still captures most of the effects. For example, without the health channel, EPHI coverage accelerates husbands' retirement by 0.9 years, which accounts for about 82% of the effects of EPHI when both channels

Figure 7.4: Job Exit Rates without the Health Channel

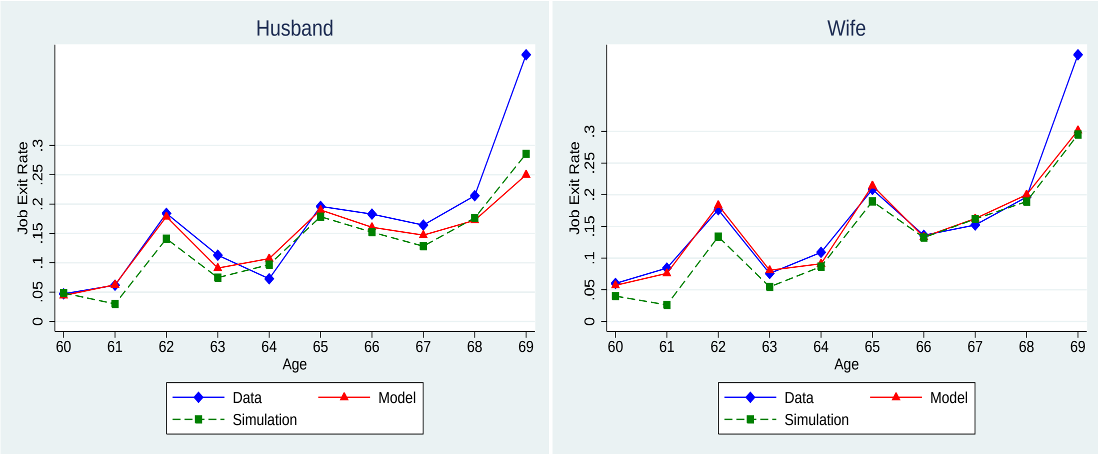

are considered. Thus, the health channel accounts for only 18% of the effects of EPRHI.[82] In summary, most of the value that workers place on health insurance comes through the medical expense channel, which reduces both the amount and the volatility of medical expenses.

Figure 7.5: Effects of EPRHI on Participation Rates without the Health Channel

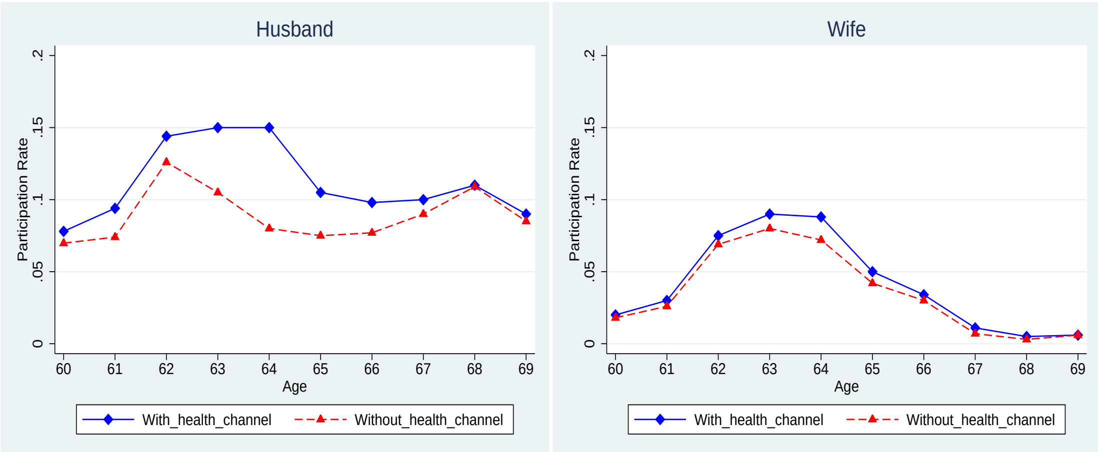

Figure 7.6 compares effects of EPRHI on job exit rates with and without the health channel. Without the health channel, the effects of EPHI on job exit rates become smaller at almost every age before 65, while the effects become bigger after 65. This might be because, without the health channel, spouses in my model would expect worse health. Therefore, before they are eligible for Medicare, worse health makes them less likely to retire, in order to keep the employer-provided insurance. After they are eligible for Medicare, worse health makes them more likely to retire in order to have more leisure without losing insurance coverage.

[82] Similarly, without the health channel, EPRHI coverage accelerates wives' retirement by 0.4 years, which accounts for about 80% of the effects of EPRHI when both channels are considered.

Figure 7.6: Effects of EPRHI on Exit Rates without the Health Channel

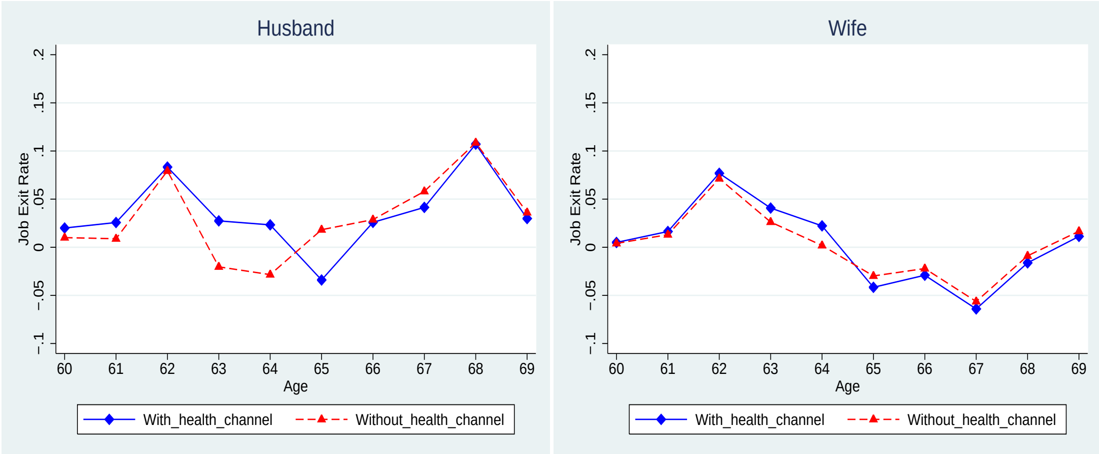

7.1.3 Spousal Coverage and Household Retirement

This paper includes spousal coverage as a factor that motivates spouses' coordinated retirement. To evaluate the causal effects of spousal coverage on couples' coordinated retirement, I rerun the two simulations from subsection 7.1.1 but take out spousal coverage.[83] Figure 7.7 compares the distributions of retirement age for husbands with tied insurance under two circumstances: with and without spousal coverage. The comparison shows that, with spousal coverage, the proportion of husbands who retire around age 65 (between 64.5 and 65.5) decreases from 65% to 20%. The decrease comes from two sources. This is because, if spousal coverage is available, a small group of husbands choose to accelerate their retirement before 65, and a large group of husbands choose to delay their retirement after 65 when spousal coverage exists.[84] The variation in husbands' responses to spousal coverage captures the heterogeneity in plan characteristics and individual preference for leisure included in my model. For example, with spousal coverage, many husbands choose to delay their retirement to provide insurance to their younger wife because husbands' employer-provided plans have better quality, while some husbands choose to switch to their wife's employer-provided plan and retire earlier, either because their wife's employer provides a better plan or because they have a high preference for leisure. Similarly, the proportion of wives who retire around age 65 decreases from 56% to 17% when spousal coverage exists. The decrease is due to the fact that most wives choose to retire earlier (before 65) when spousal coverage exists.[85] Although spousal coverage accelerates most wives' retirement, the upper-right part of the distribution functions show that a very small group of wives choose to delay their retirement.

Figure 7.8 compares the distributions of the time interval (years) between husbands'

[83] In subsection 7.1.1, I assume every spouse has spousal coverage in both simulation experiments.

[84] Figure 7.7 shows that spousal coverage increases the proportion of husbands who retire before 65 from 17% to 20%, and it increases the proportion of husbands who retire after 65 from 18% to 60%.

[85] Figure 7.7 shows that spousal coverage increases the proportion of wives who retire before 65 from 32% to 75%, and it decreases the proportion of wives who retire after 65 from 11% to 8%.

Figure 7.7: Distribution of Retirement Age

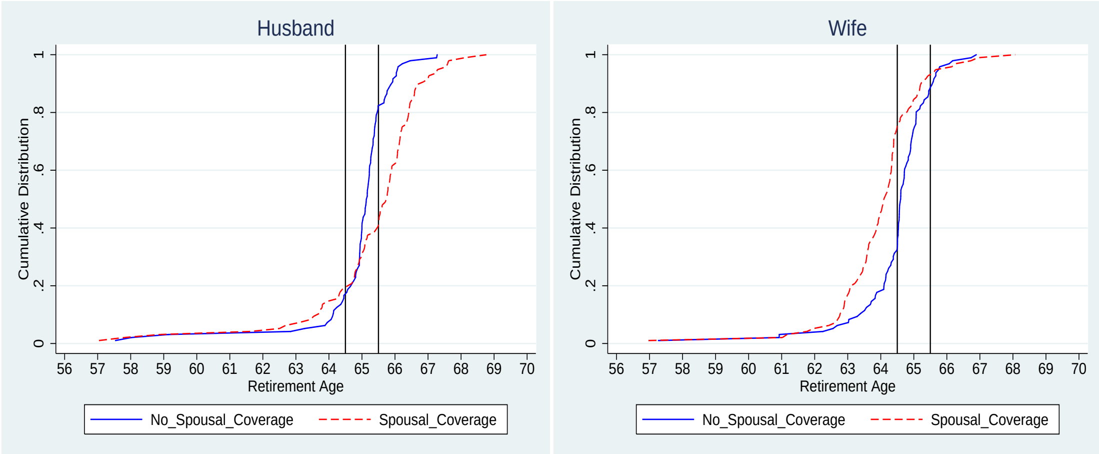

and wives' retirement for households with tied insurance. For most households, the time interval between the husband's and the wife's retirement decreases when spousal coverage exists. For example, for some of these households, the husbands retire more than two years earlier than their wife when there is no spousal coverage, whereas, when spousal coverage exists, these husbands retire within two years either before or after their wife retires. This increases the proportion of households that choose simultaneous retirement (husband and wife retire within one or two years of each other) from 18% to 33%. Recall that, for many households, spousal coverage delays husbands' retirement and accelerates wives' retirement. Additionally, since the husband is usually a few years older than the wife, pushing their retirement age apart further actually reduces the years between their retirement. This makes households more likely to choose simultaneous retirement.

Figure 7.8: Distribution of the Difference in Retirement Year

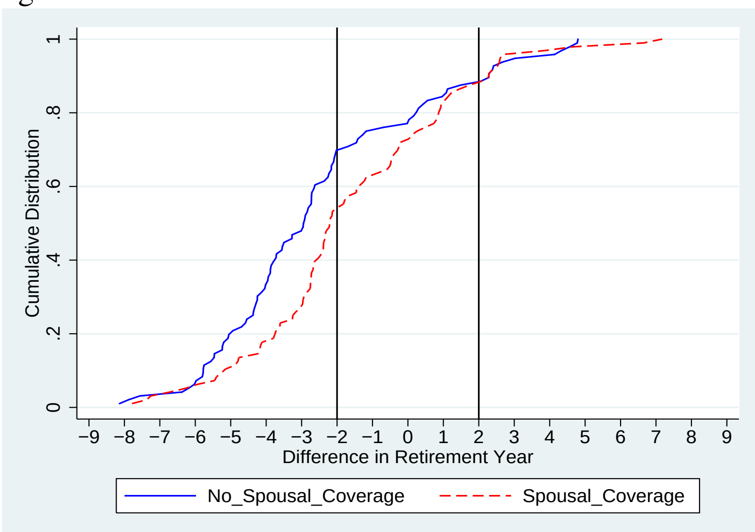

In summary, spousal coverage motivates simultaneous retirement by delaying husbands' retirement and accelerating wives' retirement. To evaluate the effects of spousal coverage on simultaneous retirement for my sample, I take out spousal coverage from spouses whose employer provides it. Then, I simulate the propensity of husbands' and wives' to retire in the same period (simultaneous retirement rate). The results show that taking out spousal coverage decreases the simultaneous retirement rate by 7 percentage points (from 0.29 to 0.22), which accounts for 24% of the model-predicted simultaneous retirement rate.

7.1.4 Complementarity in Leisure and Simultaneous Retirement

Complementarity in leisure (spouses' valuation of spending leisure together) could be one of the factors leading to simultaneous retirement in a household. To evaluate the importance of including spousal leisure in modeling household retirement behavior, I set the parameter of spousal leisure in an individual utility function equal to zero, and then I simulate the simultaneous retirement rate. The results show that omitting complementarity in leisure decreases the simultaneous retirement rate by 10 percentage points (from 0.29 to 0.19), which explains 34% of the model-predicted simultaneous retirement rate. My estimated effect of spousal leisure on simultaneous retirement is slightly smaller than the effect found in Gustman and Steinmeier (2004), while my estimated effect is much bigger than the effect found in Casanova (2010).[86]

7.2 Policy Simulations

In this subsection, I use the model to predict how the Patient Protection and Affordable Care Act (PPACA, or, in short, ACA) and the changes in Medicare and Social Security rules would affect retirement behavior.

7.2.1 The Implementation of the ACA

A key component of the ACA, which was signed into law on March 23, 2010, is the Health Insurance Marketplace. The Health Insurance Marketplace, also called the Health Insurance Exchange, "is the place where people without health care insurance can find information about health insurance options and also purchase health care insurance.[87] The law requires people without health insurance to purchase insurance from the Marketplace. Thus, the ACA helps to make health insurance independent of employment status. Four plans are available in the Marketplace for people who are facing retirement: Bronze, Silver, Gold, and Platinum.[88] To simplify the computation of the simulation, I use the Silver plan as the

[86] Gustman and Steinmeier (2004) find that omitting complementarity in leisure decreases the simultaneous retirement rate by 11 percentage points. Casanova (2010) finds that omitting complementarity in leisure decreases the simultaneous retirement rate by 3.8 percentage points.

[87] Information can also be found regarding eligibility for help with paying premiums and reducing out-of-pocket costs." https://www.irs.gov/affordable-care-act

[88] The Bronze plan has the lowest premium and coinsurance rate, while the Platinum plan has the highest premium and coinsurance rate.

representative plan in the Marketplace because it is the most popular insurance option by far.[89]

Because the implementation of the ACA makes health insurance independent of employment status, the ACA reduces the degree of the job-lock generated by employer-provided insurance for workers with tied insurance. To evaluate the effects of the ACA on labor supply of spouses with tied insurance, I assume that every spouse in my sample is eligible for tied health insurance, and I give everyone the average pension accrual rates of workers who have tied health insurance. Then I simulate spouses' labor supply in two simulation experiments: with the implementation of the ACA and without it. Figures 7.9 and 7.10 compare the simulated job participation rates and job exit rates, respectively, in these two simulation experiments. The figures show that, for spouses with tied insurance, the ACA decreases participation rates and increases exit rates, mostly between age 61 and 65. Note that the effects of the ACA on participation rates decrease significantly after 65. This might be because spouses are eligible for Medicare from age 65 onward and Medicare provides better plans than the ACA does. Thus, the ACA is not very attractive after spouses are eligible for Medicare. The sum of the differences in participation rates shows that, over a 10-year period (for ages 60-69), the ACA accelerates retirement by 0.5 and 0.3 years for husbands and wives with tied insurance, respectively.

Figure 7.9: Effects of the ACA on Job Participation Rate: Tied Insurance

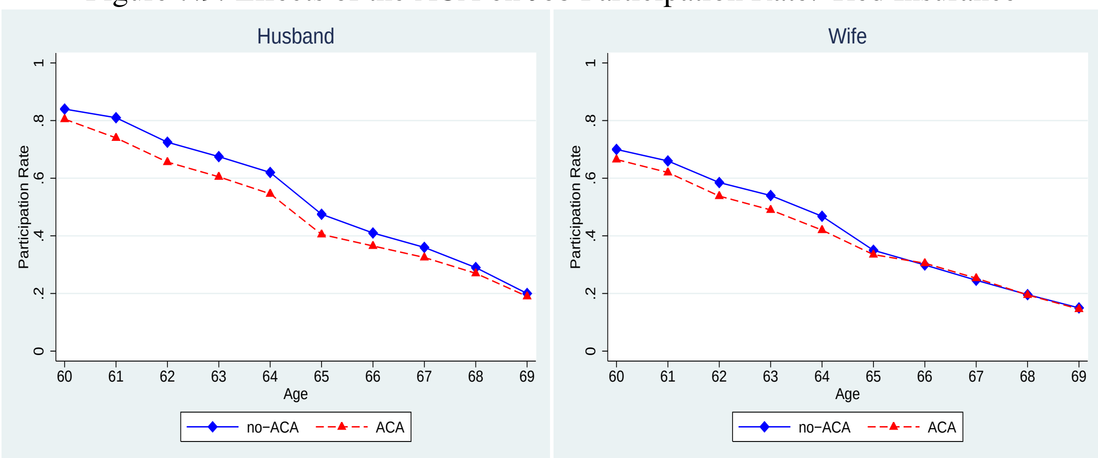

To predict the consequences of recent health care reform, I also simulate labor supply for spouses in my sample assuming the implementation of the ACA. Over a 10-year period, implementing the ACA accelerates retirement by around 0.12 and 0.09 years for husbands and wives, respectively.

Two features of my model enable me to predict the effects of the implementation of the ACA. First, my model differentiates insurance plans by their quality (plan characteristics). Second, my model allows spouses to choose an insurance plan from available

[89] In 2015, about 70 percent of consumers enrolled in Silver plans. Source: "2016 Marketplace Affordability Snapshot," CMS.gov

Figure 7.10: Effects of the ACA on Job Exit Rate: Tied Insurance

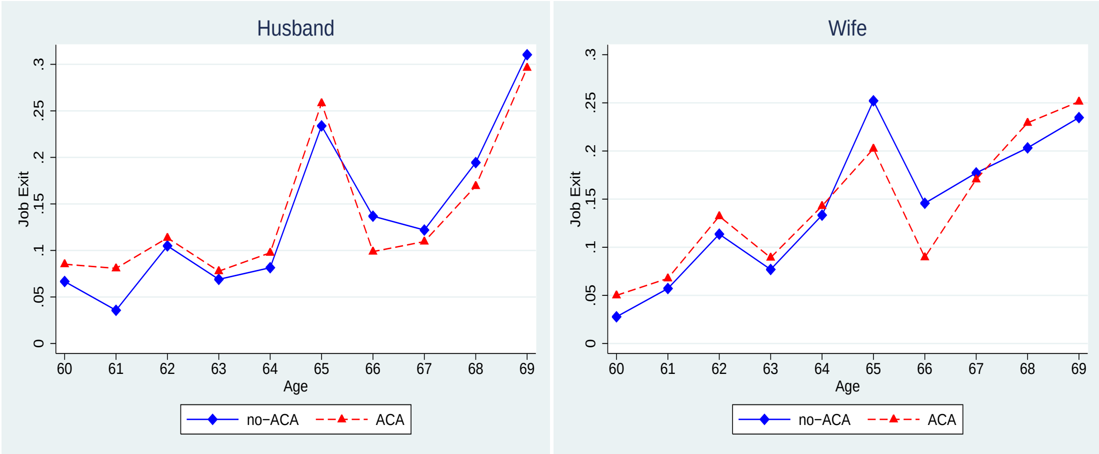

plans by comparing their qualities. Previous papers about health insurance and retirement do not consider the heterogeneity in plan qualities (e.g., Madrian et al. (1994) and Gruber and Madrian (1995)). Thus, these papers would largely overpredict the effects of the ACA on retirement because the ACA plan quality is much worse than the quality of employer-provided plans.

My model may still overpredicts the effects of the ACA because my model excludes the choice of being uninsured. I assume that workers always choose to be covered by health insurance when they are eligible.[90] However, even after the implementation of the ACA, a large amount of people still choose to be uninsured. Without capturing the possibility of choosing to be uninsured, my model may overpredict the effect of the ACA on.

7.2.2 Raising the Medicare Eligibility Age to 67

Medicare is a public health insurance program that is available to, in general, all persons age 65 and older.[91] Thus, Medicare mitigates the degree of job-lock generated by employer-provided insurance for workers with tied insurance. The Medicare eligibility age is not scheduled to rise, but there are proposals to raise it to 67. To evaluate the effects of raising the Medicare eligibility age on labor supply of spouses with tied insurance, I assume that every spouse in my sample has tied insurance and has the average pension accrual rates of workers who have tied health insurance. Then I simulate spouses' labor supply with the Medicare eligibility age set at 65 and 67, respectively. Figures 7.11 and 7.12 compare the simulated job participation rates and job exit rates, respectively, in these two simulation experiments. For spouses who have tied insurance, raising the Medicare eligibility age to 67 increases the job participation rates and decreases the job exit rates, mostly between age

[90] I make this assumption for two reasons: (1) before the implementation of the ACA, there was no private health insurance available, other than EPHI, to spouses in my sample; and (2) no spouse in my sample chooses to be uninsured when he is eligible for some EPHI.

[91] People under 65 and receive Social Security Disability Insurance (SSDI) benefits is eligible for Medicare. Source: https://www.ssa.gov

64 and 67.[92] In addition, the spike in exit rates moves from age 65 to 67 as the Medicare eligibility age increases from 65 to 67.

Figure 7.11: Effects of Medicare on Participation Rates: Tied Insurance

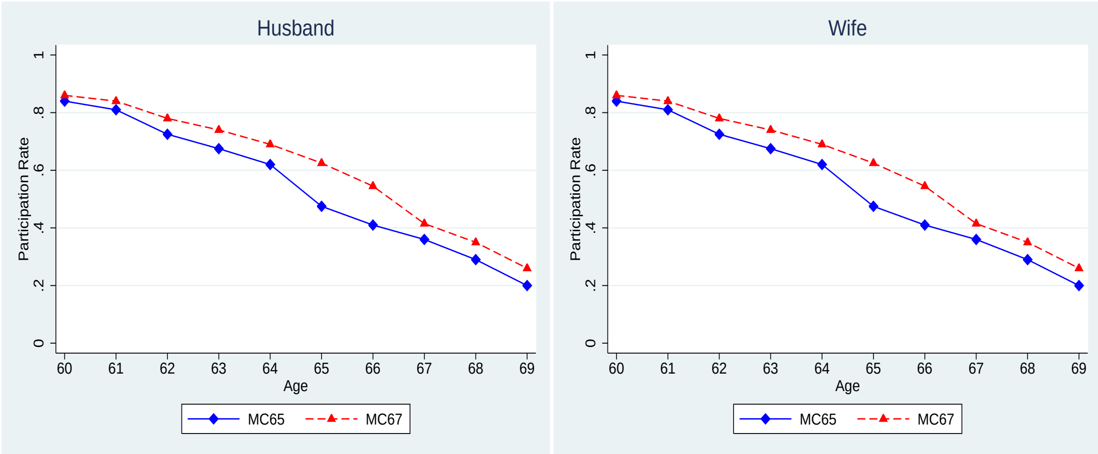

Figure 7.12: Effects of Medicare on Exit Rates: Tied Insurance

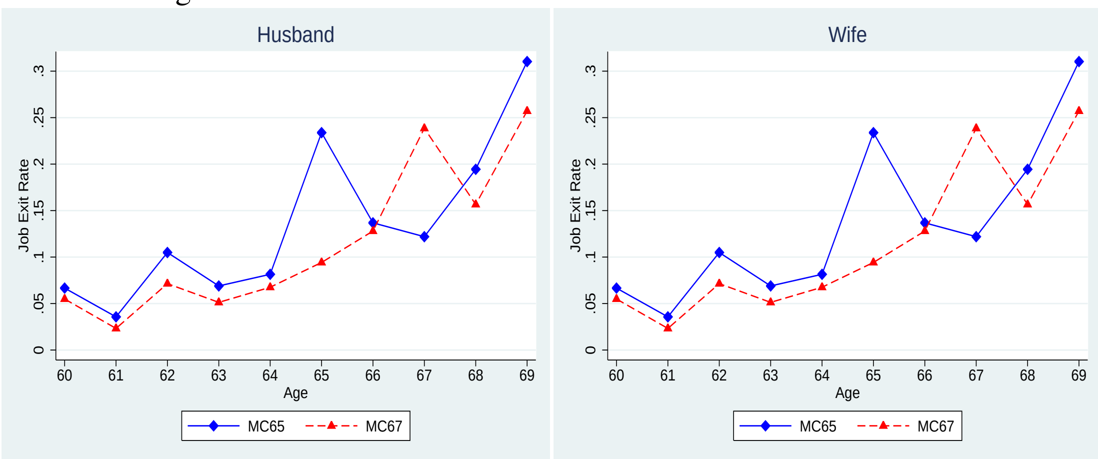

Over a 10-year period (for ages 60-69), raising the Medicare eligibility age to 67 delays retirement by 0.8 and 0.4 years for husbands and wives with tied insurance, respectively. Recall that EPRHI delays retirement by 1.1 years for husbands and 0.5 years for wives. The implementation of the ACA accelerates retirement by 0.5 years for husbands and 0.3 years for wives. Therefore, the effects of Medicare (or raising the Medicare eligibility age to 67) are bigger than the effects of the ACA but are smaller than the effects of EPRHI. We can, perhaps, attribute these differences to plan quality: the employer-provided insurance plan has the best plan characteristics (high co-insurance rate and low paid premium and deductible), and the plans in the Marketplace have the worst plan characteristics.

To predict the consequences of raising the Medicare eligibility age from 65 to 67, I simulate labor supply for spouses in my sample assuming that the Medicare eligibility age is

[92]After 67, raising the Medicare eligibility age to 67 still increases husbands' participation rates. This might be because some husbands delay their retirement to provide insurance to their wife.

67. Over a 10-year period, raising the age to 67 delays retirement by 0.17 years for husbands and 0.12 years wives. These amounts are larger than the effects found in the literature. For example, Rust and Phelan (1997), Blau and Gilleskie (2006), and French and Jones (2011) find that increasing the Medicare eligibility age to 67 delays men's retirement by 0.13, 0.01 and 0.074 years, respectively. The reason why I find a larger effect of Medicare is that my model accounts for the interdependence of both spouses' insurance.[93]

7.2.3 Raising the Normal Retirement Age to 67

Finally, I conduct simulation experiments to predict how the normal retirement age (NRA) affects labor supply. The normal retirement age is scheduled to increase from 65 to 67 by 2022, and I simulate labor supply with the normal retirement age set at age 67. Figures 7.13 and 7.14 shows how spouses' labor supply change if the normal retirement age increases from 65 to 67.

Figure 7.13: Effects of the Normal Retirement Age on Participation Rates: Sample

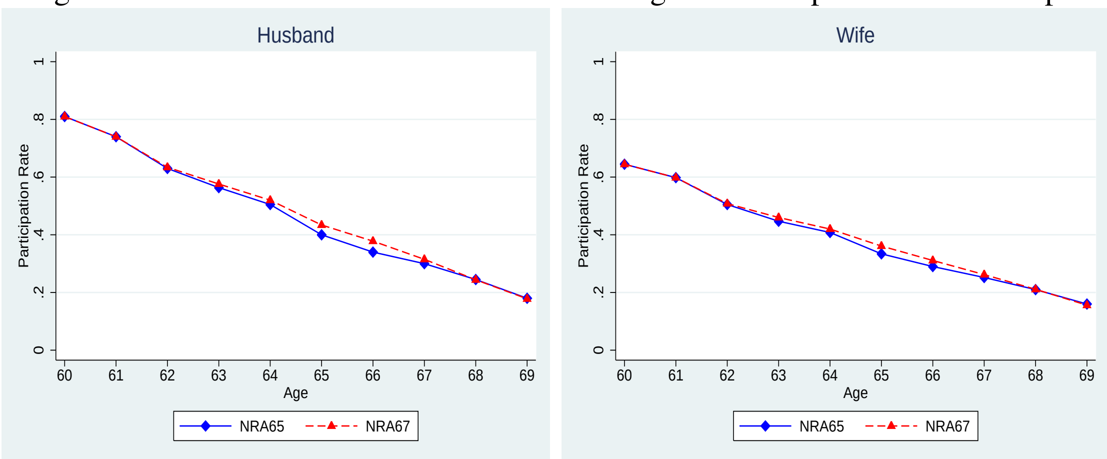

Figure 7.14: Effects of the Normal Retirement Age on Exit Rates: Sample

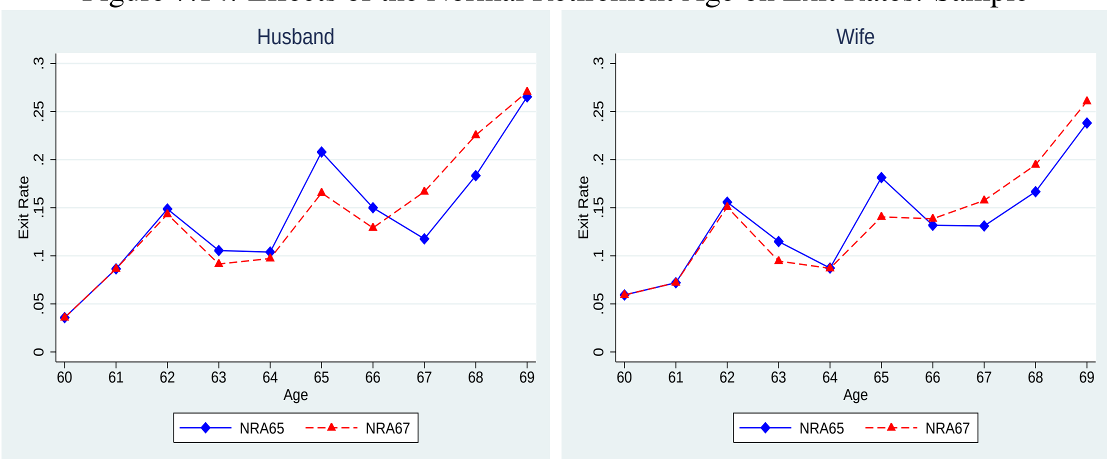

[93]This is the same reason why my model finds big effects of EPRHI on retirement (see 7.1.1).

Raising the normal retirement age to 67 increases the job participation rates from age 62 onward, mostly between age 62 and 67. In total, over a ten-year period, this delays retirement by 0.11 years for husbands and 0.08 years wives. These numbers are similar to (though slightly larger than) the effects found by French and Jones (2011), but they are smaller than the effects documented by Rust and Phelan (1997).[94] The reason why Rust and Phelan (1997) find a bigger effect may be because their sample consists of low-income males, whose Since Social Security benefits are the main source of their retirement income. Note that, after increasing the normal retirement age from 65 to 67, the spikes at age 62 and 65 still persist, but the height of the spike at age 65 drops.

8 Conclusion

The absence of health insurance coverage could be one of the biggest obstacles to retiring before age 65. The link between health insurance and retirement is especially complicated in a household in which two spouses coordinate their retirement decisions–both because they share economic resources and because spouses' health insurance coverages are interdependent. To examine how health insurance affects household joint retirement decisions, I develop a dynamic programming model of household retirement in which married couples jointly decide when to retire, how to use available insurance, and how much to save.

I estimate my model with Maximum Simulated Likelihood estimation using data from the Health and Retirement Study and the Medical Expenditure Panel Survey. The estimates of preference parameters are reasonable and the model fits the data well. I find that moving workers from retiree to tied health insurance coverage delays retirement by 1.1 years for husbands and 0.5 years for wives. In decomposing the employment response to EPHI coverage, I find that over 80% of the response reflects the valuation of the consumption smoothing effects of health insurance, and less than 20% reflects the valuation of the health improvement effects. Furthermore, I find that spousal coverage motivates simultaneous retirement by delaying husbands' retirement and accelerating wives' retirement, and it explains about 24% of observed simultaneous retirement. Lastly, I find that husbands and wives enjoy spending leisure time together, which explains 34% of observed simultaneous retirement.

My policy simulations show that health insurance-related policies would have big effects on retirement. The implementation of the ACA is predicted to accelerate retirement by 0.12 years for husbands and 0.09 years for wives. Raising the Medicare eligibility age from 65 to 67 is predicted to delay retirement by 0.17 years for husbands and 0.12 years for wives. Comparatively, increasing the Social Security full retirement age from 65 to 67 delays retirement by 0.11 years for husbands and 0.08 years for wives.

[94] French and Jones (2011) find that raising the normal retirement age to 67 increases years of work by 0.08 years. Rust and Phelan (1997) find a larger effect of raising the normal retirement age.

References

BANKS, J., R. BLUNDELL AND M. CASANOVA, "The dynamics of retirement behavior in couples: Reduced-form evidence from England and the US," *mimeo* (2010).

BENITEZ-SILVA, H., M. BUCHINSKY AND J. RUST, "How large are the classification errors in the social security disability award . . . ," *NBER* (2004).

BLAU, D. M., "Labor Force Dynamics of Older Married Couples," *Journal of Labor Economicsr* 62 (1998), 117–156.

BLAU, D. M. AND D. B. GILLESKIE, "Health insurance and retirement of married couples," *Journal of Applied Econometrics* 21 (2006), 935–953.

———, "The role of retiree health insurance in the employment behavior of older men," *International Economic Review* 49 (2008), 475–514.

BÖRSCH-SUPAN, A. AND V. A. HAJIVASSILIOU, "Smooth unbiased multivariate probability simulators for maximum likelihood estimation of limited dependent variable models," *Journal of Econometrics* 58 (1993), 347–368.

BRIEN, M. J., L. A. LILLARD AND S. STERN, "Cohabitation, marriage, and divorce in a model of match quality," *International Economic Review* 47 (2006), 451–494.

BUNTIN, M. B. AND A. M. ZASLAVSKY, "Too much ado about two-part models and transformation? Comparing methods of modeling Medicare expenditures," *Journal of Health Economics* 23 (2004), 525–542.

CAMERON, S. V. AND J. J. HECKMAN, "Life cycle schooling and dynamic selection bias: Models and evidence for five cohorts," Technical Report, National Bureau of Economic Research, 1998.

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

CHETTY, R., "A new method of estimating risk aversion," *American Economic Review* 96 (2006), 1821–1834.

CLARK, R. L., "Portability of benefits, job changes, and the role of government policies," *Taskforce on Reconstructing America's Labor Market Institutions. MITInstitute for Work and Employment Research, Cambridge, MA* (1999).

COILE, C. C., "Health shocks and couples' labor supply decisions," *National Bureau of Economic Research* (2004).

CURRIE, J. AND J. GRUBER, "Health Insurance Eligibility, Utilization of Medical Care, and Child Health," *The Quarterly Journal of Economics* 111 (1996), 431–466.

CUTLER, D. M., "Employee Costs and the Decline in Health Insurance Coverage," *Forum for Health Economics & Policy* 6 (2003).

DE NARDI, M., E. FRENCH AND J. B. JONES, "Why Do the Elderly Save? The Role of Medical Expenses," *Journal of Political Economy* 118 (2010), 39–75.

DE NARDI, M. C., "Wealth inequality and intergenerational links," *Review of Economic Studies* 71 (2004), 743–768.

ELBERS, C. AND G. RIDDER, "True and Spurious Duration Dependence: The Identifiability of the Proportional Hazard Model," *The Review of Economic Studies* 49 (1982), 403–409.

ELBNER, C. AND M. S. MARQUIS, "Employers' health insurance cost burden, 1996-2005," *Monthly Lab. Rev.* 131 (2008), 28.

FIELDS, G. S. AND O. S. MITCHELL, "The effects of social security reforms on retirement ages and retirement incomes," *Journal of Public Economics* 25 (1984), 143–159.

FREES, E. W., J. GAO AND M. A. ROSENBERG, "Predicting the frequency and amount of health care expenditures," *North American Actuarial Journal* 15 (2011), 377–392.

FRENCH, E. AND J. B. JONES, "The Effects of Health Insurance and Self-Insurance on Retirement Behavior," *Econometrica* 79 (2011), 693–732.

FRIEDBERG, L. AND S. STERN, "Marriage, divorce, and asymmetric information," *International Economic Review* 55 (2014), 1155–1199.

FRIEDBERG, L. AND A. WEBB, "Determinants and consequences of bargaining power in households," *NBER Working Paper Series* Working Pa (2006), 1–43.

GRUBER, J. AND B. C. MADRIAN, "Health-Insurance Availability and the Retirement Decision," *American Economic Review* 85 (1995), 938–948.

GUO, D., "Appendix to the job market paper–Happy Together or Home Alone: A Structural Model of The Role of Health Insurance in Household Joint Retirement," (2016), here is the link: http://dinaguo.weebly.com/.

GUSTMAN, A. AND T. STEINMEIER, "Retirement and the stock market bubble," *NBER working paper* (2002).

GUSTMAN, A. L., O. S. MITCHELL AND T. L. STEINMEIER, “The Role of Pensions in the Labor Market: A Survey of the Literature,” *Industrial and Labor Relations Review* 47 (1994), 417–438.

GUSTMAN, A. L. AND T. STEINMEIER, “Integrating retirement models,” , 2009.

GUSTMAN, A. L. AND T. L. STEINMEIER, “Employer-Provided Health-Insurance and Retirement Behavior,” *Industrial & Labor Relations Review* 48 (1994), 124–140.

———, “Retirement in Dual-Career Families: A Structural Model,” *Journal of Labor Economics* 18 (2000), 503–545.

———, “Social security, pensions and retirement behaviour within the family,” *Journal of Applied Econometrics* 19 (2004), 723–737.

———, “Integrating Retirement Models: Understanding Household Retirement Decisions,” in *Factors Affecting Worker Well-being: The Impact of Change in the Labor Market* (2014), 79–112.

GUVENEN, F., “An empirical investigation of labor income processes,” *Review of Economic Dynamics* 12 (2009), 58–79.

HAJIVASSILIOU, V., *Some practical issues in maximum simulated likelihood* (Cambridge University Press, 2000).

HECKMAN, J. J. AND B. SINGER, “Social science duration analysis,” *Longitudinal analysis of labor market data* (1985), 39–110.

HOCHGUERTEL, S., “Self-employment around retirement age,” in *Entrepreneurship, Self-Employment and Retirement* (Springer, 2015), 209–258.

HUBBARD, R. G., J. SKINNER AND S. P. ZELDES, “Precautionary Saving and Social Insurance,” *Journal of Political Economy* 103 (1995), 360.

HURD, M., “The Joint Retirement Decision Of Husbands And Wives,” *Issues In The Economics of Aging* (1990), 231–258.

KAPUR, K. AND J. ROGOWSKI, “The role of health insurance in joint retirement among married couples,” *Industrial and Labor Relations Review* 60 (2007), 397–407.

KEANE, M. P. AND K. I. WOLPIN, “The Career Decisions of Young Men,” *Journal of Political Economy* 105 (1997), 473–522.

LEVY, H. AND D. MELTZER, “The impact of health insurance on health.,” *Annual review of public health* 29 (2008), 399–409.

LUMSDAINE, R. L., J. H. STOCK AND D. A. WISE, "Pension plan provisions and retirement: men and women, medicare, and models," in *Studies in the Economics of Aging* (University of Chicago Press, 1994), 183–222.

———, "Retirement incentives: The interaction between employer-provided pensions, Social Security, and retiree health benefits," in *The economic effects of aging in the United States and Japan* (University of Chicago Press, 1996), 261–293.

MADRIAN, B. C., G. BURTLESS AND J. GRUBER, "The effect of health insurance on retirement," *Brookings Papers on Economic Activity* 1994 (1994), 181–252.

MICHAUD, P., "Joint labour supply dynamics of older couples," *IZA Discussion paper Nr. 832* (2003).

MROZ, T. A., "Discrete factor approximations in simultaneous equation models: Estimating the impact of a dummy endogenous variable on a continuous outcome," *journal of Econometrics* 92 (1999), 233–274.

POHLMEIER, W. AND V. ULRICH, "An Econometric Model of the Two-Part Decisionmaking Process in the Demand for Health Care," *The Journal of Human Resources* 30 (1995), 339–361.

RUST, J., "Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher," *Econometrica: Journal of the Econometric Society* (1987), 999–1033.

———, "Maximum Likelihood Estimation of Discrete Control Processes," *SIAM Journal on Control and Optimization* 26 (1988), 19.

RUST, J. AND C. PHELAN, "How Social Security and Medicare Affect Retirement Behavior In a World of Incomplete Markets," *Econometrica* 65 (1997), 781–831.

VAN DER KLAAUW, W. AND K. I. WOLPIN, "Social Security and the Retirement and Savings Behavior of Low Income Households.," *Journal of econometrics* 145 (2008), 21–42.

ZISSIMOPOULOS, J., L. KAROLY AND N. MAESTAS, "Retirement Transitions of the Self-employed in the United States and England," *MRRC WP2007-155. Mercer Select* (2007).