The Effect of Social Security Auxiliary Spouse and Survivor's Benefits on the Household Retirement Decision

David Knapp
University of Michigan

1 Introduction

In 2011, 12.9 million age-qualifying Americans received \$112 billion in spouse and survivor's benefits from Social Security based on their husband or wife's earnings history. The Spouse's Benefit alone, while representing less than 4% of annual Social Security old-age expenditures, amounts to \$24 billion, which is larger than the individual 2012 budgets of 27 states, Canada's total military expenditures (\$22.5b, 2013), and the entire Federal budget for assistance to families with dependent children (TANF - \$17.6b, 2012).[1] Initially called the "wife's benefit", these benefits were introduced in 1939 when only 15% of households had two earners, compared to over 72% for households retiring after 1992.[2] No study has examined the effect of both the Spouse and Survivor's Benefits on household retirement behavior because of the complexity associated with estimating a structural model of interconnected household decisions. This study answers the question: *how responsive are husbands' and wives' retirement decisions to Spouse and Survivor's Benefits?*

This paper builds on the growing structural life-cycle retirement literature, which captures the dynamic interplay in people's choices, to model the household's decisions regarding savings, labor supply, and benefit claiming. I model the complex Social Security rules that reward and penalize spousal work choices, and allow them to interact with other key determinants of the household problem including household savings, private pension plans, and uncertain health, mortality, and medical expenses. I conduct counterfactual experiments that show households respond sharply to changes in the Survivor's Benefit, but little to changes in the Spouse's Benefit. Reducing both benefits between 50% and 100% cause women to work 0.47 to 1.27 years longer. The effect is nonlinear for men: increasing work by 0.29 years when both benefits are reduced by half, but decreasing work by 0.53 years when they are eliminated. This result suggests the annuity provided by the Survivor's Benefit, even if reduced, creates a strong incentive for the couple's high earner to continue working. Finally, I find nonlinear savings to Social Security from reducing Spouse and Survivor's Benefits amongst the married, non-disabled population in my sample: when these benefits are reduced by half, it achieves 74.1% of the savings from eliminating these benefits. The model demonstrates these nonlinear savings arise primarily due to the structure of Social Security benefits, not from changes in labor supply.

---

[1] Social Security figures are derived from SSA (2012), while the other information came from the U.S. Census (state funding), SIPRI (military expenditures), and the U.S. Department of Health and Human Services (TANF expenses).

[2] In the early days of Social Security, lawmakers from opposite sides of the political spectrum feared either that the program would generate savings that would dwarf federal debt to that point, while others feared low individual benefit levels. This provided the political opportunity to reduce the program's savings while expanding the social safety net to wives and widows, thus leading to the expansion of Social Security benefit payments through old-age spouse and survivor's benefits (Altmeyer, 1966). The expansion of Social Security Old-Age Insurance to include spouse and survivor's benefits meant that the Social Security Administration would begin to pay benefits to individuals who were not contributing, weakening the notion of Social Security as an earned benefit.

Before introducing where this paper's contribution fits into the retirement literature, it is important to understand how auxiliary benefits tie the household's retirement decisions together and the magnitude of these benefits. The Social Security Spouse's Benefit specifies that a worker's spouse is eligible to claim an additional 50% of the worker's Social Security benefits, but the net gain is reduced based on the spouse's own earnings history. For example, consider a single income household where the husband is individually entitled to monthly benefits of $1,200. The wife, in this household, would receive an auxiliary benefit of $600 to bring her to 50% of her husband's monthly benefit level, yielding a combined $1,800 in household benefits. In a dual income household, alternatively, if each person is entitled to a benefit of $600 (the same baseline entitlement of $1,200 as above), then the spouse's benefit is zero. Despite the equivalent baseline entitlements, the single earning household would receive $600 more in household benefits. Additionally, the survivor's benefit specifies that the surviving member of a marriage is entitled to the greater amount of her own benefit, or the deceased's benefit. Therefore, if the husband died in our example, the single income household would have $1,200 in monthly benefits, while the dual income household would only receive a total of $600 in monthly benefits. In addition, the worker's spouse cannot claim the Spouse's Benefit until the worker has claimed his or her benefit.

In 2011, 5.16 million people received an old-age spouse's benefit, and 7.78 million people received an old-age survivor's benefit, most of whom were women. The average monthly benefit for a wife who was not entitled on her own earnings history was $608, and for a widow or widower, it was $1,185. Approximately half of women who receive the Spouse's Benefit are dually entitled, meaning that they are entitled to a benefit on their own earnings record, but that it is less than 50% of their husband's benefit. Consequently, these women receive the difference between their own benefit and the Spouse's Benefit (i.e. in the end, they receive the same amount as an individual who was not entitled to a benefit on her own earnings record). The average monthly Spouse's Benefit portion for these dually-entitled women is $243.64. While the fraction of women entitled to auxiliary benefits has fallen from 61.2% in 1960 to 52.5% in 2011, these benefits still affect the majority of the households over the age of 62 in the United States.[3]

This is the first paper to use a structural retirement model to estimate the effect of the Spouse and Survivor's Benefits on the household's retirement decisions. Using a structural model is important for understanding these benefits, because they have remained largely unchanged since their introduction in 1939, preventing a natural experiment. Furthermore, modeling the choices

---

[3]A study conducted by the AARP (2011) indicates that 97% of people surveyed were aware of the survivor benefit, while 51% of people who had not claimed Social Security benefits were aware of the spouse's benefit. Using my own calculations from the AARP's data, I examined groups most likely to gain from the existence of the spouse's benefit. I find that 62% of women with less than 20 years of work who have not claimed their Social Security benefit are aware of the spouse's benefit. I also find that 60% of men whose wives have less than 20 years of work and who have not claimed their Social Security benefit are aware of the spouse's benefit.

of each household member is important because households are becoming increasingly comprised of two income earners. Past studies have focused on models of individual decision-making, ignoring the possibility that married couples may have correlated preferences or derive benefit from each other's company. Focusing only on individuals misspecifies the impact of any entitlement or pension program. A weighted sample of the Health and Retirement Study indicates that 92.82% of men and 95.17% of women have been married, divorced, or widowed, implying an analysis based on men alone does not represent a complete picture of retirement decisions. Studies of individual retirement decisions, however, highlight issues and explanations that are important for understanding the effect of the Spouse and Survivor's Benefits.

The existing retirement literature on Social Security focuses on understanding the role of Social Security's primary earner benefit in explaining the decrease in male labor force participation and explaining spikes in retirement at ages 62 and 65 (Social Security's early and normal retirement ages, respectively). Explanations include (1) actuarial unfairness to benefit adjustments for delayed claiming, (2) borrowing constraints, (3) other beneficiary programs such as Medicare, and (4) uncertainty surrounding future income and health expenses (Gustman and Steinmeier, 1986; Rust and Phelan, 1997; French, 2005; French and Jones, 2011). My model will reflect this literature by including medical expenses and health uncertainty, variation in healthcare coverage, and limited savings (i.e. an individual will not be able to borrow against Social Security or her pension).

More recently, the structural retirement models mentioned above have been extended to capture the interconnected decisions of households. Gustman and Steinmeier (2000, 2004) provide a framework for household decision-making that accounts for interdependence of preferences, but abstracts from uncertainty and allows households to perfectly smooth consumption by borrowing without limit across time. Blau and Gilleskie (2006) create a household model of labor supply and introduce uncertain medical expenditures and employment, but do not allow for savings and do not separate labor supply and claiming decisions. More recently, van der Klaauw and Wolpin (2008) made an important contribution by modeling household labor supply while permitting savings and heterogeneity in preference for consumption.

Relative to other retirement models, such as van der Klaauw and Wolpin (2008), I solve my model separately for each household so that it captures how Spouse and Survivor's Benefits interact with the couple's age difference, private pensions, and unique earnings histories. Solving my model separately by household allows the model to parse preference heterogeneity from heterogeneity in a couple's earnings histories and a couple's age difference. I highlight here three differences from previous retirement models that are important for identifying the effects of Spouse and Survivor's Benefits: (i) households differ at baseline by their preference for individual and joint leisure, (ii) households respond to each individual's unique pension incentives as part of the household labor supply decision, and (iii) household members can claim benefits separately from each

other and independent of their labor supply decision.

(i) My model is estimated on the 1992 cohort of Health and Retirement Study (HRS), which first observes a household when one member is between age 51 and 61, implying that many of the long-term decisions of the household are established (i.e. who works, how much is saved, how much time is spent together). Since I do not model household formation and bargaining prior to when it is first observed in 1992 (baseline), I allow for households to vary by how its members value their own and joint leisure. Some marriages involve a substantial amount of shared time because the couple places a high value on that interaction. Other marriages may be characterized by one member specializing in work, and the other specializing in home production. Close relationships and household specialization are characteristics of a social structure that was developed a long time before this paper's analysis begins, and so these individuals must be treated differently from couples who enjoy separate activities or both work.[4] Similar to van der Klaauw and Wolpin (2008) and French and Jones (2011), I account for these initial conditions by allowing households to belong to one of a finite number of types. Each household is assigned to a time-invariant type that reflects its preference for individual and joint leisure. The preference parameters of the model then differ by type, leading to different outcomes for otherwise equivalent households.

(ii) In my HRS sample, 33% of households have at least one current defined benefit pension. Private pension plans will often have sharp financial incentives to delay retirement until an early retirement age, and to retire by no later than a normal retirement age. Failing to account for these incentives would bias the parameter estimates and any predictions made using the model. Figure 1 shows, by age, the substantial variation in the growth rates of annual pension benefit payments. At ages 55 and 60, there are peaks in the 95th percentile of benefit payment growth rates, which is due to these ages being common early retirement dates for defined benefit pension plans.

Heterogeneity in benefit payment growth is also common in Social Security, particularly for individuals without the maximum 35 years of earnings history, which is common for women. For one-fifth of individuals aged 62 with less than 35 years of earnings history, the Social Security benefit payment growth rate exceeds 5% for an additional year of work.[5] In order to avoid the retirement incentives induced by defined benefit pension plans and to simplify the model's estimation, many authors restrict their samples to households that are without pension plans and do not keep track of Social Security earnings histories (Rust and Phelan, 1997; van der Klaauw and Wolpin, 2008). By omitting work histories and pension plan details, these papers focus on a portion of the population that has lower incomes and for which Social Security benefits represent a very

---

[4]The Health and Retirement Study reports that of the married individuals in the 1992 cohort, 17% somewhat or do not look forward to retirement with his or her spouse, and 18.6% somewhat or do not enjoy time spent with his or her spouse.

[5]When I refer to benefit growth rates, I am referring to the growth rate in annual benefit payments once the beneficiary has claimed. I am not referring to the change in the expected present discounted value of pension wealth.

important part of retirement wealth. These household will be more likely to claim their benefit as soon as they are eligible, and the implications drawn from these models are not representative of the effects that a change in the Social Security program would have on the broader U.S. population.

(iii) Benefit claiming and retirement are not equivalent, as indicated by the fact that while more than 50% of individuals in my sample claim Social Security benefits at the early retirement age, the majority continue to work. People with small incomes or poor health may find it optimal to claim Social Security benefits as early as possible. Single income couples may find it optimal for the earner to claim as soon as possible, so the nonworker can access the Spouse's benefit. The choice of when to claim annuitized benefits, like Social Security and defined benefit pensions, is dependent on each couple's unique incentives stemming from their health and earnings history, their accumulation of non-annuitized liquid assets, and their opportunity cost of delayed claiming.

Authors have often linked claiming of benefits with an individual's retirement, but benefit claiming is becoming more strategic as Social Security incentivizes delayed claiming, couples live longer, and phased retirement or unretirement becomes more common (Shoven and Slavov, 2013). Using the HRS, Maestas (2010) showed that 18.2 - 23.8% of workers who initially exit the labor force with the intent of retiring return to full or part-time work within six years. Furthermore, she finds that, of the individuals who exit their job with the intent to fully retire, only 33.9% of individuals claimed their pension at the time they exited the job. Other studies point to greater

Figure 1: Growth Rates in Annual Benefit Payment of Defined Benefit Pensions, by Age (Multiply vertical axis by 100 for percent growth rates)
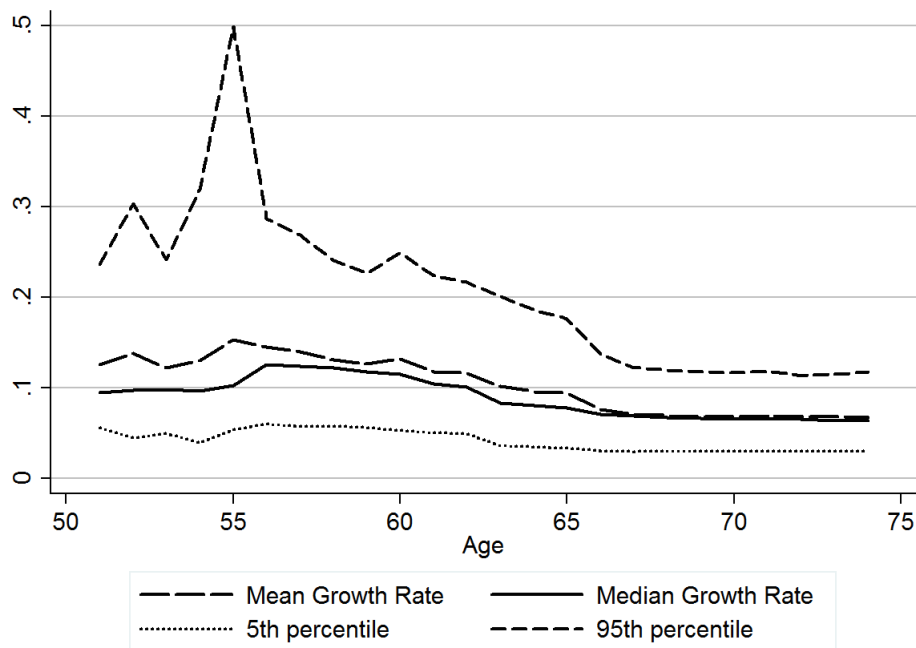
early claiming rates for Social Security than are predicted by a typical life-cycle model (Hurd et al., 2002; Coile et al., 2002; Sass et al., 2013). The puzzle surrounding high early claiming rates of Social Security, and the more arbitrary claiming rates of pensions, can not be captured by previous structural models because most do not separate the benefit claiming decision from the labor supply decision, and those that do only model the husband's decision.

In section 2, I introduce a simple model to build intuition for the effect of the Spouse's Benefit on the high and low earner's work decisions. Section **??** introduces the dynamic, life-cycle model, while section 4 describes the data selection from the HRS. Section 5 describes the estimation method, the baseline results, and the ability of the model to replicate empirical regularities. Section 6 conducts three policy experiments on Social Security benefits and discusses the implications of these changes for individual labor supply, benefit claiming, and average lifetime benefits received from Social Security. I conclude in section 7 by summarizing the key results and discussing the implications of my model.

2   Simple Two Period Model

The discussion in the first section demonstrated that a household's Social Security primary and auxiliary benefits can be a complicated result of household leisure choices. In this section, I provide a simple two period model to help the reader understand the impact of spouse benefits on the household's labor supply. In the next section, I will introduce a more realistic model that is meant to capture the complexities associated with the entire life-cycle of a household.

2.1   Setup

To simplify the discussion, suppose that a household lives for two periods. The first period represents the timeframe where the household chooses to work, and the second period represents retirement. The household is comprised of two agents who choose their joint consumption in both periods and how much to work in the first period. The household's utility is derived from joint consumption and individual leisure in both periods, where δ represents how the household discounts future utility, as in:

$$U = u(C_1, L_{H,1}, L_{W,1}) + \delta \cdot u(C_2, 1, 1) \tag{2.1}$$

where C_t represents the household joint consumption in period t, and $L_{H,1}$ and $L_{W,1}$ represent the husband and wife's leisure in period one, respectively.[6] I assume that consumption and leisure are normal goods.

---

[6]The assumptions that I place on the household preferences described in (2.1) are that the utility function is convex, monotonic, and inter-temporally separable.

The budget constraint is determined by each household member's income, which is a function of his or her first period leisure and potential income (e.g. $Y(1 - L_{H,1}, Y_H^*)$), as well as Social Security old-age and auxiliary benefits. An individual's primary benefit is determined by a three bracketed formula based on his or her indexed monthly earnings. For illustrative purposes in figures 2 and 3, I use the 1994 beneficiary rules, where an earner would receive 90% of the first $4,440, 32% of the next $22,320, and 15% of the remaining $33,840, for a maximum annual benefit of $16,214. A household can save earnings from period 1, but cannot borrow from the Social Security benefits due in the second period. The auxiliary benefits for the low earner are equal to 50% of the high earner's benefit level or her own benefit level, whichever is greater.[7] The household's budget constraint in period 1, assuming no preexisting assets, is

$$C_1 + A_1 = Y(1 - L_{H,1}, Y_H^*) + Y(1 - L_{W,1}, Y_W^*), \qquad (2.2)$$
$$A_1 \geq 0.$$

The budget constraint in period 2 is

$$C_2 = (1 + r) \cdot A_1 + SSB\left(Y(1 - L_{H,1}, Y_H^*), Y(1 - L_{W,1}, Y_W^*)\right), \qquad (2.3)$$

where A_1 is the household's assets saved in period 1 and $SSB(\cdot, \cdot)$ represents the household's Social Security benefit, which is a nonlinear function of the husband and wife's leisure decisions in period 1 and their potential incomes, Y_H^* and Y_W^*.[8] For ease of exposition, I will assume, only in this section, that the husband is the high earner in the household (i.e. $Y_H^* > Y_W^*$).

2.2 Effect of Spouse Benefit on Low Earner

Figure 2 provides an illustration of the wife's budget constraint with the spouse benefit kink point, assuming the husband works full-time $(L_{H,1} = 0)$, and that the household has nonnegative savings in the first period $(A_1 > 0)$.[9] In this figure, point A represents the outcome for households that

---

[7]In the simple two period model, I do not include delayed claiming increments or early claiming penalties.

[8]The Social Security benefit is function of each household member's income. The Social Security Benefit is defined by:

$$SSB\left(Y\left(L_{H,1},Y_H^*\right),Y\left(L_{W,1},Y_W^*\right)\right) = max\{1.5 \times SSB_H\left(Y\left(L_{H,1},Y_H^*\right)\right), 1.5 \times SSB_W\left(Y\left(L_{W,1},Y_W^*\right)\right),$$
$$SSB_H\left(Y\left(L_{H,1},Y_H^*\right)\right) + SSB_W\left(Y\left(L_{W,1},Y_W^*\right)\right)\},$$

where for $i \in \{H,W\}$,

$$SSB_i(Y) = \begin{cases} 0.9 \times Y & \text{if Y<\$4,440} \\ 0.32 \times Y + 0.9\,(4,440) & \text{if \$4,440}\leq\text{Y<\$26,760} \\ 0.15 \times Y + 0.32\,(22,320) + 0.9\,(4,440) & \text{if \$26,760}\leq\text{Y<\$60,600} \\ \$16,124 & \text{if \$60,600}\leq\text{Y.} \end{cases}$$

Also, note that legally households cannot borrow against their Social Security benefits, therefore Social Security benefits only become available in the second period as in (2.3).

[9]Note that I am only considering the returns to working in the first period relative to consumption in the second period, because the choice of consumption in the second period determines the consumption in the first period through

Figure 2: Example of a Household Budget Constraint
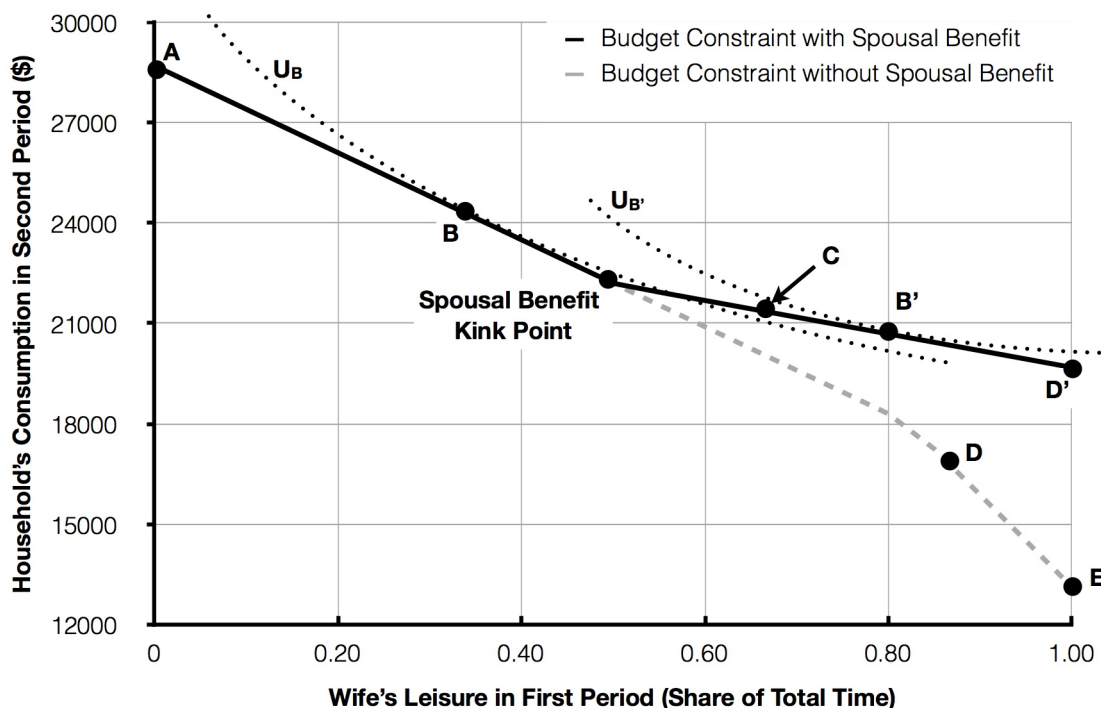
Notes: The baseline budget constraint (solid line) assumes the husband earns $40,000 and saves none of his income, and the wife's potential income was $25,000 if she worked full time and the wife saves 20% of her income.

find it optimal for the wife to work full-time. The indifference curve, U_B, represents a set of preferences for a household where the wife would optimally supply a level of labor corresponding to point B without the spouse benefit, but with it, she reduces her labor supply significantly to point B'. The set of preferences described by U_B represent an example where the spouse benefit results in the wife's leisure discontinuously jumping to a higher level, an issue discussed in greater detail below. Point E represents the outcome for households with a high preference for the wife's leisure, indicating that the wife would not work regardless of the spouse benefit's existence.

Using this figure, the intuition for the effect of the spouse benefit on labor supply can be seen by the wife's decision if she is a little to the right of the spouse benefit kink point (point C in figure 2). In this case, each additional hour of leisure sacrificed increases second period consumption by only the marginal savings from the wife's earnings, because her Social Security benefit is based only on her husband's earnings history. Alternatively, if she works enough to be to the left of the kink point, then her return from each additional hour of leisure sacrificed is the change in Social Security benefits based on her own earnings history plus the marginal savings from her earnings. Thus, the household's budget constraint becomes steeper.

For a household with strictly convex preferences over consumption and leisure, the existence of the spouse benefit kink point will cause the wife to work less in certain circumstances, because it reduces her return to work. I consider three cases, represented by the letters in figure 2.

A The wife continues to maximize household utility by working the same amount regardless of the spouse benefit's existence. This can only occur where the wife optimally supplies labor to the left of spouse benefit kink point in figure 2.

B The household would optimally supply labor to the left of the kink point without the spouse benefit, but then jump to a higher level of leisure, to the right of the kink point, with the spouse benefit. This is illustrated by a household with preferences represented by U_B at point B without spouse benefits, jumping to point B' with a higher level of utility, $U_{B'}$, with the inclusion of spouse benefits (as in figure 2).

D The wife would have optimally worked a positive amount at a level of leisure to the right of the spouse benefit kink point. With the spouse benefit, she will now find that her optimal choice is working less or not working (D'). This is because each additional hour of leisure sacrificed increases second period consumption by only the fraction of her income that is saved. At the extreme of this case, the household maximizes utility by the wife not working, as in point E in figure 2. In this case wife's work behavior is unaltered by the existence of the spouse benefit, but consumption increases from E to D'.

---

the typical Euler equation: $\frac{\partial u}{\partial C_1} = \delta(1+r)\frac{\partial u}{\partial C_2}$ if assets are nonnegative.

To summarize, the spouse benefit weakly discourages the low-earning spouse from working by reducing her return from work because the income effect (i.e. receiving more benefits in retirement increases demand for leisure) and substitution effect (i.e. lower returns from working increases demand for leisure) act in the same direction.

2.3 Effect of Spouse Benefit on High Earner

The spouse benefit also impacts the husband's decision (i.e. the high earner) to work by increasing his return to work if his wife earns a sufficiently low income. As represented in figure 3a, the spouse benefit increases both the husband's return from work and increases the household's income if the husband's first period earnings are sufficiently high relative to his wife's earnings. Much like a change in wage, the spouse benefit induces an income effect that discourages work, but a substitution effect that encourages it. Figure 3 shows the impact of the spouse benefit on the husband's first period leisure decision holding constant the wife's leisure decision. Similar to the impact on the wife's budget constraint discussed above, there are four possible cases for the husband that correspond to points labeled in figure 3.

A Husband works full-time and the introduction of the spouse benefit increases income but does not alter his labor supply - pure income effect.

B The income effect from the spouse benefit dominates the increase returns from work leading to an increase in the husband's leisure, as in figure 3a.

C The increase returns from work dominate the income effect from the spouse benefit leading to a decrease in the husband's leisure, as in figure 3b.

D For some original leisure choices to the right of the spouse benefit kink point, the husband either chooses never to work or does not make enough income relative to his wife for the spouse benefit to change his first period decision.

Unlike the low earner, the high earner is impacted by offsetting income and substitution effects, making the final impact on his labor supply ambiguous.

2.4 Summary

The combined impact of the spouse benefit on the high and low earner is to discourage the low earner from work but has ambiguous incentives on the high earner's labor supply. The existence of the spouse benefit will matter more to households where the difference in potential earnings are the greatest.

Figure 3: Household budget constraint relative to Husband's Leisure Choice

(a) Dominant Income Effect

(b) Dominant Substitution Effect
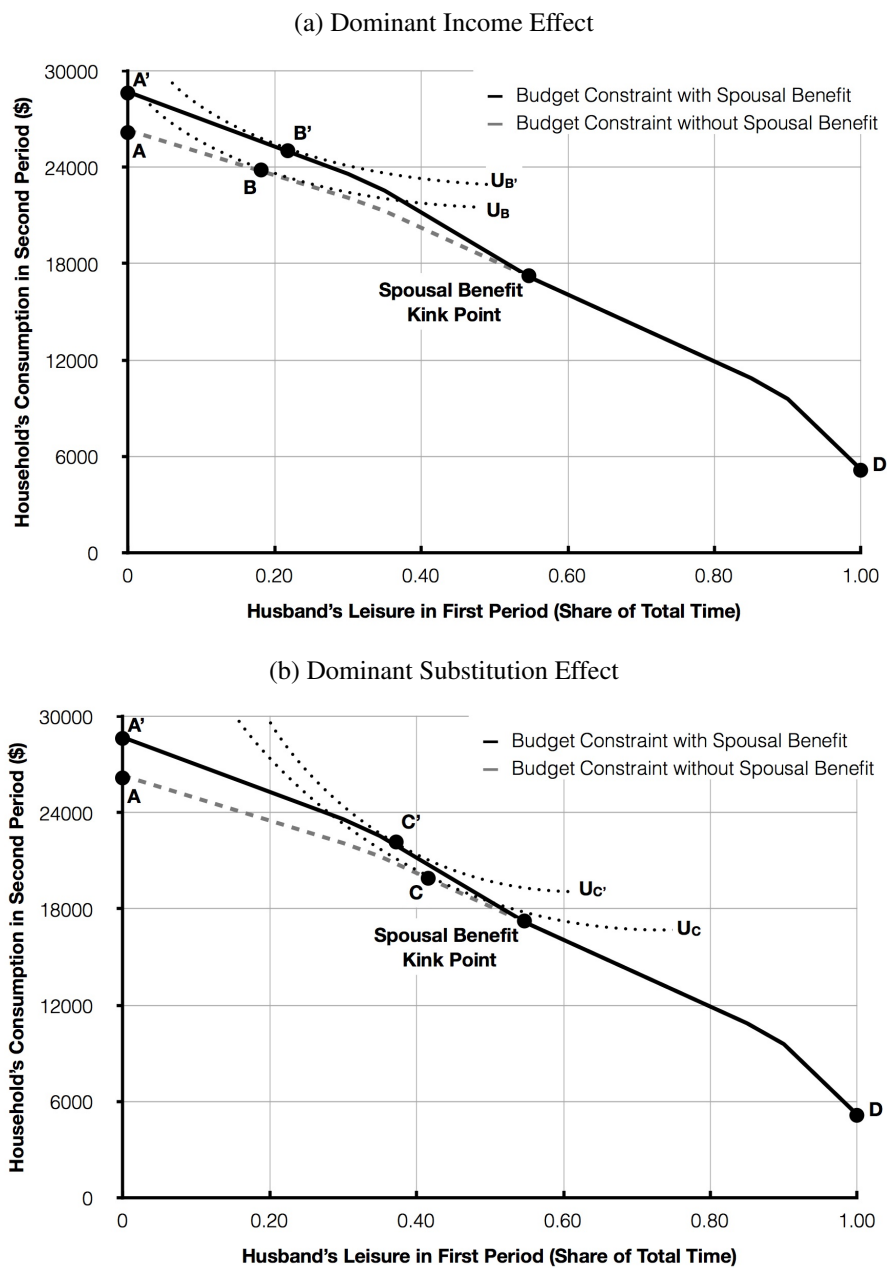
Notes: The baseline budget constraint (solid line) assumes the wife earns $5,000 and saves 20% of her income, and the husband's potential income was $60,000 if he worked full time and the husband saves 20% of his income.

In a model that includes more decision periods, which can capture the fact that Social Security benefits are based on lifetime earnings histories, the appropriate comparison would be households where the earnings histories are more disparate. A wife who has an earnings history that is substantially lower than her husband's earnings history would not benefit from Social Security based on her own earnings history, and so earns no additional retirement benefits from continued work. In the context of the life-cycle model presented in the next section, this implies that spouse benefits help single earning households and discourages the low earner from returning to work because she receives no retirement benefit from further work. The impact on the husband is more ambiguous because it provides a lifetime income effect discouraging work while increasing the marginal return from work.

I will find, in the policy experiments of §6, that the husband's substitution effect will dominate for my sample, implying that men would work 0.11 years less without the spouse benefit (intuitively $C' \to C$ in figure 3b).

3 Model

In this section I introduce a dynamic life cycle model of labor supply and benefit claiming for married couples who maximize their utility based on state variables in year t (X_t), preference parameters (θ), and parameters of the data generating process (χ). This model differs substantively from most structural retirement models by considering the choices of a couple instead of just the male head of household. Uncertainty arises from random mortality, health changes, and medical expenses, while further permanent heterogeneity is based on variation in households' preference for work, leisure, and future consumption.

3.1 Choice Set

Every individual, $i \in \{H\,(\text{husband}), W\,(\text{wife})\}$, is part of a household, h, and each period (year) the household decides (i) whether each individual works, (ii) whether each individual claims his or her Social Security or other claimable pension benefits, and (iii) how much income to consume, $C_{h,t}$.[10]

Individual decisions are made via household decisions. As a result, I will abstract away from strategic decision making between household members. Intra-household bargaining is assumed to be fixed at baseline and is reflected in permanent differences in households' preference for own and spousal leisure (discussed in greater detail in §5.1.4). Household preferences reflect the externality of each person's leisure on the other member of the couple, and the relative weight each individual has in the decision making process.

Retirement can be an ambiguous concept, with many workers retiring and then proceeding to un-retire or return to the labor force within a few years (Ruhm, 1990). As a result, I do not define retirement explicitly, rather, I focus only on the per period labor supply decision. In this model, all individuals will eventually opt out of work, given a sufficiently advanced age. Each household participant's labor supply, $N_{i,t}$, is restricted to one of four states:

$$
N_{i,t} = \begin{cases}
1 & \text{if working full-time in baseline job} \\
1 & \text{if working full-time in non-baseline job} \\
0.5 & \text{if working part-time in non-baseline job} \\
0 & \text{if not working.}
\end{cases}
$$

I distinguish between baseline and non-baseline jobs both because the assumptions regarding how earnings evolve over-time will differ between these jobs and because only baseline jobs will have

---

[10] All consumption in this model is joint consumption because the HRS is unable to distinguish between joint and individual consumption.

pensions associated with them.

Assuming the household member is eligible to claim benefits, the household can also choose to claim benefits, $B_{i,t} = \{1 \text{ for claim, } 0 \text{ for no claim}\}$. Depending on the types of benefits an individual is eligible to claim, this can include both a defined benefit pension and a Social Security benefit, just one of these benefits, or neither. These benefits do not have to be claimed in conjunction with leaving the labor force, but current and future benefit levels may vary with the household's labor force decision (see §4.2 and §4.3 for a discussion on claimable benefits). There is no "claiming" of defined contribution plans, because these funds are treated as savings. All benefit claiming decisions are treated as absorbing states.

3.2 Preferences

A household, h, maximizes its expected present value of lifetime utility by choosing their consumption, labor participation and whether or not to claim benefits. The household instantaneous utility function in year t is given by:

$$U\left(C_{h,t}, L_{H,t}, L_{W,t}\right) = \frac{C_{h,t}^{1-\alpha} - 1}{1-\alpha} + \frac{D_{H,t}L_{H,t}^{1-\gamma_H} - 1}{1-\gamma_H} + \frac{D_{W,t}L_{W,t}^{1-\gamma_W} - 1}{1-\gamma_W} \quad , \tag{3.1}$$

where the parameter $\alpha > 0$ captures the household's diminishing returns from joint consumption.

Each individual's leisure, $L_{i,t}$, is defined as:

$$L_{i,t} = L - N_{i,t} \quad , \tag{3.2}$$

where L is the endowment of leisure. Note that the relative value of part-time to full-time work changes based on the parameter γ_i. I fix γ_i across time, thus only permitting age to affect the marginal rate of substitution for identification purposes.[11] I do not include a specific leisure cost for reentry into the labor force.

The coefficient $D_{i,t}$ represents a modifier for each individual's marginal rate of substitution between leisure and consumption. It changes based on state variables, including a constant term for the husband or wife, the age of the husband or wife, the health of the husband or wife, and additional variables meant to reflect the change in the individual's substitution between consumption and leisure. In the case of the husband ($i = H$), it takes the form

$$D_{H,t} = exp\left(\beta_H + \beta_{H,age}age_{H,t} + \beta_{H,health}health_{H,t}\right. \tag{3.3}$$
$$\left. + \beta_{H,SP}\mathbf{1}\left[N_{W,t} > 0\right] + \beta_{H,SFT}\mathbf{1}\left[N_{W,t} = 1\right] + \varepsilon_H\right) \quad ,$$

---

[11] Alternatively, Gustman and Steinmeier (2005) allow their equivalent of γ_i and $D_{i,t}$ to vary across time, make identification harder to argue.

where the last two terms on the right-hand side represent how the wife's participation in the labor force and whether she works full or part-time affect the husband's preferences over consumption and leisure. Analogously, the wife's modifier, $D_{W,t}$, is determined by

$$
\begin{aligned}
D_{W,t} \;=\; & exp\left(\beta_W + \beta_{W,age}age_{W,t} + \beta_{W,health}health_{W,t}\right. \\
& \left. + \beta_{W,SP}\mathbf{1}\left[N_{H,t} > 0\right] + \beta_{W,SFT}\mathbf{1}\left[N_{H,t} = 1\right] + \varepsilon_W\right) \quad .
\end{aligned}
\tag{3.4}
$$

$D_{H,t}$ and $D_{W,t}$ capture the complementarity of spousal leisure time, and how it differs between part and full-time work. This setup, where I distinguish the impact of health, age, and joint marital time on the rate of substitution between consumption and leisure will help identify the effect of changes in joint benefit programs like Social Security.

After controlling for age, health status, and leisure complementarities, there may still exist a permanent level of heterogeneity across the population in the relative value of leisure (see Gustman and Steinmeier (2004)). This individual fixed effect for higher value of retirement to an individual, $\varepsilon_i \sim N(0, \sigma_{\varepsilon_i})$, is treated as permanent component of the individual's utility. If $\varepsilon_i > 0$, then the individual receives greater returns from leisure, and is thus likely to leave the labor force sooner. Additionally, ρ_{HW} represent the correlation between ε_H and ε_W. If households sort based on preference for leisure, then $\rho_{HW} > 0$.

These preferences are non-homothetic. Homotheticity constrains the share of income to be spent on consumption to remain unchanged. If income doubles, so will the share of income spent on consumption. While this often provides a reasonable baseline from which to examine long-run behavior, it oversimplifies the relationship between retirement and savings. Hypothetically, if household were to receive a surprise endowment in one period, it might choose to save the entire sum and retire a year sooner. Homothetic preferences would not permit this choice. I chose to allow preferences to be non-homothetic, similar to most of the retirement literature (see van der Klaauw and Wolpin, 2008; Rust and Phelan, 1997; Gustman and Steinmeier, 2005). The non-homothetic preferences allowed for in my model permit the household's willingness to substitute leisure across time to differ from its willingness to substitute consumption across time.

Individuals have a probability $s_{t+1}^i = s(age_{i,t}, health_{i,t}, i)$ of surviving until period $t + 1$, discussed further in §5.1.2, and households discount the future at rate δ. Households that become single through widowhood are assumed to receive a 50% greater return from \$1 of consumption than a two person household, $C_{widow} = 1.5 \times C_{married}$, and the deceased individual i is assumed to not participate in the labor force, $N_{i,t} = 0$, and does not contribute to household utility, $D_{i,t} = 0$.[12] As in De Nardi (2004); De Nardi et al. (2010), households where both members are deceased value

---

[12]This is equivalent to the implicit returns to scale assumed by the Social Security spouse and survivor's benefits.

their bequests from assets, $A_{h,t}$, according to the function

$$b(A_{h,t}) = \theta_B \left(\frac{(A_{h,t} + \kappa)^{1-\alpha} - 1}{1 - \alpha} \right) \quad . \tag{3.5}$$

This a standard "warm-glow" bequest, where the household gets non-negative utility from leaving assets to future generations. The bequest shifter, κ, and the bequest intensity, θ_B, determine the value of the additional assets, in terms of utility, relative to the other states where one or both members of the household are alive.

3.3 Budget Set

The household is able to accumulate assets, $A_{h,t}$, over its lifetime subject to the following equation

$$A_{h,t+1} = A_{h,t} - C_{h,t} - M_{h,t} + Y_{h,t} + tr_{h,t} \quad , \tag{3.6}$$

where $C_{h,t}$ is per period household consumption, and $M_{h,t}$ is stochastic health expenses. Additionally, $Y_{h,t}$ is per period income and $tr_{h,t}$ are government transfers, which are defined more explicitly below.

A household's per period income can come from a number of sources: household interest income, $rA_{h,t}$, a household Social Security benefit, $ssb_{h,t}$, and each individual's annual earnings, $\omega_i(N_{i,t}, age_{i,t})$, and defined benefit pension income, $db_{i,t}$, where all of these sources of income are subject to tax, tx:

$$Y_{h,t} = Y \left(rA_{h,t} + ssb_{h,t} + \sum_{i \in h} \left(\omega_i(N_{i,t}, age_{i,t}) + db_{i,t} \right), tx \right) \quad . \tag{3.7}$$

Taxation in this model is handled using the Internal Revenue Service rules for taxation in 1992 and assumes that individuals do not experience the changes to the tax code since 1992. Further details on how taxes are calculated are included in Appendix A.

Finally, households are borrowing constrained based on their flow of income. Following past work (e.g. Hubbard et al. (1995)) I include a minimum level of consumption that determines government transfers. Government transfers guarantee a minimal, positive consumption level, even if a household is uninsured and experiences a severe medical expense shock. Government transfers are defined by

$$tr_{h,t} = \max \left\{ 0, C_{min} - A_{h,t} - Y_{h,t} \right\} \quad , \tag{3.8}$$

so that an individual will always be able to consume at least C_{min} (i.e. $C_{h,t} \geq C_{min}$).[13]

---

[13] C_{min} will depend on whether the household is single (i.e. widowed) or married. As mentioned in §??, I set

The household Social Security benefit and private pensions are described in §4.2 and §4.3, after the data source is introduced. The evolution of an individual's annual earnings and stochastic medical expenses, are described in §5.1.1 and §5.1.3, respectively, following the description of the data and estimation strategy.

3.4 Recursive Formulation

Each period, a household chooses its consumption and each individual's level of labor force participation, Social Security claiming decision, and pension benefit claiming decision (if applicable). The decision to claim benefits is irreversible, can only be done once the individual reaches early retirement age (62 for Social Security and as early as 55 for some pension plans), and must be done no later than age 70. The household's maximization problem is

$$V_t(X_t) = \max_{C_{h,t}, L_{h,t}, B_{h,t}} \left\{ U\left(C_{h,t}, L_{h,t}\right) + \delta \left(1 - s_{t+1}^H\right)\left(1 - s_{t+1}^W\right) b(A_{h,t+1}) \right. \tag{3.9}$$
$$+ \delta \left(1 - s_{t+1}^H\right) s_{t+1}^W \mathbb{E}\left[V_{t+1}\left(X_{t+1} \mid X_t, t, C_{h,t}, B_{h,t}, L_{h,t}, \text{ wife survives}\right)\right]$$
$$+ \delta s_{t+1}^H \left(1 - s_{t+1}^W\right) \mathbb{E}\left[V_{t+1}\left(X_{t+1} \mid X_t, t, C_{h,t}, B_{h,t}, L_{h,t}, \text{ husband survives}\right)\right]$$
$$\left. + \delta s_{t+1}^H s_{t+1}^W \mathbb{E}\left[V_{t+1}\left(X_{t+1} \mid X_t, t, C_{h,t}, B_{h,t}, L_{h,t}, \text{ both survive}\right)\right] \right\} \quad,$$

subject to a non-negative borrowing constraint and the consumption floor in equation (3.8). Let $C_{h,t}$, $L_{h,t}$, and $B_{h,t}$ represent the set of each household's bundle of choices for consumption, leisure, and benefit claiming, respectively.

The solution to the recursive formulation in equation (3.9) requires solving for each household's consumption, labor force participation, and benefit claiming choices at every age at and after baseline (1992), collectively referred to as the decision rules. These decision rules are calculated by backward induction using the above mentioned model. I describe my choice of recursive and numerical methodology in Appendices B and C, respectively.

$1 of consumption in a two person household to be equivalent to $1.50 of consumption in a widowed household, $C_{single} = 1.5 \times C_{married}$. This is done because households may benefit from economies of scale, and this ratio reflects the implicit economies of scale assumed by the Social Security Administration when handling single versus dual household benefits through the Supplemental Security Income program.

4 Data

4.1 Health and Retirement Study

The model in §**??** is estimated using the original cohort of the Health and Retirement Study (HRS), which was born between 1931 and 1941, and has 12,652 respondents and 7,704 households. The HRS follows these households every two years, and in this study I use data from 1992 through 2010. It collects information on income, work, assets, pension plans, health insurance, disability, individual health, and health care expenditures. It has an impressive retention rate, with approximately 80.5% of the original, surviving cohort responding as of the 9th wave (2008).[14]

The HRS is well suited to estimating my model because it also collects individual Social Security Administrative data and detailed pension data from respondents' employers. The Social Security administrative data includes individual earnings histories for 79.77% of the original cohort. Moreover, the HRS also contacted employers of respondents who reported having employer-provided retirement plans. If the individual consented, then the HRS contacted the employer to obtain a copy of the summary plan description of each plan the employer offered, and then extracted information about the plan or plans relevant to the respondent from these documents. This information was then used in designing a pension calculator for projecting a respondent's benefit levels based on any future retirement date, as described in §4.3.[15]

From the original HRS sample, I keep households that (1) are married in wave 1, (2) are not missing information on their labor force participation in wave 1, (3) have never applied for Social Security disability benefits, (4) are not missing pension or Social Security information, (5) have a spousal age difference of less than 10 years, (6) are not missing information on individual earnings if household members report working, and, for computational reasons, (7) households

---

[14]"A total of 13,687 individuals are in the HRS sample since the baseline interviews in 1992. Over two-thirds (67.1%) of the respondents in this sample have complete interview histories from their initial entry through 2008. The remaining 32.9% have missed at least one interview: an average of 2.7 missed interviews (7.3 average attempts)" (HRS Sample Sizes and Response Rates, 2011). This number is larger than 12,652 because new spouses are added to the sample if a respondent marries after baseline.

"The HRS cohort rate of 80.5% retention at 16 years of survey duration is slightly better than the National Longitudinal Surveys (NLS)-Older Men (76.3%) and Mature Women studies (73.1%), but somewhat below the record levels of the National Longitudinal Survey of Youth 1979 (NLSY79) cohort, which stood at 89% among survivors after 16 years."

[15]A number of studies have examined the selectivity of the Social Security and pension samples (Haider and Solon, 2000; Gustman and Steinmeier, 1999; Kapteyn et al., 2006). For Social Security, sample selection occurs because individuals not permitting their earning profiles to be linked are different from those who do permit their earnings histories to be linked (95% of those who give permission are matched). Individuals who are non-white, in the highest asset or education groups, and who never expect to retire or do not report a retirement date are the least likely to give permission. For pensions, selection may also occur on the ability of HRS to obtain a SPD, conditional on the person giving permission. Individuals who are in the highest asset and earnings groups, are at firms with less than 100 workers, are in management professions, and have a defined contribution plan are the least likely to successfully have their plans matched, conditional on giving permission. In this paper, I use matched SPD only for defined benefits plans. I rely on individual reports from defined contribution plans.

Table 1: Statistics from HRS sub-sample used in estimation (§5)

		Men	Women
Age	Mean	60.8	57.8
	Median	60.5	58.2
	Standard Dev.	3.1	3.4
Earnings*	Mean	$27,431	$11,858
	Median	$20,236	$7,100
	Standard Dev.	$33,112	$20,136
AIME	Mean	$2,013	$675
	Median	$2,207	$446
	Standard Dev.	$894	$675
Predicted Annual Pension Benefit	Mean	$20,205	$6,998
	Median	$9,577	$3,107
	Standard Dev.	$37,409	$11,444
% with Current Pension Benefit*		21.3	26.6
% Working		70.7	52.9
% Working Full-time*		80.2	60.5
% in Self-Reported Bad Health*		12.1	10.6
% White		89.4	89.0
Average Years of Education		12.5	12.4

		Household
Assets*	Mean	$339,267
	Median	$182,558
	Standard Dev.	$569,046
% with Retiree Health Insurance		67.3
% with Tied Health Insurance		15.2
% with No Health Insurance		17.4
Preference Type (Own Leisure, Spousal Leisure)	Out	17.4%
	Low, Low	22.9%
	High, Low	18.0%
	Low, High	20.9%
	High, High	20.8%
Fraction of Women Eligible for Spousal Benefit	Overall	56.4%
	1st Asset Quantile	56.0%
	2nd Asset Quantile	56.7%
	3rd Asset Quantile	56.7%
Number of Households		948

*Notes: Sample consists of only households where one member is born between 1931 and 1935. Individual earnings is conditional on participating in the labor force in 1992. Predicted Annual Pension Benefit is defined benefit pensions that are vested and is conditional on having a pension. The percentage with current pension is conditional on participating in the labor force in 1992. The percentage working full-time is conditional on participating in the labor force in 1992. Self-reported bad health is based on an individual reporting his or her overall health status as being fair or poor at the time of the HRS interview. Assets are comprised of non-annuitized assets, including net housing wealth and defined contribution plans. Assets do not include Social Security wealth, defined benefit pensions, or defined contribution plans that were converted to annuities prior to 1992.

where no more than one member has a defined benefit pension.[16] After this sample selection, I am left with 1,728 married households. I use the Social Security Administrative data for earnings and respondent reports for periods not covered by the Social Security data. Doing so yields an average of 14.95 annual observations per household (out of a maximum possible of 20), providing a long history of observations. My HRS sample will not exactly reflect participation patterns observed from a cross-section of ever-married individuals from the U.S. census, or similar sample. The omission of divorced, separated, and previously-widowed households increases the sample's labor force participation slightly, but eliminating those households that ever apply for Social Security disability benefits increases the sample's labor force participation at all ages by approximately 10%. This result is not surprising since individuals who credibly apply for disability will likely have a reduced ability to participate in the labor force.

I use a subsample to estimate my model, consisting of all households with one member born between 1931-35. This results in a final sample size of 948. I use the subsample born between 1937-41 for testing the out-of-sample fit of the model after it is estimated. Table 1 shows the descriptive statistics of the subsample used in the estimation of the model. A more detailed version of the sample selection and sample statistics for the entire sample as well as the out-of-sample fit cohort are included in Appendix E.

Given that I am looking at households where at least one member was born between 1931-35, it is not surprising that the median age of men and women is 60.5 and 58.2, respectively. The average difference in age of a married couple is 3 years. The sample is primarily white, with slightly more than a high school education on average. Assets are heavily skewed, as expected, with mean assets being $339,267 and median assets being only $182,558. Perhaps the most surprising feature of the data is that the fraction of women eligible for the spouse's benefit is roughly equal across the asset distribution. This will be particularly noteworthy when I discuss the reaction to changes in the spouse's benefit by asset quantile in §6.4.

4.2   Social Security

This paper's core research question is the effect of Social Security's benefit structure on the retirement decision of a couple. Therefore, in this section I carefully detail the incentives created by Social Security's benefit structure, and which are included in my model. The HRS has detailed earnings histories from the Social Security Administration, which permits using true earnings histories to calculate an individual's financial alternatives based his or her own claiming and labor supply decision as well as the claiming and labor supply decision of his or her spouse. Social

---

[16]Additionally, I drop annual observations if employment or health status of either household member is not reported, and if health insurance status cannot be determined when the household is less than age 65 (Medicare age). Households with two defined benefit pensions are dropped (170 households) because calculation of their decision rules takes the same time as the remainder of the sample.

Security is based on a worker's best 35 years of earnings, but similar models of life-cycle labor supply do not incorporate that benefit growth rates differ by individual because of the variation in individuals' earnings histories. The model includes the specific Social Security rules as they apply to the primary earner and the earner's spouse and survivor, as well as the special tax treatment of Social Security, the earnings test, and each worker's unique earnings history.

An earner is defined as someone who contributes to the Old Age and Disability Social Insurance Program, which I will refer to as Social Security. This program has three major parts: (i) a pension benefit for the earner, (ii) auxiliary benefits for an earner's spouse, survivor, and in some cases children and parents, and (iii) a disability benefit. In the next two subsections I will focus on the first two parts of the Social Security program. I leave the additional complexity of integrating spousal decisions with disability application decisions to future work. In the final subsection, I describe how Social Security benefits can be taxed or reduced due to work.

4.2.1 Primary Earner Benefits

An earner qualifies for a Social Security benefit (i.e. becomes insured) if he or she has 40 qualifying quarters of coverage (QC).[17] His or her benefit is computed using a multistep formula. First, the earner's average indexed monthly earnings (AIME) is calculated by taking the average of the best 35 years of earnings since 1950, where earnings before age 60 are indexed by the average annual wage at age 60 (earnings after age 60 are not indexed).[18] Second, the earner's primary insurance amount (PIA) is based on a progressive calculation, where the earner receives 90% of his or her first $761 of AIME, 32% of the next $3,825 of AIME, and 15% of AIME over $4,586 (assuming reaching age 62 in 2010). The PIA bend points change every year based on the average U.S. annual wage. For an earner, they are calculated using the bend points in the year the worker reached age 62. Third, the AIME is increased each year by a cost of living adjustment based on the consumer price index.

Finally, if the earner claims his or her Social Security Benefit (SSB) in the month he or she achieves the full retirement age, then the benefit is equal to the PIA. The full retirement age is 65 for workers born before 1938 and increases gradually to age 67 for any earners born after 1959. Alternatively, earners may choose to claim their benefits as early as age 62. An early claimer's benefit, however, is reduced by 6.67% for the first three years before the full retirement age and then reduced by an additional 5% for any additional years. Earners may also choose to claim their

---

[17]In some cases, a worker can become qualified if he or she has less than 40 QCs. These include earners who were born before 1928 who only need the the difference between the year they reach age 62 and 1950 to qualify (e.g. an earner born in 1926 will only need 38 QCs). An earner must have a minimum of 6 QCs at any point to qualify for coverage.

[18]Alternatively, for those born before 1928, the number of years used in this calculation is only the difference between the year they reach age 62 and 1955 (e.g. an earner born in 1926 will only use their best 33 years of earnings since 1950)

benefits after the full retirement age, in which case these benefits are increased by up to 8% for each year of delayed claiming up to age 70. The delayed retirement credit has been gradually increasing over the sample period in order to avoid disincentivizing work - previously the delayed retirement credit was only 1% annually.

4.2.2 Auxiliary Benefits

An earner's spouse or survivor, and in some cases children and parents, may be eligible for a SSB based on the earner's earnings history. In this paper, since I am looking at older couples, I ignore the child and parent benefits since they are unlikely to apply.

A spouse's primary benefit amount is 50% of the earner's PIA. If the spouse claims the benefit before his or her full retirement age, this amount is reduced by 8.33% per year for the first 3 years and an additional 5% per year for any earlier years. The spouse is not credited for delayed claiming. A spouse is eligible to claim a benefit on the earner's earnings history, only if the earner has also claimed Social Security benefits and the spouse is at least 62. Therefore, for a spouse who claims the benefit at age 62 and has a full-retirement age of 67, the maximum reduction is 35%. A spouse can only have the better of the spouse's benefit and his or her own benefit.

A survivor's primary benefit amount is the greater of 82.5% of the earner's PIA or the SSB the earner would be eligible for if he or she was alive.[19] A survivor may claim the SSB as early as age 60. If the survivor claims the benefit before his or her full-retirement age, this amount is reduced by the number of months before his or her full-retirement age divided by the total number of months between his or her full retirement age and age 60 times 28.5%. Therefore, regardless of full-retirement age, the maximum reduction is 28.5% for a widow who claims the benefit at age 60.

There are complex ways of claiming benefits that can increase the lifetime benefit levels of dual income couples that involve suspending benefits. For simplicity, I do not model the choice to suspend one's benefits after the normal retirement age.[20]

4.2.3 Benefit Taxation and Reduction

Social Security benefits can be taxed or reduced in four ways: the earnings test for early claimers who continue work, income taxation, the windfall elimination provision, and the government pension offset. In calculating the decision rules and simulating my model, I will only account for the

---

[19]If the earner was entitled to delayed retirement credits, then the survivor would receive the higher benefit level after accounting for these credits. Alternatively, if the earner had claimed his or her benefit early, then the survivor would receive the lower benefit level. The ability of the benefit reduction to impact the survivor is capped at 17.5% (the equivalent of claiming 31.5 months before the earner's normal retirement age).

[20]Claim and suspend is only available for individuals born after 1937, which is a small portion of my estimation sample since his or her spouse would have to be born in 1931-35 to be included in the estimation sample.

earnings test and income taxation.

The earnings test applies to anyone who works after claiming Social Security benefits. Prior to 2000, anyone between the age of 62 and the normal retirement age who had claimed benefits would have their benefit reduced by $1 for every $2 earned above an exempt amount (the exempt amount was $14,160 in 2010). Between the normal retirement age and age 70, the reduction factor was $1 for every $3 earned. In 2000, the earning test was eliminated for earners above the normal retirement age.[21] Any benefits that are reduced or eliminated by the earnings test, are returned to the worker at his normal retirement age. This is best illustrated with an example. Suppose a worker claims his benefits at age 62, but continues working until his normal retirement age. If he earns enough income to have his benefits eliminated before his normal retirement age due to the earnings test, then his benefits at normal retirement age would be equivalent to the benefits he would receive had he claimed at his normal retirement age. The earnings test has been shown in previous studies using structural models to have a significant impact on older workers' incentives to work (French, 2005; van der Klaauw and Wolpin, 2008).[22]

Depending on the household's adjusted gross income, part of its Social Security benefits may be subject to the standard U.S. income tax. In 2012, for married individuals, incomes below $32,000 were exempt, 50% of the total SSB was taxable for incomes from $32,001 to $44,000, and 85% of the total SSB was taxable for incomes above $44,000. These taxable amounts, unlike most of Social Security's provisions, are not indexed to inflation, implying they will become more binding over time. Further detail surrounding the taxing of Social Security benefits is included in Appendix A, with the discussion of how all taxes are accounted for in my model.

The windfall elimination provision and the government pension offset pertain to benefit reductions for individuals who have non-covered pensions. Currently, I am not able to distinguish between covered and non-covered pensions, and, therefore, do not include the windfall elimination provision and the government pension offset as part of my estimation.

4.3 Pensions

There are two major types of pension plans made available to employees in the United States, defined benefit and defined contribution retirement plans. Defined benefit plans (DB) pay a monthly benefit once the earner has claimed benefits and the investment risk is borne by the employer. Alternatively, defined contribution plans (DC) are accounts that an employer and employee can pay

---

[21]It was also changed to $1 for every $3 earned in the year in which the earner reaches full retirement age, with a higher exempt amount in that year. This was done for individuals, if their birthday occurred late in the year, who would have their Social Security benefit eliminated by the earnings test if they claimed on their birthdate, because of high earnings in the months prior to claiming the benefit.

[22]Using a reduced form model, Gruber and Orszag (2003), find no robust influence of the labor supply decision on men, but some suggestive evidence for women.

into (e.g. IRAs, 401(k) or 403(b) accounts), and then the employee is able to manage the account, and the investment risk is borne by the employee. Many employers have developed combination plans which have both a DB and DC component, but these are generally managed separately, the DC plan by an external investment agency (e.g. ING, Fidelity, etc.) and the DB plan managed by the employer or someone contracting with the employer (e.g. State Teacher Retirement Systems, unions, etc.) that absorbs a portion of the investment risk.

The HRS collects information from study respondents about whether or not they have a pension plan and, upon an affirmative response, will approach the respondent's employer to collect the pension plan's summary plan description (SPD). SPDs were coded into a pension calculator produced by the HRS and made available to researchers.

In Appendix F, I describe the two types of pension plans, as well as a few additional technical assumptions. For the purposes of the model presented here, the HRS pension calculator is used to predict the benefit level upon leaving the firm for any period following baseline. DB pension benefits are treated as income for tax purposes. Individuals who reported having a DB plan but for whom there was not a SPD are dropped from the sample. Defined contribution plans are converted to post-tax savings at baseline, and are treated as post-tax savings in subsequent periods.[23]

---

[23]Appendix F describes this conversion in greater detail. This is done primarily as a simplification because I do not retain separate state variables for pre- and post-tax savings.

5   Estimation

In this section, I introduce the estimation strategy for the model. I use a two-step estimation method that is increasingly common in the life-cycle literature (Gourinchas and Parker, 2002; French, 2005). First, key parameters that can be identified from the data are estimated (i.e. health transition rates, mortality transition rates, and earnings profiles), and others, such as the growth rate of assets, are calibrated.

In the second step, using the first step estimates and calibrations, $\hat{\chi}$, I estimate the preference parameters, θ, using method of simulated moments (MSM). Due to complexity of the model, I can not solve for θ directly, but instead calculate the optimal decision rules for a given "guess" of θ, which I will refer to as $\hat{\theta}$. Using the optimal decision rules for $\hat{\theta}$, I simulate life cycle profiles of households' labor supply, benefit claiming, and savings decisions. I then match moments observed from the data (generated by the true θ), to their counterpart moments from the simulation model (generated by $\hat{\theta}$). I iterate on this process until the model matches the data moments as closely as possible. Identification of the model's parameters is heavily dependent on the choice of moment conditions, which are discussed in §5.2. Further details about the econometric and computational procedures are specified in Appendix B - D.

5.1   First Step

The model presented in §**??** describes how a household makes choices across time and between consumption and leisure, but does not specify how individuals' earnings are determined and evolve over time, nor how households transition between uncertain states of health, mortality, and medical expenses. In this section I describe how I estimate part-time and full-time earnings paths for each member of the household, how I use observed HRS data to estimate transitions between uncertain states, and how I use subjective questions to estimate discrete preference types that capture unobserved differences in household's preferences for own and joint leisure. These intermediate "sub-models" are assumed to be true when solving the decision rules to estimate the preference parameters in the second step.

5.1.1   Annual Earnings

Earnings are known to the individual (i.e. there is no wage uncertainty). Assuming an individual is working at baseline, he or she may continue to receive the same level of nominal annual earnings in perpetuity.[24] The assumption of constant wage growth is necessary to remain consistent with

---

[24]This implies, given the assumption that inflation is 2%, that the real value of annual earnings fall by 2% per year. This is equivalent to the observed (negative) real wage growth rate of continuing workers from the sample used in the model's estimation.

the predicted defined benefit paths in the HRS pension calculator.

Every individual, regardless of whether he or she is working at baseline, may choose to work in a full-time or part-time non-baseline job. The evolution of earnings for full-time non-baseline jobs is determined from using a fixed-effect regression on a quartic in age and quadratic in firm tenure. The initial non-baseline earnings are determined from the residual of the fixed-effect regression, or, if that information is missing, is estimated from the individual's lifetime earnings (via the AIME), education, race, and baseline wage (if it exists). A separate, but similar, procedure is followed for estimating part-time earnings. A detailed description of the non-baseline earnings estimation process is included in Appendix G.

5.1.2 Health and Mortality

Following other papers in the literature (Rust and Phelan (1997); Blau and Gilleskie (2008, 2006)), I assume that health takes one of two discrete states: good or bad. I consider an individual in good health if he or she reports being in either good, very good, or excellent health; otherwise, if he or she reports poor or fair health, I treat the individual as being in poor health.

I estimate per period transitions using a logit model, where the probability of transitioning, π_{ij}, from state $i \in \{\text{good,bad}\}$ to state $j \in \{\text{good,bad}\}$ is a function of the individual's age, and previous health status. Obviously future health depends on current health, and it is well known that different ages and genders have higher propensities for poorer health.

Similarly, I estimate per period transitions from life to death using a logit model, where individual i's probability of surviving to period $t+1$ conditional on surviving to period t, s_{t+1}^i, is a function of the individual's age and previous health status. Since individuals have information about their health when making their labor supply decision, the estimated probability of mortality must be accounted for when making forward-looking projections of income flows.

The transitions between health states as well as from life to death are as expected: rising in age, and more favorable for women. A more detailed graphical analysis is provided in Appendix H.

5.1.3 Medical Expenses and Insurance

Respondents report whether or not they have access to health insurance through their current employer and whether that insurance continues into retirement. The HRS also identifies if the respondent's spouse has insurance coverage and whether that persists in retirement. Therefore, I identify three possibles states for health insurance: retiree coverage, no coverage, and tied coverage (i.e. insurance coverage that only exists as long as the employee continues to work). I assume that if one household member has health insurance, then they both have health insurance.

In the model, stochastic medical expenses are realized after the household's labor supply choice. Medical expenses are assumed to be log normally distributed. The mean and standard deviation of medical care expense are estimated conditional on the household's health status, access to health insurance, work status, and age, with a discontinuity at age 65 to capture Medicare eligibility. I include details about how the medical distribution is calculated in Appendix I.

Due to computational concerns, I model medical expenses only as a transitory shock to income, which will have the effect of biasing the precautionary savings incentive downward, thus reducing an individual's attachment to the labor force. There will be some persistence in medical expenditures because I model persistence in health status, which will affect medical expenses.

5.1.4 Preference Types

Households can vary based on characteristics that will be reflected in their preference for consumption versus leisure, but are not otherwise captured by the typical state variables. For this reason I include a finite number of discrete preference types, as in Keane and Wolpin (1997), van der Klaauw and Wolpin (2008), and French and Jones (2011), to capture heterogeneity in preference for own and joint leisure.

My model is estimated on the HRS cohort of households that are married in 1992, with one member born between 1931 and 1935, implying that many of the long-term decisions of the household are established (i.e. who works, how much is saved, how much time is spent together). I allow for households to vary by how its members value their own and joint-leisure to account for the fact that the model will not capture household formation and bargaining prior to when the household is first observed at baseline (i.e. 1992).

The preference for own-leisure is determined, as in French and Jones (2011), by questions such as "Even if I didn't need the money, I would probably keep on working" and questions about how much each individual enjoys his or her job. The second source of heterogeneity is likewise determined by questions regarding if the couple enjoys time together, looks forward to joint retirement, and who controls the family finances. I convert the responses to these questions, asked in 1992, into binary measures and include them in predicting the husband and wife's labor force participation after 1998, while controlling for the state variables in the model (i.e. age, health, assets, earnings, health insurance, Social Security benefit level, private pension levels, and marital status). For each individual, the own-leisure preference index is the sum of the work preference coefficients multiplied by their respective independent variables, and similarly for the spousal (or joint) preference index. The household's work or spousal preference index is simply the equally weighted sum for each household member's respective preference indices. By partitioning the indices at each measures' median, the index is converted into a binary measure (i.e. high and low) of the household's preference for own-leisure or joint-leisure.

I observe that a high preference for own-leisure is positively correlated with earnings, assets, AIME, defined-benefit pension flows, and negatively correlated with health. A high preference for spousal leisure is positively correlated with assets and health, but negatively correlated with earnings and AIME. An "out' preference index is created for households who were not asked the work questions in the first period because they were not working. As noted in Table 1, the initial distribution consists of 17.4% of the "out" preference type and a relatively even distribution between the four other preference types. In Appendix J, I describe the questions in detail and provide additional information on how the preference index is calculated.

The subscript τ represents different preference types based on preferences for own and joint leisure. If model parameters vary only based on preference for joint leisure, they are denoted $\tau(s)$. Since household preference heterogeneity is expected to affect consumption, time, own-leisure, and joint leisure, I allow the parameters that directly augment these to vary my preference type (i.e. α_τ, δ_τ, $\gamma_{i,\tau}$, $\beta_{i,\tau(s)}$, $\beta_{i,SP,\tau(s)}$, and $\beta_{i,SFT,\tau(s)}$). As $\beta_{i,SP,\tau(s)}$, and $\beta_{i,SFT,\tau(s)}$ reflect the effects of joint leisure, I allow these to only vary based on household preference types pertaining to spousal leisure in order to ease the computational burden.

5.1.5 Remaining Calibrations

I calibrate the real growth rate of assets, r, to 4%, and normalize the endowment of leisure, L, to 4. I choose this endowment of leisure because it implies that full time work is equivalent to a quarter of the leisure endowment. A quarter of leisure endowment falls between the annual equivalent, $\frac{2000 \text{ hours}}{8760 \text{ hours}}$, and the daily equivalent $\frac{8 \text{ hours}}{24 \text{ hours}}$ for full-time work. Finally, I set \$1 of consumption in a two person household to be equivalent to \$1.50 of consumption in a widowed household, $C_{single} = 1.5 \times C_{married}$. This is done because households may benefit from economies of scale, and this ratio reflects the implicit economies of scale assumed by the Social Security Administration when handling single versus dual household benefits through the Supplemental Security Income program.

Similar to other papers in this literature, I set a maximum age for claiming benefits and working, age 70, and a maximum lifespan of 110 to reduce the computational burden.[25]

5.2 Second Step (Moment Conditions & Identification)

The purpose of the MSM is to find the simulated moments that approximately match the same moments calculated from the observed data. In this section, I specify which simulated moment conditions I match to moment conditions from the observed data in the HRS sample, and discuss

---

[25] Age 70 corresponds to the last age where Social Security benefits are adjusted for delayed retirement. According to the U.S. Bureau of Labor Statistics (Toossi (2012)), 2010 male [female] labor force participation between 70-74 was 22.0% [14.7%], and between 75-79 was 14.5 [8.2%].

how they will identify the model's parameters. The full set of preference parameters include: $\theta = \{\alpha_\tau, \delta_\tau, \kappa, \theta_B, c_{min}, \gamma_{i,\tau}, \sigma_H, \sigma_W, \rho_{HW}, \ \beta_{i,\tau(s)}, \beta_{i,age}, \beta_{i,health}, \ \beta_{i,SP,\tau(s)}, \beta_{i,SFT,\tau(s)}\}$, where $\theta \in \Theta$ and $\Theta \subset \mathbb{R}^{48}$.

I divide any moments using household assets into thirds to capture the dispersion of assets in the data. The moment conditions which are matched include:

1. Mean assets by tertile, for the first two "thirds" (thirds $\times$ age $= 2 \times 12$ moments),

2. Share of households within each asset tertile by preference type, for the first two "thirds" ($\tau \times$ thirds $\times$ age $= 5 \times 2 \times 12$ moments)

3. Labor force participation by preference type, ($\tau \times$ sex $\times$ age $= 5 \times 2 \times 12$ moments)

4. Percent working full-time, conditional on working, excluding first preference type which does not work in the first period, $((\tau - 1) \times$ sex $\times$ age $= 4 \times 2 \times 12$ moments)

5. Labor force participation by health (health status $\times$ sex $\times$ age $= 2 \times 2 \times 12$ moments)

for a total of $34 \times 12 = 408$ moments.[26] The technical details of how these moments are calculated, the MSM, the optimization algorithm, and the calculation of the standard errors are included in Appendix D.

Households vary at baseline by their potential earnings, accumulated assets, spousal age difference, race, and many other factors fixed at baseline based on previous decisions. While there is not space to discuss the identification for each of the model's 48 preference parameters, I provide an argument for identification, using α_τ and δ_τ as examples. For the remaining parameters, I indicate where I expect the primary sources of identification.

Consider households A and B, identical except for the fact that in household A the couple is the same age, and in household B the wife is 10 years younger. Variation in these two households' savings will identify the willingness of the household to substitute consumption across time (i.e. α_τ), because household B will find it necessary to consume less and save more to account for the extended lifetime of the wife (moment cases (1) and (2)). If households highly value a smooth rate of consumption over time, then we would expect a large α_τ. Alternatively, the discount rate (δ_τ) affects the instantaneous utility, a composite of household consumption and the husband and wife's leisure, so it is identified by variation across time from households with the same consumption and leisure choices. If households' instantaneous utility decreases over time, then $\delta_\tau < 1$. Unlike models of infinitely lived households, the discount rate can exceed 1 if the household values higher levels of future instantaneous utility.

---

[26]I exclude the highest asset tertile because these households, with an average of over $800,000 in combined assets, are likely to be very sensitive to the rate of return, which is fixed in this model.

The preference parameter for leisure, $\gamma_{i,\tau}$, for gender $i \in \{H, W\}$ and household preference τ, is identified by variation in participation and full-time work (moment cases (3) and (4)). The time-invariant household bargaining parameter, $\beta_{i,\tau(s)}$, weights i's leisure relative to household consumption and is identified by variation in how households weight each member i's leisure relative to consumption when making decisions. Variation in how the household members weight consumption versus leisure over time identifies $\beta_{i,age}$. Finally, the joint retirement parameters, $\beta_{i,SFT,\tau(s)}$ and $\beta_{i,SP,\tau(s)}$, are identified by time-invariant variation in husband and wife's preference for own leisure based on the other's leisure choice.[27]

The bequest parameters and the minimal consumption level are determined by the upper and lower asset quantiles respectively from moment cases (1) and (2), because they are treated as both time and preference invariant.

The last 48 moment conditions help to identify the impact of health by gender on the relative value for leisure, $\beta_{H,health}$ and $\beta_{W,health}$. Finally, the variance and covariance of the fixed effects by gender, σ_H, σ_W and ρ_{HW}, are identified by time-invariant individual variation not otherwise described by the model.

5.3 Parameter Estimates

Using the procedure specified above, I estimate the model using the subsample of the married households from §4.1, specifically those households where one member was born between 1931 and 1935. The remainder of the sample is used in §5.5 to provide an out-of-sample test of the model based on the parameter estimates. Individual labor supply varies across the life-cycle due to changes in preference for leisure, $\beta_{i,age}$, increased risk of falling into bad health or dying, and spousal labor supply decisions.

Table 2 presents the parameter estimates and their standard errors. Recall that α_τ represents constant relative risk aversion with respect to consumption. High values of α_τ imply that a household is highly risk averse and hence does not want to substitute consumption across time. As a result, it is willing to consume less today if it can be guaranteed the same level of consumption tomorrow. Conditioning on the discount rate, a high α_τ can shift consumption across time and lead to precautionary savings. I would expect this to be particularly important for this sample because older individuals are at risk for substantial medical expenses, and risk averse agents would stockpile assets to guarantee a specific level of consumption in every period. The estimates in table 2 show values for α_τ between 2.81 and 3.15, which is consistent with estimates for the CRRA coefficient with respect to consumption, commonly found in the macro literature on consumption

---

[27]Consider the husband's return from his wife working full-time, $\beta_{H,SFT,\tau(s)}$, within joint preference type $\tau(s)$. All else constant, $\beta_{H,SFT,\tau(s)}$ is identified by variation in the husband's willingness to work when his wife moves from full-time work to either part-time or no work.

smoothing. It is lower than 3.72-7.27 found by French and Jones (2011), and much greater than estimates typically found in the structural retirement literature, such as 1.072 (Rust and Phelan, 1997), 1.26 (Gustman and Steinmeier, 2005), and 1.59-1.67 (van der Klaauw and Wolpin, 2008).[28] I believe this results from my choice to match asset holdings across time and modeling both husband and wife: the data indicates that many households accumulate assets over their 60s, building up large stockpiles of assets. This pattern is hard to match without significant risk aversion. Of the papers that match moments based on asset measures (e.g. van der Klaauw and Wolpin (2008); French and Jones (2011)), my estimates fall in between. Furthermore, since I use respondent data from 10 interview waves of the HRS, my estimation method will put more weight on the ability to describe asset accumulation at older ages (as compared to 3 waves in van der Klaauw and Wolpin (2008)).

Households discount future flows of expected instantaneous utility, described in (3.1), by δ_τ. Specifically, δ_τ acts as a temporal weight on combined utility in period t relative to period $t+1$. For example, if $\delta_\tau = 1$, then the individual values utility today the same as the utility tomorrow (conditioning on survival). Alternatively, if $\delta_\tau < 1$, then the household's utility will decrease over its life-cycle if it does not face liquidity constraints because utility (and hence consumption) is valued more today. Alternatively, if $\delta_\tau > 1$ this implies that a household weights utility tomorrow more relative to today, which is possible if a finitely lived household demands more utility in old age. The estimates in table 2 show that δ_τ ranges from 0.890 to 0.942, which is consistent with existing values found in the literature and implies significant heterogeneity in the population in rates of time preference.

Each individual in a household earns diminishing returns from leisure based on $\gamma_{i,\tau}$, which can be interpreted as the willingness to spread leisure across time. A lower $\gamma_{i,\tau}$ implies an individual is more responsive to changes in earnings, and is more willing to substitute leisure across time. Alternatively, a higher $\gamma_{i,\tau}$, common for men, indicates labor supply is unresponsive to changes in earnings. My results support that women's labor supply is more responsive than men's labor supply, as in Blundell and MaCurdy (1999).

In a model with no joint leisure ($\beta_{i,SFT,\tau(s)} = 0$ and $\beta_{i,SP,\tau(s)} = 0$), the Frisch elasticity of labor supply is given by $\frac{1}{\gamma_{i,\tau}}$. I consider what my estimates imply about the Frisch elasticity of labor supply if I ignore the joint leisure term, since the rareness of a factor that accounts for non-separability in spousal leisure prevents a more meaningful comparison. My estimates of $\gamma_{i,\tau}$ for men [women] indicate that it falls between 1.56 [1.01] and 1.77 [1.23], implying that the Frisch Elasticities without joint leisure are between 0.56 [0.81] and 0.64 [0.99]. For men, these

---

[28]There is also a separate literature using behavioral questions from the HRS to determine the CRRA coefficient. Barsky et al. (1997) find risk aversion to be very heterogenous across the population, with many people being very risk averse. Correcting for measurement error, they find the mean CRRA to be 12.1.

Table 2: Preference Parameter Estimates

Parameters based on type

Preference Type (Own L.,Joint L.)	Type 0 (Out)	Type 1 (High , Low)	Type 2 (Low , Low)	Type 3 (High , High)	Type 4 (Low, High)
α_τ Consumption	3.1480 (0.0924)	2.8592 (0.0085)	2.8193 (0.0096)	2.9502 (0.0102)	2.8736 (0.0082)
δ_τ Discount Rate	0.9072 (0.0205)	0.8903 (0.0079)	0.9242 (0.0095)	0.9414 (0.0089)	0.9013 (0.0083)
$\gamma_{H,\tau}$ Leisure	1.7676 (0.1173)	1.5762 (0.0521)	1.6042 (0.0666)	1.7080 (0.0492)	1.5685 (0.0440)
$\gamma_{W,\tau}$ Leisure	1.2338 (0.0913)	1.0051 (0.0682)	1.0065 (0.0246)	1.0595 (0.0343)	1.1624 (0.0518)
$\beta_{H,\tau(s)}$ Leisure Weight	-18.8057 (0.6725)	-19.8134 (0.1032)		-19.9252 (0.1237)	
$\beta_{W,\tau(s)}$ Leisure Weight	-19.7558 (1.4704)	-19.7589 (0.1018)		-20.2805 (0.1207)	
$\beta_{H,SP,\tau(s)}$ Participation	-0.0910 (0.8783)	-0.0203 (0.0015)		-0.0201 (0.0010)	
$\beta_{H,SFT,\tau(s)}$ Full-time work	-0.0661 (0.7060)	-0.1411 (0.0089)		-0.0817 (0.0039)	
$\beta_{W,SP,\tau(s)}$ Participation	-0.0698 (0.0023)	-0.0055 (0.0005)		-0.0222 (0.0014)	
$\beta_{W,SFT,\tau(s)}$ Full-time work	-0.0845 (0.2974)	-0.0857 (0.0071)		-0.1224 (0.0042)	

Row labels at far left (rotated): "Externality of Spouse's Decision on" — "Husband" (rows β_{H}) and "Wife" (rows β_{W}).

Parameters common to all types

$\beta_{H,age}$ Husband's Age-60	0.1852 (0.0039)	κ Bequest Shifter	297,050 (3465)	
$\beta_{W,age}$ Wife's Age-60	0.1904 (0.0046)	θ_B Bequest intensity	114,364 (2708)	
$\beta_{H,health}$ Husband's Health	1.1037 (0.0262)			
$\beta_{W,health}$ Wife's Health	0.9233 (0.0367)	c_{min} Consumption Floor	5667 (70.59)	
σ_{ε_H} Husband's Leisure shock	0 n.e.	ρ_{HW} Corr. between	0 n.e.	
σ_{ε_H} Wife's Leisure shock	0 n.e.	Husband and Wife's Leisure shock		

Degrees of Freedom   363

$q\left(\hat{\theta},\hat{\chi}\right)$ 2552.6

Notes: n.e. = not estimated. Degrees of Freedom = 408 moments in MSM procedure - 45 preference parameters. Parameters are estimated using method of simulated moments - see Appendix D for technical details.

are generally higher than estimates found using panel studies of male labor supply in the micro-labor literature, which typically fall between 0 and 0.5. Additionally, these values are less than the range of Frisch elasticities usually necessary to capture aggregate volatility in macro labor models, which usually fall between 2 and 4 (see Peterman (2012) and Blundell and MaCurdy (1999) for useful surveys). For women, the estimates for $\gamma_{W,\tau}$ are similar to U.S. studies using the Panel Study of Income Dynamics, which have found values closer to 1 (Trieste, 1990; Hausman, 1981). The differences for men may be explained in part by the older sample used here, since previous studies have typically focused on younger men who tend to always work.

Beyond the Frisch elasticity of labor supply, gender variations within a household occur based on age and health status (i.e. $\beta_{i,age}$ and $\beta_{i,health}$). If older and sick individuals value leisure more, then we would expect that $\beta_{i,age}, \beta_{i,health} > 0$, which is confirmed by my results.

Additionally, I allow each spouse to exhibit an external influence on the individual's return from leisure based on whether the spouse is participating in the workforce, $\beta_{i,SP,\tau(s)}$, or, working full-time, $\beta_{i,SFT,\tau(s)}$. If $\beta_{i,SP,\tau(s)}, \beta_{i,SFT,\tau(s)} < 0$, then the individual considers his or her spouse's leisure as complementary, implying if one's spouse takes more leisure-time, then the individual will also take more leisure-time. The point estimates indicate that men in general find their wives' leisure time to be complementary for their own leisure. Women from households with a low joint leisure preference type find their husbands' labor force participation to have little effect on their own preference for leisure. Similarly, Gustman and Steinmeier (2004), using a similar model and assuming no uncertainty, find that while the husband values joint retirement, the wife is indifferent.

Finally, permanent and unobserved changes arising from differences between individuals in a marriage are captured by the realization of ε_i, but as of this draft, estimates for the variance and correlation between ε_H and ε_W have not been completed. These will be included in a future revision.

The consumption floor, c_{min}, is primarily identified by the lower asset quantiles, and is considered time-invariant. It represents the household's guaranteed per period consumption as a result of government welfare plans. The estimate of \$5,667 is below \$7,687, which is the annual value of 2012 SSI benefits for a couple (discounted to 1992 dollars). It is not surprising for this value to be lower. As in Hubbard et al. (1995), the consumption floor affects all portions of the asset distribution because households fear the uncertainty of substantial medical expenses late in life that would make this constraint binding. A consumption floor below SSI levels may indicate loss aversion, an additional disutility of ending up in a bad state due to significant medical expenses, that is not otherwise captured by the model.

Finally, as in De Nardi et al. (2010), the bequest parameters θ_B and κ represent the bequest intensity and a bequest shifter, respectively. Using these parameters, and comparing the marginal utility of consumption in the last period to the discounted marginal utility from the bequest, I derive

a marginal propensity to bequeath of 0.98. Moreover, $\kappa \cdot (\theta_B \cdot \delta)^{-\frac{1}{\alpha}}$, represents the minimal flow of per period assets where the bequest motive begins to impact the individual's consumption choices. This implies a very low level where the bequest motive becomes effective of around \$5,850. Taken together, this implies that the bequest motive in this model is very strong, driving many households to save.

5.4 Model Fit (In-sample)

The over-identification test in table 2 rejects the model at the 1% level (428.6). The model, however, is able to capture important details of the household lifecycle such as the gradual decline in labor force participation among both sexes, with pronounced labor force exit at age 62. It is also able to capture asset accumulation across the population when the husband is in his 60s. Finally, the model captures phenomena observed in the data that are not matched as part of the estimation: the twin peaks of labor force exit at ages 62 and 65 (as in Gustman and Steinmeier (1986)), the large claiming of Social Security benefits at age 62, and the joint retirement of dual-career households.

In this section, I report the moment profiles from both actual and simulated data to develop an understanding of how well the model matches the moments specified in the estimation process. For moments that are matched for all preference types, I only include graphical illustrations of households that have a high preference for their own-leisure and low preference for joint-leisure (the rest are included in Appendix K). Additionally, I analyze non-matched moment profiles for SSB claiming based on HRS's linked Social Security claims histories, male labor force exit rates, and the prevalence of joint retirement.

5.4.1 Matched Profiles

Figure 4 reports the data and simulated moments for asset quantiles. The data indicate that household assets rise from age 58-69, as individuals save assets for retirement. The mean asset level for a household in the lower third of the asset distribution at age 62 is approximately \$80,000, while the average asset level for someone in the second third of the asset distribution is slightly more than double that amount, at about \$200,000. Note that the highest third have assets that grow very quickly, from an average of approximately \$800,000 at age 62 to over \$1.2 million at age 69. The model is able to match the means of the first two asset quantiles well. These are the only ones that are matched because the third quantile is sensitive to extreme asset values in the HRS data, something the model is not well-equipped to capture because the real rate of return is fixed at 4%.[29]

Household types have different propensities to save as noted in Table 2. As a result, the

---

[29]Figure 17 in Appendix K provides a closer view of the first two thirds to verify that they do indeed match well.

Figure 4: Asset Quantiles (by thirds) by Male Age

Asset Quantiles: Data and Simulations
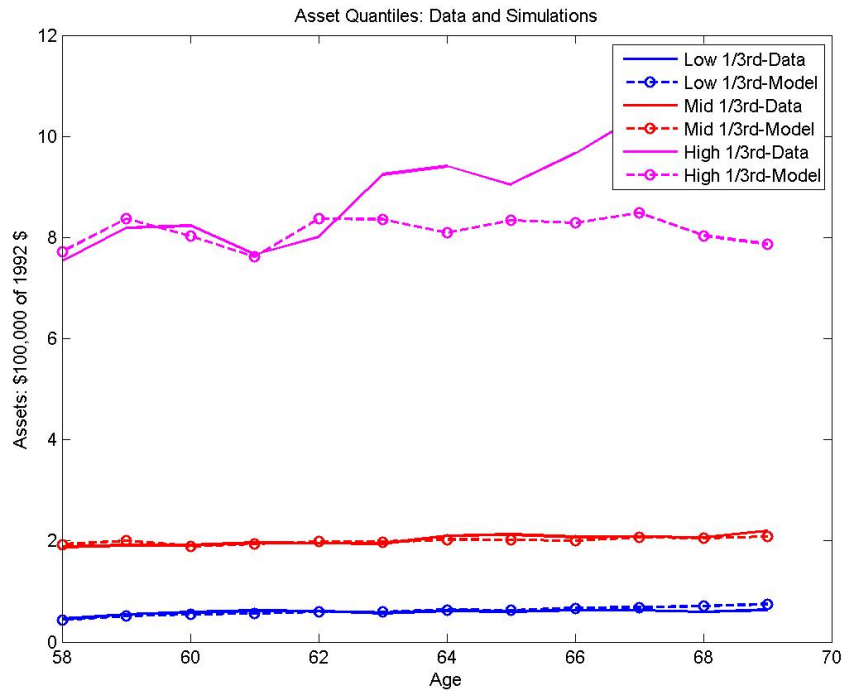
share of each household type in a given quantile varies. For households with a high preference for their own-leisure and low preference for joint-leisure (i.e. type 1), the household needs to accumulate assets in order to retire early. A disproportionate share of households therefore fall into the middle and top thirds of the asset distribution. Figure 5 reveals that the share of this household type in a given asset quantile is matched well by the simulation. Similar descriptive relationships exist for the other preference types, such as households that have high preference for joint-leisure (i.e. types 3 and 4) are more likely to have lower incomes. Figure 18, in Appendix K, reveals that the model does a good job of matching the asset distribution of the overall sample as divided by preference type.

Next, I consider the impact of preference type on individual labor force participation. In figure 6, both men and women's labor force participation by age is charted for type 1 households. For men in households with a high preference type for own-leisure, the men are far more likely to exit the labor force after the age of 61. This results in a sharp 40-45% decrease in labor force participation between age 62 and 69 for these types. For women, the decline in labor force participation is less dramatic, because fewer women are working to begin with. The model over predicts initial participation, and predicts a sharper than observed exit between ages 60-64. Figures 19 and 20, in Appendix K, reveal that, among all the preference types, the model is able to predict the downward trends in participation for men and women over their 60s.

Figure 5: Asset Quantile Shares by Preference Type
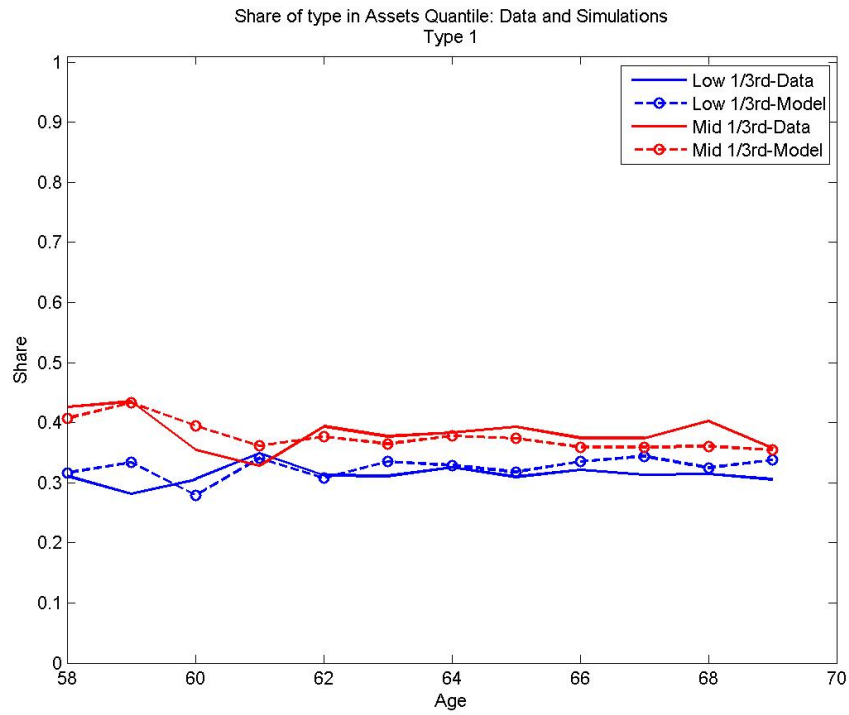Type 1 (High Preference for Own Leisure, Low Preference for Spousal Leisure)

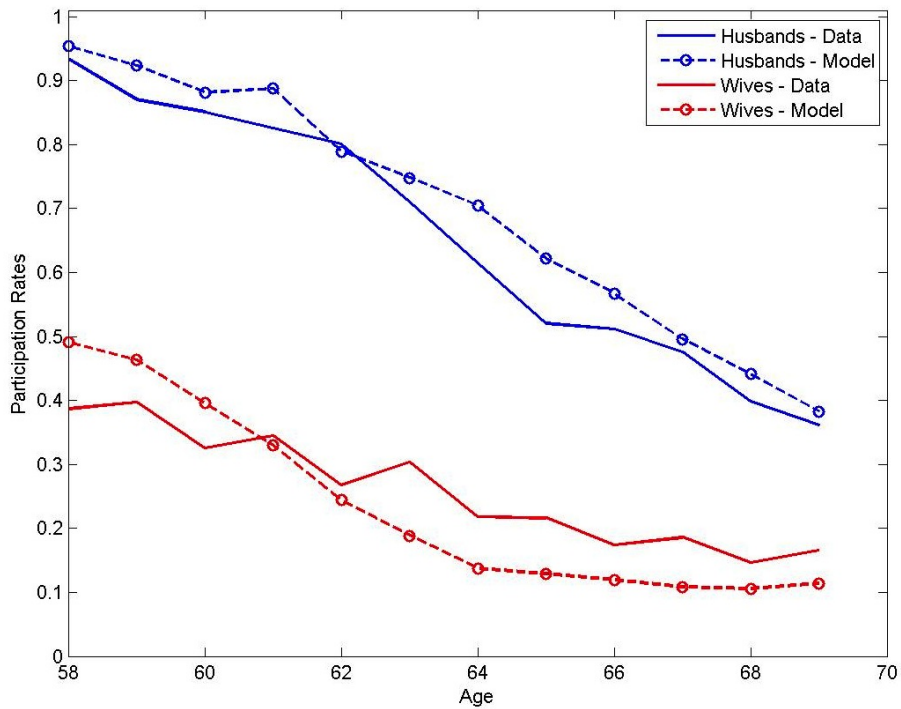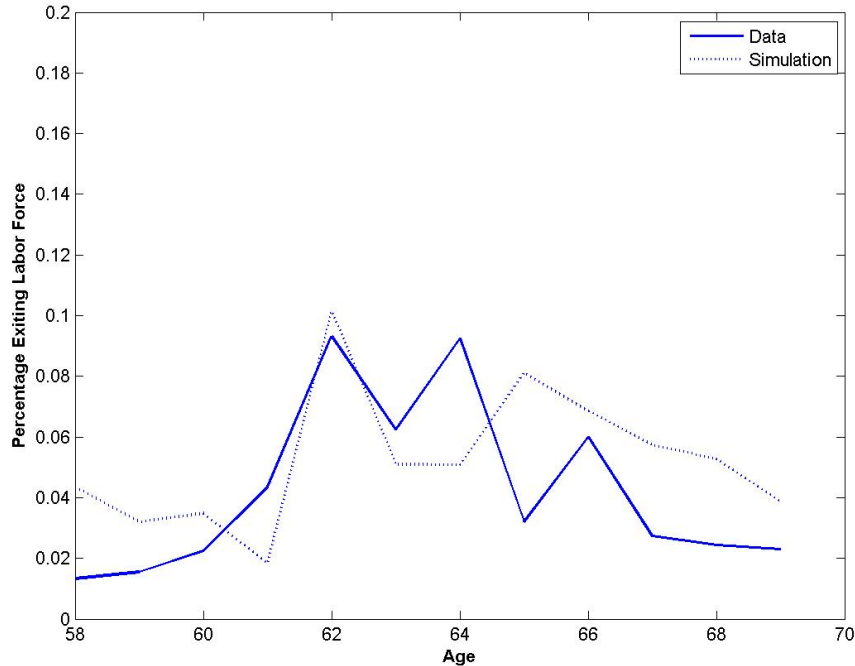Figure 6: Participation by Preference Type

Figure 7: Male Labor Force Exit by Age

The remaining moments on full-time work and health's influence on participation are captured in figures 21, 22, and 23 in Appendix K. The model has the most difficulty reproducing the employment composition of the workforce (i.e. part-time versus full-time), which could reflect my coarse discretization of labor supply - using 2 discrete states instead of hours. As expected, figure 23, in the appendix, indicates that individuals in bad health are more likely to not work.

5.4.2 Unmatched Profiles

Specifically excluded from the estimation procedure was matching any moments directly related to labor force exit, benefit claiming, or joint retirement. As mentioned in the introduction, a puzzling aspect of retirement behavior is that there exist spikes in retirement at ages 62 and 65, the early and normal retirement ages for Social Security (Gustman and Steinmeier, 1986). The model does not include an age-specific preference parameters, meaning that spikes in retirement ages can only arise from the structure of the constraints. As demonstrated in figure 7, my model is able to reproduce the spikes in labor force exit at these ages.

The HRS has access to administrative data from the Social Security Administration, which I use to judge the performance of the model in matching Social Security claiming rates. Figure 8 shows the benefit claiming history of the men and women in our sample. In comparing the benefit claiming rates in figure 8 to figure 6, it is clear that male participation rates remain high, despite

over 50% of the male population claiming their benefit at age 62. The model correctly predicts that a substantial amount of the population claims at 62, and that women are more likely than men to claim at age 62. This means that my model can rationalize claiming at age 62 through economic incentives, contrary to most of the literature on Social Security benefit claiming (i.e. Coile et al., 2002; Shoven and Slavov, 2013; Sass et al., 2013). It also correctly captures that very few individuals in my sample actually delay claiming beyond age 65.

Figure 8 shows that the model over-predicts the number of individuals who claim benefits at age 62. The model also captures some benefit claiming at age 65, but not nearly to the degree observed in the data. This is due to the disproportionate number of people claiming at age 62. It is possible that the differences are the result of the earnings test, which may discourage people with imperfect knowledge of Social Security from claiming until it expires at age 65. This would not happen in the model, because individuals perfectly understand the Social Security rules, including the fact that if benefits are taxed away through the earnings test, then they are returned in future benefits in an actuarially fair way.[30] The model overpredicts the number of men and women affected by the earnings test, which is expected because of the large number of early claimers in the simulations.

Finally, figure 9 shows the year difference in labor force exit rates for the subset of households where both individuals are working full-time at baseline. A positive value in the figure indicates that the husband retires after the wife. The model reproduces a spike where husband and wife exit the labor force at the same time, but it cannot capture the entire magnitude of joint labor force exit in the data. This may be due in part to the biannual nature of the HRS survey (from which I imperfectly approximate annual labor force participation), whereas the model predicts annual decisions.

Overall, the model is able to capture an impressive array of empirical regularities which are not included in the model's estimation but are able to be reproduced by the structure of the model.

5.5 Model Fit (Out-of-sample)

Similar to French and Jones (2011), I use a subsample of the HRS population that was not used in the estimation to confirm the out-of-sample fit of predicted results. The results presented in §5.3 use the sample of the HRS where one individual in the household was born between 1931 and 1935. In this validation exercise, I use the sample of the HRS population where at least one member of the household was born between 1937 and 1941. Some households may overlap between the

---

[30]Individuals who claim at 62 but continue to work until after 65 will have their benefits reduced or eliminated until age 65. The reduced or eliminated benefits do increase the final benefits in such a way as to make the final benefits of someone claiming and continuing to work from 62-65 to have the same benefits as if they had not claimed. Since the model defaults to claiming if it is indifferent, this may increase the claiming rate at age 62 relative to the normal retirement age.
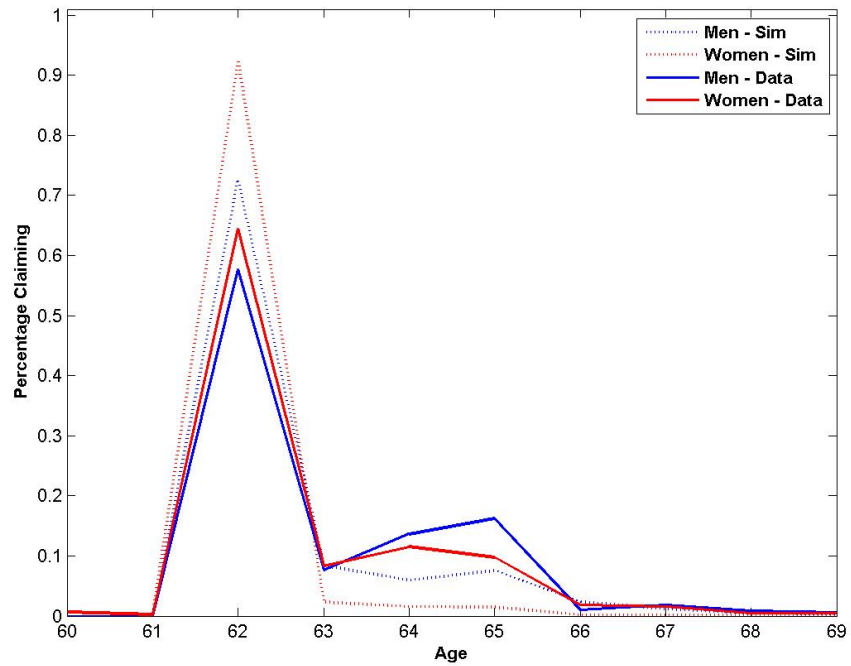
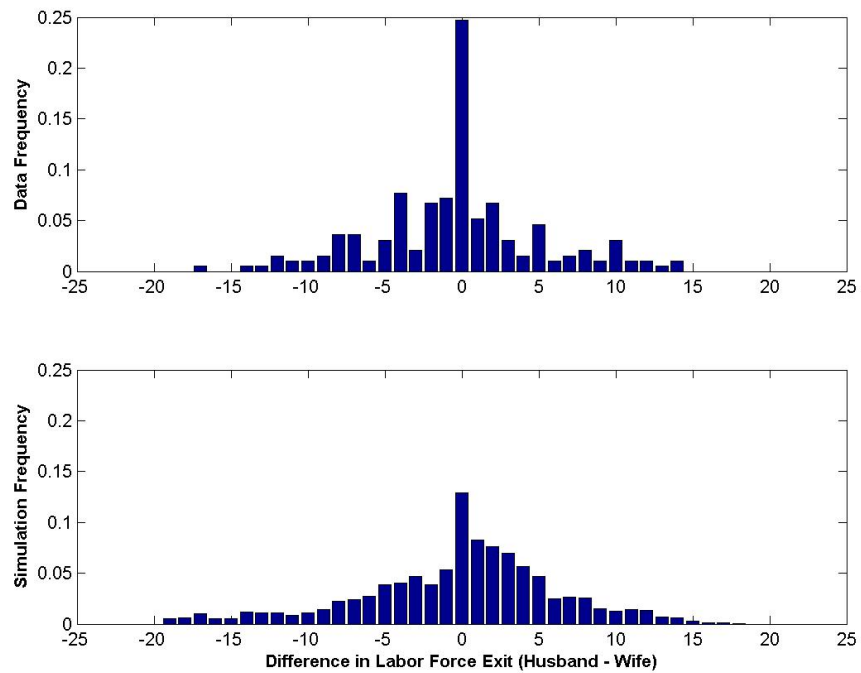Figure 8: Social Security Claiming Rates

Figure 9: Joint Labor Force Exit for Dual Career Households

two samples, but the two samples generally do not use the same households. Since the estimation sample was between the ages of 66 and 70 in 2001, when the earnings test was eliminated, and had smaller delayed benefit increments, their labor force participation rates are lower relative to the out-of-sample cohort.

Table 3a reports labor force participation rates for both the estimation sample and the out-of-sample group. The model columns report the predicted labor force participation of husbands and wives respectively for each sample. The *over-prediction of the model* columns report the difference between the model's prediction and the rate observed in the data. The model over-predicts labor force participation for husbands, and under-predicts labor force participation for wives. My model captures the higher participation rates of both men and women in the out-of-sample group. In the out-of-sample cohort, the over-prediction of male labor force participation relative to the data is amplified (i.e. $0.63 > 0.31$) and similarly for the under-prediction of female labor supply (i.e. -0.50 < -0.35). The differences are primarily from households with a low preference for joint leisure (results available from author). The amplification could be driven by differences in how cohorts value own and joint-leisure since the differences are driven primarily by specific preference types.

Table 4b reports Social Security benefit claiming rates for both the model and the data. The baseline columns reports the predicted benefit claiming of husbands and wives respectively for the estimate sample. The *Difference with Baseline* columns report the difference between the benefit claiming by the estimation sample (i.e. 1931-35 cohort) and the out-of-sample group (i.e. 1937-41 cohort). Husbands in the younger cohort delay claiming at age 62 by approximately 4 more percentage points relative to the older cohort. Wives in the younger cohort increase claiming at age 62 by approximately 2 percentage points. While the model is able to replicate the decrease in husband's claiming at age 62 and the increase in wives' claiming at age 62, the changes are significantly smaller than what is observed in the data.

Table 3: Comparison of model fit in-sample versus out-of-sample

(a) Participation Rates by Subsample and Sex

	Husbands				Wives			
	Baseline (1931-1935)		Out-of-Sample (1937-41)		Baseline (1931-1935)		Out-of-Sample (1937-41)	
Age	Model	Over-prediction of model	Model	Over-prediction of model	Model	Over-prediction of model	Model	Over-prediction of model
58	0.86	0.04	0.90	0.03	0.59	0.03	0.67	0.07
59	0.83	0.03	0.89	0.06	0.55	0.00	0.60	0.01
60	0.80	0.03	0.87	0.08	0.49	0.00	0.58	0.03
61	0.77	0.05	0.85	0.10	0.44	-0.02	0.51	-0.01
62	0.67	0.02	0.75	0.06	0.37	-0.05	0.41	-0.06
63	0.63	0.03	0.71	0.07	0.31	-0.05	0.35	-0.08
64	0.59	0.08	0.64	0.06	0.27	-0.07	0.32	-0.08
65	0.51	0.05	0.57	0.04	0.24	-0.04	0.27	-0.10
66	0.45	0.03	0.53	0.06	0.22	-0.03	0.25	-0.07
67	0.39	0.00	0.48	0.03	0.19	-0.03	0.22	-0.07
68	0.35	-0.02	0.45	0.04	0.17	-0.05	0.20	-0.06
69	0.31	-0.03	0.38	0.00	0.15	-0.05	0.16	-0.08
Total (58-69)	7.16	0.31	8.03	0.63	4.00	-0.35	4.53	-0.50

	Husbands				Wives			
	Model		Data		Model		Data	
Age	Baseline	Difference with Baseline	Baseline	Difference with Baseline	Baseline	Difference with Baseline	Baseline	Difference with Baseline
62	0.727	-0.002	0.577	-0.040	0.923	0.001	0.645	0.026
63	0.083	-0.001	0.076	0.004	0.023	-0.007	0.084	-0.031
64	0.060	0.005	0.136	0.022	0.015	-0.01	0.115	-0.001
65	0.075	0.006	0.162	0.047	0.015	-0.006	0.097	0.009
66	0.024	0.004	0.010	-0.008	0.002	0.002	0.018	-0.012
67	0.012	-0.007	0.018	-0.009	0.001	0.001	0.016	-0.007
68	0.004	-0.001	0.008	-0.006	0.001	0.001	0.005	0.004
69	0.004	-0.002	0.006	-0.004	0.001	0.002	0.005	-0.002

(b) Claiming Rates by Subsample and Sex

Notes: Table A reports the model's predicted labor force participation rates for the stated samples. The "over-prediction of model" columns represent how much greater the respective rate is compared to the sample's observed rate. Table B reports the cohort differences in benefit claiming rates predicted by the model and observed in the data. The "Difference with Baseline" columns represent how much greater the benefit claiming rate is in the in-sample cohort relative to the out-of-sample cohort.

6 Policy Experiments

Using the preference parameter specification estimated in §5.3, I conduct three counterfactual policy experiments: (1) the impact of eliminating or reducing the spouse's benefit, (2) the impact of eliminating or reducing both the spouse and survivor's benefits, and (3) the impact of making Social Security more progressive as proposed by the 1994-96 Social Security Advisory Council. In each case, I will examine the policy's predicted effect on (i) household labor supply, (ii) household benefit claiming, and (iii) the amount of Social Security benefits paid on average per contributing earner.

6.1 Elimination of the Spouse Benefit

Using the preference parameters in §5.3, I first, using the original sample of households (i.e. those where at least one individual was born between 1931-35), simulate eliminating the Social Security spouse's benefit. I then repeat the same exercise reducing the spouse's benefits by 50%. In tables 4a and 5a, I report the change in labor force participation from ages 58-69 for men and women. Figures 10a and 10b show this information graphically. The average gain in labor force participation for women with the elimination of the spouse's benefit is an additional 0.078 years of work (4 weeks), over ages 58-69. The effect is smaller but still positive, 0.062 additional years, when the benefit is reduced by 50%. Surprisingly, the effects for men are larger and negative. If the spouse's benefit is eliminated [reduced by 50%] then male labor force participation decreases by 0.11 [0.07] years or 5.5 [3.4] weeks. This suggests that the substitution effect from lower returns to work dominates the income effect (see figure 3).

The reduction or elimination of the spouse's benefit has a differential impact by household types. For women, elimination of the benefit causes households with a low preference for own-leisure and high preference for joint-leisure to increase their labor supply by 0.14 years of work. For men, the effect is largest for households with a low preference for own and joint leisure (type 2) where work falls by 0.21 years.

Households also change their claiming behavior in response to the elimination or reduction of the spouse's benefits (see tables 4b and 5b). Following the elimination of the spouse's benefit, a small percentage of men (3.4%) and women delay claiming (5.3%), mostly to the latest possible retirement age. Alternatively, when the benefit is reduced by 50%, claiming behavior is largely unaffected.

Since benefit claiming and work histories change because of this policy experiment, I can estimate the change in lifetime Social Security benefits. I estimate the average benefits paid per contributing earner by simulating each household through the end of life. Using the simulated claim and work histories, I can estimate the direct reduction in benefits paid through lower benefit

Table 4: Policy Experiment Results - Men

(a) Change in Average Labor Supply

Age	Baseline	Reduce Spouse's Benefits by 100%	by 50%	Reduce Sp. & Surv. Benefits by 100%	by 50%	Increase SS Progressivity
58	0.8644	0.8593	0.8595	0.7871	0.8606	0.8647
59	0.8301	0.8253	0.8253	0.7604	0.8377	0.8413
60	0.7953	0.7895	0.7928	0.7014	0.8029	0.8032
61	0.7701	0.7659	0.7671	0.6857	0.7756	0.7744
62	0.6708	0.6680	0.6694	0.6049	0.6988	0.6986
63	0.6325	0.6261	0.6296	0.5854	0.6602	0.6612
64	0.5856	0.5764	0.5823	0.5428	0.6182	0.6209
65	0.5084	0.4922	0.5003	0.4782	0.5488	0.5559
66	0.4494	0.4401	0.4438	0.4384	0.4998	0.4998
67	0.3905	0.3772	0.3835	0.3939	0.4297	0.4347
68	0.3525	0.3405	0.3427	0.3442	0.3873	0.3874
69	0.3058	0.2890	0.2935	0.3002	0.3307	0.3349
Avg. Years Worked (58-69)	7.1554	7.0495	7.0897	6.6228	7.4501	7.4770
Difference		-0.1059	-0.06570	-0.5326	0.2946	0.3216
Average Years Worked between 58-70 (Difference with Baseline)						
Type 0		0.0190	-0.0003	0.1194	0.0266	0.0164
Type 1		-0.0732	-0.0293	-0.8957	0.3484	0.3454
Type 2		-0.2063	-0.1204	-0.7149	0.1402	0.2499
Type 3		-0.1940	-0.1271	-0.9649	0.3952	0.4358
Type 4		-0.0669	-0.0534	0.0625	0.4639	0.4708
Asset Quantile 1		-0.0483	-0.0405	-0.9304	0.1564	0.1675
Asset Quantile 2		-0.2028	-0.1273	-0.4900	0.2961	0.4018
Asset Quantile 3		-0.0589	-0.0235	-0.2483	0.3723	0.3500

(b) Change in Percentage Claiming at a given age

Age	Reduce Spouse's Benefits by 100%	by 50%	Reduce Sp. & Surv. Benefits by 100%	by 50%	Increase SS Progressivity
62	-0.0338	-0.0099	-0.0489	-0.0665	-0.0532
63	-0.0094	-0.0104	-0.0318	-0.0188	-0.0088
64	-0.0062	0.0015	-0.0193	-0.0084	-0.0067
65	0.0186	0.0124	0.0481	0.0572	0.0478
66	0.0024	-0.0014	0.0096	0.0125	0.0107
67	0.0064	0.0065	0.0152	0.0154	0.0037
68	0.0015	0.0015	0.0056	0.0067	0.0051
69	-0.0008	-0.0003	0.0008	0.0017	0.0009
70	0.0213	0.00010	0.0208	0.00030	0.0003

Notes: The first column of table (a) reports the percentage of the observed sample participating in the labor force at each age (58-69) as predicted by the simulated model. The standard errors for men at each age are less than 0.0020. The average years worked is calculated by summing the averages of each age between 58 and 69. As a result, standard errors cannot be produced for the average years worked, or the differences with the baseline model for each of the policy experiments.

Table 5: Policy Experiment Results - Women

(a) Change in Average Labor Supply

Age	Baseline	Reduce Spouse's Benefits		Reduce Sp. & Surv. Benefits		Increase SS
		by 100%	by 50%	by 100%	by 50%	Progressivity
58	0.5943	0.5936	0.5929	0.6899	0.6171	0.6142
59	0.5468	0.5500	0.5496	0.6440	0.5824	0.5740
60	0.4943	0.5011	0.5005	0.6067	0.5345	0.5219
61	0.4436	0.4449	0.4477	0.5527	0.4898	0.4761
62	0.3731	0.3840	0.3807	0.4858	0.4212	0.4099
63	0.3130	0.3287	0.3255	0.4365	0.3643	0.3463
64	0.2723	0.2798	0.2756	0.3857	0.3177	0.3012
65	0.2405	0.2501	0.2469	0.3600	0.2873	0.2663
66	0.2192	0.2261	0.2254	0.3251	0.2534	0.2387
67	0.1880	0.1946	0.1939	0.2889	0.2190	0.2093
68	0.1684	0.1731	0.1724	0.2657	0.2031	0.1929
69	0.1462	0.1517	0.1505	0.2282	0.1756	0.1667
Avg. Years Worked (58-69)	3.9997	4.0776	4.0616	5.2693	4.4653	4.3174
Difference		0.0779	0.0619	1.2697	0.4657	0.3178
Average Years Worked between 58-69 (Difference with Baseline)						
Type 0		-0.0059	-0.00020	1.5332	0.0424	-0.0022
Type 1		0.1008	0.0765	1.6577	0.6056	0.3958
Type 2		0.0772	0.0560	0.5913	0.4124	0.2918
Type 3		0.1472	0.1128	1.8492	0.6341	0.4198
Type 4		0.0565	0.0549	0.6780	0.5245	0.4059
Asset Quantile 1		0.1338	0.1014	0.9549	0.4602	0.3060
Asset Quantile 2		0.1015	0.0820	1.2185	0.6398	0.4332
Asset Quantile 3		0.0088	0.0091	1.5345	0.1991	0.1385

(b) Change in Percentage Claiming at a given age

Age	Reduce Spouse's Benefits		Reduce Sp. & Surv. Benefits		Increase SS
	by 100%	by 50%	by 100%	by 50%	Progressivity
62	-0.0526	-0.0073	-0.057	-0.023	-0.0166
63	0.0014	-0.0003	0.0119	0.0099	0.0073
64	0.0033	0.0031	0.0063	0.0041	0.0007
65	0.0067	0.0044	0.0091	0.012	0.0069
66	0.0013	0	0.0035	0.0013	0.0004
67	0.0011	0	0.0041	0.0017	0.0009
68	0.0008	0	0.0046	-0.0004	-0.0005
69	0.0007	0	0.0048	0.0009	0.0006
70	0.0373	0	0.0127	-0.0064	0.0004

Notes: The first column of table (a) reports the percentage of the observed sample participating in the labor force at each age (58-69) as predicted by the simulated model. The standard errors for women at each age are less than 0.0017. The average years worked is calculated by summing the averages of each age between 58 and 69. As a result, standard errors cannot be produced for the average years worked, or the differences with the baseline model for each of the policy experiments.

Table 6: Policy Experiment Results - Change in Benefits paid over lifetime

| Sex | Baseline | Reduce Spouse's Benefits | | Reduce Sp. & Surv. Benefits | | Increase SS |
		by 100%	by 50%	by 100%	by 50%	Progressivity
Men	$175,312.88	-$1,211.43	-$1,052.08	-$1,764.32	-$103.07	-$30,246.92
Women	$134,245.69	-$19,837.99	-$14,651.02	-$42,314.88	-$32,577.96	-$21,396.43

%Δ Due to

Men	Reduced Benefits	n.a.	n.a.	n.a.	n.a.	104.04%
	Changed Labor Supply					-4.04%
Women	Reduced Benefits	97.82%	99.27%	108.58%	104.38%	101.61%
	Changed Labor Supply	2.18%	0.73%	-8.58%	-4.38%	-1.61%

Notes: n.a. = not applicable because the change in lifetime benefits is too small to be informative. All monetary values in 1992 dollars.

levels, and the indirect effects of increased benefits from longer work histories and reduced benefits from the earnings test. While the elimination of the spouse's benefit should induce women to work more, it is only binding for lower earners implying the replacement rate on the spouse's own earnings will be high, mitigating the combined benefit of eliminating the spouse's benefit. The budgetary benefit will be largest if the women for whom the spouse's benefit is binding also work very little.

Table 6 presents the change in benefits paid over men and women's lifetimes. In the baseline case, I find that the average male will receive $175,313 in benefits, while the average female will receive $134,246 in lifetime benefits. The elimination of the spouse's benefit has a negligible effect on lifetime benefits for men, but decreases women's lifetime benefits by 14.8% if spouse's benefits are eliminated and by 10.9% if spouse's benefits are reduced by 50%.

Recall from the analysis of claiming and labor supply that women are induced to work more and claim later if the spouse's benefit is eliminated, but only work longer if it is reduced by 50%. In the second part of table 6, I separate the reduction in benefits into the direct effect from reducing benefits, and the indirect effect from changes to claiming and work behavior. I find that while most of the change is attributable to the direct impact of reducing benefits, the indirect effect varies based on whether benefits are reduced or eliminated. Regardless of how much of the benefit is reduced, the longer work history, induced by the benefit change, further reduces lifetime benefits.

6.2 Elimination of the Spouse and Survivor's Benefits

Similar to above, I first simulate eliminating the Social Security spouse and survivor's benefits. I then repeat the same exercise, but instead reduce these benefits by 50%.

The elimination of the survivor's benefit has a significant effect on lifetime benefits, particularly for women, and has large labor supply effects on both men and women. Eliminating the

Figure 10: Labor Supply Response to Social Security Benefit Changes

(a) Men

(b) Women
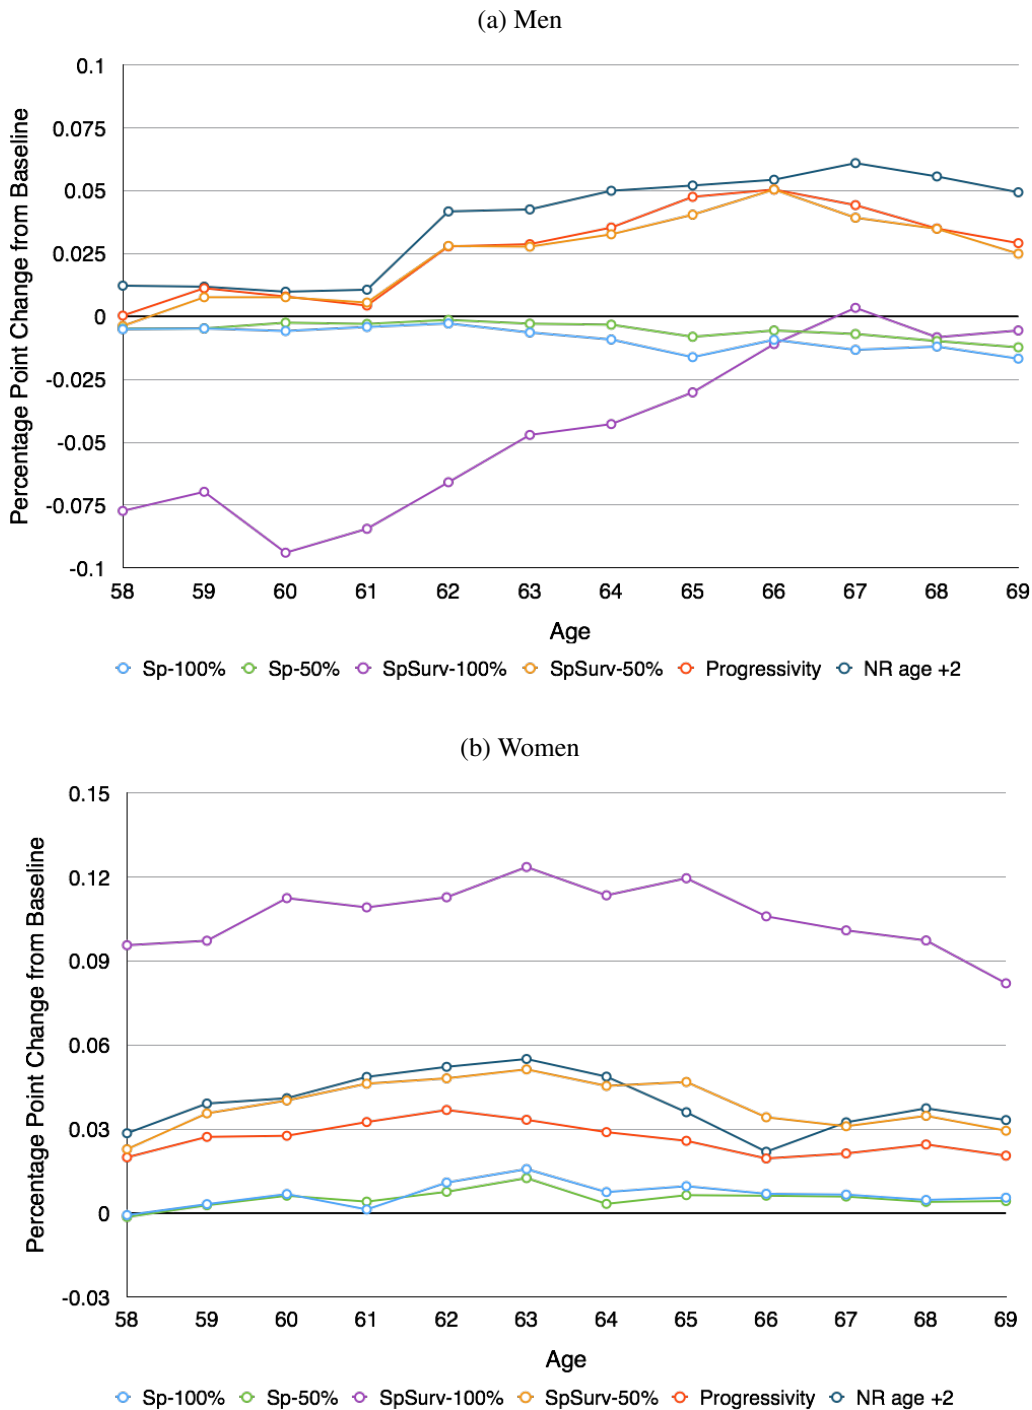
Notes: This figure presents results from counterfactual experiments of changes in Social Security's benefit structure based on the model of household savings, labor supply, and benefit claiming. The model incorporates uncertain longevity, health, and medical expenses. Sp-100% refers to eliminating the Spouse's Benefit. Sp-50% refers to reducing the Spouse's Benefit by 50%. SpSurv-100% refers to eliminating the Spouse and Survivor's Benefits. SpSurv-50% refers to reducing the Spouse and Survivor's Benefits by 50%. Progressivity refers to increasing the progressivity of the Social Security Benefit formula from 90-32-15 to 90-22.4-10.5 according to one of the proposals from the 1994-6 Advisory Council on Social Security. NR age + 2 refers to increasing everyone's Normal Retirement age by two years.

spouse's benefits causes female labor force participation in increase, on average, 1.27 years (66 weeks), while men decrease their labor force participation by 0.53 years (27.7 weeks). Interestingly, the model predicts that men's labor supply response is very different if spouse and survivor's benefits are reduced by 50%. Men *increase* their labor supply by 0.29 years (15.6 weeks) and women increase their labor supply by 0.47 years (24.2 weeks). The differential effect on men is driven by whether or not they are able to increase the annuity value of Social Security for their surviving spouse (See figure 10a). If they are unable to, as in the case of benefit elimination, then they are more likely to retire sooner - similar to when just spouses' benefits are reduced or eliminated. However, if the husband still has some ability to improve his surviving wife's income security through a higher spouse's benefit, as in the case when survivor's benefits are only reduced, then he will choose to work longer (i.e. the income effect would dominate the substitution effect).

Heterogeneity plays a large role in individual labor responses. Women who are out of labor force at baseline (type 0), return to the labor force to work, on average, another 1.53 years. Women with a low preference type for own leisure work longer, but the increase is less than one year. Changes to spouse and survivor's benefits also have a heterogenous effect on men's attachment to the labor force. Removal of spouse and survivor's benefits leads men with a high preference for own leisure or low preference for joint leisure (e.g. intuitively, the more selfish) to work 0.71-0.96 years (37-50 weeks) less. Moreover, changes in claiming behavior when both spouse and survivor's benefit is eliminated look similar to when only the spouse's benefit is eliminated, with the exception that claiming of benefits is only delayed to the normal retirement age instead of age 70.

Eliminating spouse and survivor's benefits will reduce lifetime benefits by 31.5% [1.0%] for women [men], while decreasing them by 50% still reduces lifetime benefits by 24.3% [0.1%].

6.3 Increased Progressivity of Social Security Benefits

As mentioned in §2, Social Security benefits follow a progressive formula that pays out a higher fraction of average indexed earnings to low income earners. In 2013, Social Security paid out 90% of a worker's first $9,492, 32% of the next $47,724, and 15% of the rest. Following one of the recommendations of the 1994-96 Social Security Advisory Council, I consider the impact of making this system more progressive by reducing the second replacement rate from 32% to 22.4% and the third replacement rate from 15% to 10.5%. This would have the effect of reducing an individual's annual benefit by approximately $4200, or 18.7% for a worker whose average indexed earnings were equivalent to $60,000 in 2013 dollars.

In tables 4a and 5a, the result is a dominant income effect, inducing men and women to work an additional 0.32 years between ages 58-69. However, the aggregate results actually disguise some very interesting variation by type. For couples with a high preference for own and

joint leisure (i.e. type 3), the effect is to increase labor force participation by 0.44 years for men and 0.42 years for women. Alternatively, for households with a low preference for own and joint leisure (i.e. type 2), the impact on labor force participation is only an additional 0.25 years for men and 0.29 years for women.

As reported in tables 4b and 5b, there is no substantial effect to changes in the claiming age for women, and 5% of men delay benefit claiming from age 62 until 65.

Similar to the previous two analyses, I analyze the direct reduction in average benefits paid, and the indirect benefit of increased Social Security revenues from longer work histories. As would be expected, increasing the progressivity of the Social Security's primary insurance amount leads to a reduction in lifetime benefits equal to 17.3% for men and 15.9% for women. The smaller change for women is likely due to smaller lifetime incomes.

6.4 Discussion

The previous three analyses permit a comparison of the budgetary impact from changes to the Social Security System. Because my model internalizes both benefit claiming and labor supply, I am able to separate how a change to the Social Security old-age and survivor benefits are likely to alter the program's funding along these two important margins.

Relative to one of the 1994-1996 proposals to improve the solvency of Social Security through increasing the progressivity of Social Security benefits, I find that eliminating both the spouse and survivor benefit would achieve 85.4% of the savings, while eliminating the spouse's benefit would achieve 40.8% of the savings (at least among the married and aged beneficiary population that my sample reflects). Put another way, eliminating both the spouse and survivor benefits reduces average lifetime Social Security payments to a household by 14.2% while just eliminating the spouse's benefit leads to a lifetime benefit reduction of 6.8% in my HRS sample.

Eliminating or reducing the Social Security Spouse and Survivor's Benefits cause a very large increase in women's labor force participation rates and encourages delayed benefit claiming. For men, reducing both benefits increase male labor force participation, while eliminating both benefits causes a large decrease in male labor force participation. Specifically, eliminating spouse and survivor benefits increases women's labor force participation by 1.27 years and decreases men's labor force participation by 0.53 years. Reducing both of these benefits by 50% increases women's labor supply by 0.47 years, and increases men's labor supply by 0.29 years.

These results highlight the importance of structural modeling in the context of complex life-cycle programs like Social Security. Since the above analysis does not specifically account for the additional income to SSA through payroll taxes, the reduction in male labor supply is significant. If males earn on average $28,175 (in 1992 dollars), then this implies losing out on $2,445 in payroll tax income. Failing to account for the impact of the policy change on labor force

participation would paint a better financial picture of the savings from these policy changes than would actually occur.

The model also points to a very important nonlinear relationship between benefit reduction and actual savings. Reducing the spouse's benefit by 50% achieves 74.6% of the savings that are achieved by eliminating the spouse's benefit. This nonlinear relationship comes from the fact that many women have at least small Social Security earnings histories, and by reducing the spouse's benefit, only those who were never eligible for Social Security benefits on their own earnings history would receive the spouse's benefit. Similarly, reducing both the spouse and survivor benefit by 50% achieves 74.1% of the savings that would be achieved from eliminating both benefits.

Finally, allowing for heterogeneity in the model, tables 4a and 5a show that the impact of benefit changes on labor force participation depend very much on the preference type and the asset levels of the different individuals at baseline. Changing only the spouse's benefit most significantly alters the labor supply of men in the middle of the asset distribution and women in the lower third of the asset distribution. Eliminating both the spouse and survivor's benefit most significantly alters the labor supply of men in the lowest third of the asset distribution and women in the highest third of the asset distribution. Alternatively, changes to the progressivity and normal retirement age (results available from author) primarily alter the labor supply of the middle part of the asset distribution. As the population shifts towards more dual income households, the prevalence of certain preference types is likely to rise, indicating that any change to Social Security can be better informed by accounting for the heterogeneity in household responses to proposed policy changes.

7 Conclusion

Social Security provides benefits to a worker's spouse and survivor that alter the work incentives of both household members. These benefits, despite being relatively small in size when compared to the rest of Social Security disbursements, are large relative to other federal government expenditures. In this paper, I construct a life-cycle model of household savings, labor supply, and benefit claiming decisions that accounts for health, survival, and medical expense uncertainty. The model allows me to answer the question of how altering Social Security's spouse and survivor benefits would change the work and retirement decisions of each household member.

Applying my model to a sample from the Health and Retirement Study, I confirm the theoretical implication that spouse benefits encourage the household's low income earner to work less and encourage the household's high income earner to work longer. For a household's high income earner, this implies that the increased return from work (i.e. the substitution effect) dominates the effect of higher benefit levels (i.e. the income effect). Among those households nearing retirement, I find that eliminating both spouse and survivor benefits cause wives, who are statistically the household's low earner, to increase their average labor force participation by 1.27 years, while decreasing husbands' labor force participation by 0.53 years. This effect is important because it implies that the impact of spouse and survivor benefits is large for both women and men. Furthermore, if the spouse and survivor benefit is reduced by 50%, then husbands increase work 0.29 years longer. The differential response of men to the elimination versus the reduction of both benefits suggests that a household values the option of increasing guaranteed annual income over the household's lifespan with the annuity provided by a Social Security benefit. If this option is taken away, as in the case of the elimination of the survivor benefit, the incentive for the high earner to work is significantly reduced.

Additionally, I demonstrate that there are positive but diminishing savings from reducing spouse and survivor benefits. Specifically, I show that reducing these benefits by 50% achieves about 74.1% of the savings that result from eliminating these benefits. The model demonstrates these nonlinear savings arise primarily due to the structure of Social Security benefits, with only a small impact due to changes in labor supply. The non-linearity in savings from auxiliary Social Security benefit reduction is important for policymakers to account for when considering any alterations to Social Security's auxiliary benefits.

8 References

AARP, "Assessing Current and Future Beneficiaries' Knowledge of Social Security Benefits," prepared by S. Kathi Brown and Rebecca Perron, 2011.

Altmeyer, Arthur, *Formative Years of Social Security*, University of Wisconsin Press, 1966.

Barsky, Robert B., F. Thomas Juster, Miles S. Kimball, and Matthew D. Shapiro, "Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study," *Quarterly Journal of Economics*, May 1997, *112* (2), 537–579.

Blau, David M. and Donna B. Gilleskie, "Health Insurance and Retirement of Married Couples," *Journal of Applied Econometrics*, 2006, *21* (7), 935–953.

__ **and** __ , "The Role of Retiree Health Insurance in the Employment Behavior of Older Men," *International Economic Review*, May 2008, *49* (2), 475–513.

Blundell, Richard and Thomas MaCurdy, "Labor Supply: A Review of Alternative Approaches," in O. Ashenfelter and D. Card, eds., *Handbook of Labor Economics*, Vol. 3, Elsevier Science, 1999.

Coile, Courtney, "Retirement Incentives and Couples' Retirement Decisions," *Topics in Economic Analysis and Policy*, 2004, *4* (1), Article 17.

__ , **Peter Diamond, Jonathan Gruber, and Alain Jousten**, "Delays in Claiming Social Security Benefits," *Journal of Public Economics*, June 2002, *84* (3), 357–385.

De Nardi, Mariacristina, "Wealth Inequality and Intergenerational Links," *Review of Economic Studies*, 2004, *71*, 743–768.

__ , **Eric French, and John B. Jones**, "Why Do the Elderly Save? The Role of Medical Expenses," *Journal of Political Economy*, February 2010, *118* (1), 39–75.

French, Eric, "The Effects of Health, Wealth, and Wages on Labor Supply and Retirement Behaviour," *Review of Economic Studies*, 2005, *72* (2), 395–427.

__ **and John B. Jones**, "The Effects of Health Insurance and Self-Insurance on Retirement Behavior," *Econometrica*, May 2011, *79* (3), 693–732.

Gourieroux, Christian, Alain Monfort, and E. Renault, "Indirect Inference," *Journal of Applied Econometrics*, December 1993, *8*, S85–S118.

_ **and** _ , *Simulation-Based Econometric Methods* CORE Lectures, Oxford University Press, 1996.

Gourinchas, Pierre-Olivier and Jonathan A. Parker, "Consumption over the Life Cycle," *Econometrica*, January 2002, *70* (1), 47–89.

Gruber, Jonathan and Peter Orszag, "Does the Social Security Earnings Test Affect Labor Supply and Benefits Receipt," *National Tax Journal*, 2003, *56* (4), 755–773.

Gustman, Alan L. and Thomas L. Steinmeier, "A Structural Retirement Model," *Econometrica*, May 1986, *54* (3), 555–584.

_ **and** _ , "What People Don't Know about their Pensions and Social Security: An Analysis using Linked Data from the Health and Retirement Study," in "Working Paper 7368" National Bureau of Economic Research September 1999.

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

_ **and** _ , "Social Security, Pensions and Retirement Behaviour within a Family," *Journal of Applied Econometrics*, 2004, *19*, 723–737.

_ **and** _ , "The social security early entitlement age in a structural model of retirement and wealth," *Journal of Public Economics*, 2005, *89* (2-3), 441–463.

Haider, Steven and Gary Solon, "Nonrandom Selection in the HRS Social Security Earnings Sample," in "Labor and Population Program, Working Paper Series 00-01" RAND February 2000.

Hausman, Jerry, "Labor supply, how taxes affect economic behavior," in H. Aaron and J. Pechman, eds., *Tax and the economy*, Brookings Institution, 1981.

Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes, "Precautionary Saving and Social Insurance," *Journal of Political Economy*, 1995, *103* (2), 360–399.

Hurd, Michael D., James P. Smith, and Julie M. Zissimopoulos, "The Effects of Subjective Survival on Retirement and Social Security Claiming," in "NBER Working Paper 9140" National Bureau of Economic Research September 2002.

Kapteyn, Arie, Pierre-Carl Michaud, James P. Smith, and Arthur Van Soest, "Effects of Attrition and Non-Response in the Health and Retirement Study," in "Labor and Population Program, Working Paper Series WR-407" RAND July 2006.

Keane, Michael P. and Kenneth I. Wolpin, "The Career Decisions of Young Men," *Journal of Political Economy*, June 1997, *105* (3), 473–522.

Maestas, Nicole, "Back to Work: Expectations and Realizations of Work after Retirement," *Journal of Human Resources*, 2010, *45* (3), 718–748.

Peterman, William B., "Reconciling Micro and Macro Estimates of the Frisch Labor Supply Elasticity," December 2012. Federal Reserve Board of Governors.

Ruhm, Christopher T., "Bridge Jobs and Partial Retirement," *Journal of Labor Economics*, 1990 1990, *8* (4), 482–501.

Rust, John and Christopher Phelan, "How Social Security and Medicare Affect Retirement Behavior in a World of Incomplete Markets," *Econometrica*, July 1997, *65* (4), 781–831.

Sass, Steven A., Wei Sun, and Anthony Webb, "Social Security Claiming Decision of Married Men and Widow Poverty," *Economics Letters*, 2013, *119*, 20–23.

Shoven, John B. and Sita N. Slavov, "Recent Changes in the Gains from Delaying Social Security," in "Working Paper 19370" National Bureau of Economic Research August 2013.

Social Security Administration, "Social Security Bulletin: Annual Statistical Supplement 2012," United States Government Printing Office 2012.

Stockholm International Peace Research Institute, "SIPRI Military Expenditure Database," 2012.

Tauchen, George, "Finitie State Markov-chain Approximations to Univariate and Vector Autoregressions," *Economics Letters*, 1986, *20*, 177–181.

Toossi, Mitra, "Labor force projections to 2020: a more slowly growing workforce," *Monthly Labor Review*, January 2012, *135* (1), 43–64.

Trieste, Robert, "The Effect of Income Taxation on Labor Supply in the United States," *Journal of Human Resources*, Summer 1990, *25* (3), 491–516.

U.S. Census Bureau, "State Government Finances," [http://www.census.gov/govs/state/] July 2013.

U.S. Department of Health and Human Services, "Temporary Assistance for Needy Families - Justification of Estimates for Appropriations Committees," 2012.

van der Klaauw, Wilbert and Kenneth I. Wolpin, "Social security and the retirement and savings behavior of low-income households," *Journal of Econometrics*, 2008, *145*, 21–42.

Appendix

A Tax Rates

I assume households jointly file their tax return if both individuals are alive, otherwise the household files as a single. The household pays federal, state, and payroll taxes on income from both household members. Income includes earnings from each individual's job, pension income, Social Security income, and asset returns (both defined contribution and savings).

In 1992, 50% of Social Security benefit income was taxed for jointly filing household with incomes over $32,000, and single filers with income over $25,000. In December 1993, the 50% threshold was kept in place, but a second bracket, 85% of Social Security benefit income, was added for households with incomes over $44,000 for joint filers and $34,000 for single, unmarried filers. In my analysis, I assume that the 1993 rules hold for every year. For example, if single John received $10000 in Social Security benefits and earned an additional $25,000 for part-time work, then $0.5\,(32000 - 25000) + 0.85(35000 - 32000) = \$6,050$ of his Social Securiy benefit would be taxable as income. However, John will never have more than 85% of his Social Security benefits taxed, implying that if he earned $50,000$ for his part time work, then only $0.85(\$10,000) = \$8,500$ would be taxable. Note that these rules and levels have not changes since 1993, and therefore are not indexed for inflation.

I use the IRS tax rules from 1992 and reported state tax rates in NBER's TAXSIM calculator.[31] I weight state tax rates by the U.S. Census's projections of population in each state in July 1992 for ages 50 and greater.[32] I assume that all individuals are not self-employed for tax purposes, meaning that he or she only pays half of the payroll tax. In table 7, I report the tax rates for married, jointly filing households and single households. For joint households, the 3rd tax bracket (ending at $55,500) represents the maximum Social Security contribution level. In table 7b, I assume that only one individual earns a total of $55,500. If both individuals are working, then the third through fifth tax brackets could change depending on each individual's earning levels. Notice that the addition of the fifth tax bracket between the two tables is due to the correspondence between the top income tax bracket and the maximum Social Security contribution level in 1992.

---

[31] See http://users.nber.org/~taxsim/ for more details

[32] See State Characteristics, 1990 to 1999 Annual Time Series of State Population Estimates by Age and Sex, 5-Year Age Groups by Sex at http://www.census.gov/popest/data/historical/1990s/state.html

B   Recursive Methods

Recursion is commonly used in structural models, but the typical design of a decision tree taught in standard game theory can be difficult or impossible to reproduce due to finite computational time. Often this can arise when decisions are continuous (such as how much to consume or save), when the number of periods covered are large, or when choices in each period require historical variables. While this list is not extensive, it does represent all the challenges faced in a life-cycle model of labor force participation and benefit claiming. There exist significant computational tradeoffs that must be considered when developing a structural model of this variety, and these can only be understood if the reader first has an appreciation for how the backward recursion is actually conducted and approximated.

First, it is currently impossible to come close to calculating an entire decision tree. Instead, it is approximated at each decision period by a discretized set of the state variables. In the model by French and Jones (2011), this is done with 9 state variables: (1) benefit application decision, (2) preference type, (3) whether or not there is a cost of reentering the labor force for that period (i.e. un-retire), (4) health insurance transition, (5) health status, (6) health care cost transition, (7) Social Security AIME level, (8) wage change, (9) asset level. Given their discretization, this implies $2 \times 3 \times 2 \times 3 \times 2 \times 3 \times 16 \times 5 \times 32 = 552,960$ state combinations. The calculation of decision rules through backward recursion is based (in theory) on the history of choices an agent has made up until the current period's decision node, but due to the continuous nature of state

Table 7: Taxes

(a) Single

Pre-Tax Income (1992 $)	Marginal Tax Rates			
	Federal	State	Payroll	Combined
0-3,600	0.00%	0.00%	7.65%	7.65%
3,601-25,050	15.00%	4.56%	7.65%	27.21%
25,051-55,500	27.86%	5.03%	7.65%	40.54%
55,500 +	30.68%	5.37%	1.45%	37.50%

(b) Married - Filing Jointly

Pre-Tax Income (1992 $)	Marginal Tax Rates			
	Federal	State	Payroll	Combined
0-6,000	0.00%	0.00%	7.65%	7.65%
6,001-41,800	15.00%	4.51%	7.65%	27.15%
41,801-55,500	28.00%	4.79%	7.65%	40.44%
55,501-92,500	27.83%	5.05%	1.45%	34.33%
92,500 +	31.34%	5.21%	1.45%	38.00%

Notes: For each household member above the age of 65, the income threshold increases by $900 for single households and $700 for married housedholds.

variables, such as assets, and the long history required for other state variables, such as Social Security's Average Indexed Monthly Earnings (AIME) measure, it becomes impossible to permit state variables to depend on prior decisions.

The calculation of AIME provides an excellent example of the challenges presented by backward recursion when future choices depend upon histories in addition to current states. AIME is calculated using the best 35 years of earnings. Even if we coarsely discretize potential earnings into 5 levels, and assume everyone started their significant earning at age 25, then just for the calculation of AIME at age 60, we would have $5^{35} \approx 2.91 \times 10^{24}$ possibles wage combinations required to calculate the AIME at age 60 (think about how bad it gets at age 61!). Instead, French and Jones (2011) take the AIME in the period as given, abstracting from the history that led to its level. Since what is relevant for the current decision is both the current AIME and the AIME for continuing to work, the lack of wage history requires the modeler to approximate AIME if the agent continues to work. In French and Jones' work, they use estimated replacement wages for the population based on age. For example, for an agent continuing to work at age 62, they assume that the current wage replaces 58.9% of one year's wages relative to the individual's AIME:

$$AIME_{t+1} = (1 + CPI) \cdot AIME_t - \frac{1}{35} \{0, W_t N_t - (0.589) \cdot (1 + CPI) \cdot AIME_t\}$$

In this setup, 58.9% is meant to approximate the ratio of the lowest earnings year to AIME. As the population gets older, the ratio approaches 100%, such that the AIME does not grow through replacement of the lowest earnings. At younger ages (before 55), they assume that entire years are replaced (i.e. that the ratio is 0%).

This setup presents two major challenges to a life-cycle model of labor supply and benefit claiming. First, it smooths out accruals in retirement programs (both Social Security and defined benefit pensions), possibly reducing or eliminating the incentive to delay claiming for some individuals. Moreover, since French and Jones' tie pension benefits to the AIME, it becomes less clear how to separate the actual effects of Social Security from pensions on labor supply outcomes. Second, since this setup only approximates AIME in the next period, it cannot account for any possible notches in benefit calculations that might exist from delayed claiming beyond a year. For example, an agent in good health who is replacing a zero earnings year in the AIME calculation for each additional year of work (e.g. a woman who took a decade off from the labor force) will not only have a significant incentive to delay claiming at 62, but may face a much larger incentive to delay claiming beyond a single year due to her likelihood of survival relative to Social Security's delayed benefit adjustments.

In this paper, I take an approach similar Gustman and Steinmeier (2005), where, in order to capture the AIME and pension calculations, I take the wage paths for a worker as given, allowing

AIME and pension benefits to be calculated directly. This requires calculating the decision rules for each individual in sample, and thus requires a simplification in the number of states to achieve computational feasibility. Moreover, it eliminates the feasibility of a modeler incorporating wage uncertainty into the model for fear of quickly increasing the computational burden. Choosing a fixed wages allows my model to reflect the institutional details of Social Security and individual pension plans, as well as be able to appropriately account for individuals' unique earnings histories. Since I estimate each household's decision rules separately, I use the husband and wife's earnings history at baseline to determine the rate at which lowest earnings are replaced.

C Numerical Methods

The recursive formulation of a household's value function is given by:

$$
\begin{aligned}
V_t\left(X_t\right) \;=\; \max_{C_{h,t},N_{h,t},B_{h,t}} \;\Big\{ & U\left(C_{h,t},L_{h,t}\right) + \delta_\tau \left(1 - s^h_{t+1}\right)\left(1 - s^w_{t+1}\right) b(A_{h,t+1}) \\
& + \delta_\tau \left(1 - s^h_{t+1}\right) s^w_{t+1}\, \mathbb{E}\left[V_{t+1}\left(X_{t+1} \mid X_t, t, C_{h,t}, B_{h,t}, N_{h,t}, \text{wife survives}\right)\right] \\
& + \delta_\tau s^h_{t+1}\left(1 - s^w_{t+1}\right) \mathbb{E}\left[V_{t+1}\left(X_{t+1} \mid X_t, t, C_{h,t}, B_{h,t}, N_{h,t}, \text{husband survives}\right)\right] \\
& + \delta_\tau s^h_{t+1} s^w_{t+1}\, \mathbb{E}\left[V_{t+1}\left(X_{t+1} \mid X_t, t, C_{h,t}, B_{h,t}, N_{h,t}, \text{both survive}\right)\right] \Big\}
\end{aligned}
\tag{C.1}
$$

subject to a non-negative borrowing constraint and the consumption floor. The solution to the recursive formulation requires solving for each household's consumption, labor force participation, and benefit claiming choices at every age at and after baseline (1992), collectively referred to as the decision rules. These decision rules are calculated numerically, using the model detailed in the paper because no closed form solution exists. No closed form solution exists for several reasons, including that future state variables depend upon the history of those variables, and that there are several discontinuities in the budget set arising from taxation, pensions, and Social Security Benefits.

The recursive formulation above is solved using value function iteration beginning in period T, assumed to be age 110, and solved backward to the first period. The vector of possible states is discretized into 13 state variables: (1) the husband's stochastic preference for leisure, ε_H , (2) the wife's stochastic preference for leisure, ε_H, (3) household marital status (this is important for widowhood), (4) household health insurance status, (5) husband health status, (6) wife's health status, (7) husband's Social Security PIA, (8) wife's Social Security PIA, (9) husband's pension level, (10) wife's pension level, (11) household asset level. Given the discretization, this implies $3 \times 3 \times 3 \times 3 \times 2 \times 2 \times 2 \times 2 \times 2 \times 2 \times 8 = 41{,}472$ state combinations solved for each of the 948 households in the sample. The time T decision rule is found by assuming that everyone knows they will die at period T, such that $V_T = U(C_T, 0) + \delta \cdot b(A_{T+1})$. For each set of state variables, X_T, we calculate the optimal consumption (and hence savings) decision for period T. This yields the value function at time T, which can be used in calculating V_{T-1} to find the decision rules at period $T-1$ according to the Bellman equation in (C.1). This process is repeated from period $T-1$ back to period 0, which in this model corresponds to the male household member's age at baseline (i.e. 1992).[33]

The value function is evaluated at each state combination and linear interpolation is used for continuous variables (i.e. assets, AIME, pension benefit, ε_H, and ε_W). Discretization is finer at lower levels of assets since I would expect greater responsiveness at lower levels to changes in

---

[33] The appendix on *Recursive Methods* goes into greater detail about how I handle approximating next periods assets given our discretization method.

asset accumulation. In my initial estimate process I keep the number of states for Social Security and pension benefits small (2 states each), but these states reflects the individual's worst and best possible benefits based on his or her own earnings history. In robustness checks, I will investigate whether the results are sensitive to this rough discretization.

Each period, the household chooses the level of consumption, labor supply, and benefit application that maximizes their discounted lifetime utility. Consumption is a continuous choice in the model, however, implying that for each state combination the household must determine the optimal level of consumption. Given the discrete nature of the other choice variables, there is no reason to expect the value function to be globally concave with respect to consumption. I discretize the consumption space into 36 choice states and allow the household to solve for V_T based on each choice state, from which the household will choose the level of consumption that maximizes its discounted lifetime utility. When the problem is solved again for period $T - 1$, the agent will test only a local range of consumption choices. As the backward induction process continues, the range of consumption states tested will depend upon the male household members' age, with a larger range being used during periods of critical life choices (e.g. age 65 when respondent reaches normal retirement age for Social Security). If a value on the boundary of the consumption range is chosen, then the range is expanded by three choice states in the direction of increasing utility until a local optimum is found.

Once the decision rules are calculated, the rules are then used to generate simulated household histories. 200 random outcomes of health, medical expenses, mortality, and unobserved individual heterogeneity ($\varepsilon_{i,t}$) are generated per household. Using each household's period 0 state vector, the household's decisions in period 0 are determined from the appropriate decision rules. When the state vector does not precisely lie on the discretized grid of state combinations, I use linear interpolation to approximate the household's decisions. Combining the states and decisions from period 0, I use the budget constraint and asset accumulation conditions from the model, in addition to the health, mortality and medical expense shocks after period 0, and the appropriate Social Security and pension rules for the household, to calculate the state vector in period 1. This process is repeated, creating a life-cycle history for the household. In generating the shocks, the actual (not discretized) value of the shock is used. In generated the pension and Social Security benefit levels, the actual earnings history and pension rules are used in the calculation.

In order to reduce the computational burden, individuals are assumed to claim their benefits and cease work by age 70.

D Moment Conditions and Method of Simulated Moments

D.1 Moments Conditions

This section is a more detailed description of §5.2 in the main text. For expositional clarity, I reproduce the moments cases as they are in the main text and describe the technical details of the moments that are matched.

I divide any moments using household assets into three quantiles to capture the dispersion of assets in the data. I match the following moment conditions for ages 58-69 ($T = \{58, 59, ...69\}$) for a total of $34T = 408$ moments.

1. Mean assets by quantile and men's age, for the lowest two quantiles ($2T$ moments)

   I divide any moments using household assets into three quantiles to capture the dispersion of assets in the data. The v_j percentage of households (h) with assets below $Q_{v_j}(A_{ht}, t)$ is defined as

   $$Pr\left(A_{ht} \leq Q_{v_j}(A_{ht}, t) \mid t\right) = v_j$$

   where the quantile index is denoted by j. Put another way, $Q_{v_j}(A_{ht}, t)$ is the v_jth age-condition asset quantile. The model analog to $Q_{v_j}(A_{ht}, t)$ is $\hat{Q}_{v_j}(t; \theta_0, \chi_0)$ from the simulated asset distribution. Note that $t = age_i$ is individual i's age, where here it is assumed to be the male's age ($i = H$). Let $\bar{A}_j(t, \theta_0; \chi_0)$ represent the model's prediction of the mean asset level observed in asset quantile j at age t. The implied conditional moment then becomes

   $$\mathbb{E}\left[A_{ht} \mid t, Q_{v_{j-1}}(A_{ht}, t) \leq A_{ht} \leq Q_{v_j}(A_{ht}, t)\right] = \bar{A}_j(t, \theta_0; \chi_0) \ .$$

   This can then be converted into an unconditional moment that can be estimated from the simulation results by rearranging the previous equation and plugging in for the model analogs:

   $$\mathbb{E}\left[A_{ht} - \bar{A}_j(t, \theta_0; \chi_0) \mid t\right] \times \mathbf{1}\left\{\hat{Q}_{v_{j-1}}(t; \theta_0, \chi_0) \leq A_{ht} \leq \hat{Q}_{v_j}(t; \theta_0, \chi_0)\right\} = 0 \ . \tag{D.1}$$

2. Share of a preference type's household population within each asset quantile by age (lowest two quantiles only) for men ($10T$ moments)

   Let $\bar{h}_j(\tau, t; \theta_0, \chi_0)$ represent the model's prediction of share of households, h, where the husband is $t = age_i$ years old in asset quantile interval j with preference type τ. If the model is true then:

   $$\mathbb{E}\left[h \mid Q_{v_{j-1}}(A_{ht}, t) \leq A_{ht} \leq Q_{v_j}(A_{ht}, t), t, \tau\right] = \bar{h}_j(\tau, t; \theta_0, \chi_0) \ .$$

Empirically when estimating the moment vector, $m(LC_i, \theta_0; \chi_0)$ (see next section), I convert this relationship into an unconditional moment equation:

$$\mathbb{E}\left[h - \bar{h}_j(\tau, t; \theta_0, \chi_0) \mid t, \tau\right] \times \mathbf{1}\left\{\hat{Q}_{v_{j-1}}(t; \theta_0, \chi_0) \leq A_{ht} \leq \hat{Q}_{v_j}(t; \theta_0, \chi_0)\right\} = 0 \qquad \text{(D.2)}$$

for asset quantiles $j \in \{1, 2\}$.[34] I exclude the share of the third asset quantile as the shares are constrained to add to one, and so it is identified by the other two moments.

3. Percent participating in the labor force by preference type, age, and sex ($10T$ moments)

Recall that each household, h, is comprised of two members of each gender $i \in \{H, W\}$ at baseline. $LFPR_{hit}$ represents i's labor force participation at $t = age_i$. I match the following unconditional moment for men and women by age:

$$\mathbb{E}\left[LFPR_{hit} - L\bar{F}PR_i(t, \tau; \theta_0, \chi_0) \mid t, \tau\right] = 0, \qquad \text{(D.3)}$$

where $L\bar{F}PR_i(t, \tau; \theta_0, \chi_0)$ is the model's prediction of average labor force participation for each gender with household preference type τ.

4. Percent working full-time, conditional on working, by preference type and sex (excluding first preference type which does not work in the first period - $8T$ moments)

Similar to case (3), each household, h, is comprised of two members of each gender $i \in \{H, W\}$ at baseline. FT_{hit} represents i's labor force status conditional on participation at $t = age_i$. If $\bar{FT}_i(t, \tau; \theta_0, \chi_0)$ represents the model's prediction of individuals working full-time conditional on participation for preference type τ at age t, then the implied conditional moment condition becomes:

$$\mathbb{E}\left[FT_{hit} \mid LFPR_{hit} = 1, t, \tau\right] = \bar{FT}_i(t, \tau; \theta_0, \chi_0).$$

I then convert this relationship to an unconditional moment condition:

$$\mathbb{E}\left[FT_{hit} - \bar{FT}_i(t, \tau; \theta_0, \chi_0) \mid t, \tau\right] \times \mathbf{1}\left\{LFPR_{hit} = 1\right\} = 0, \qquad \text{(D.4)}$$

which is used as $8T$ of the moment conditions. Note that I exclude the type where both individuals are out of the labor force at baseline, $\tau = 0$, because the moment condition may be empty for certain ages.

5. Labor force participation by individual health status, age, and sex ($4T$ moments)

---

[34] I define $\hat{Q}_{v_0}(age_i; \theta_0, \chi_0) = -\infty$ and $\hat{Q}_3(age_i; \theta_0, \chi_0) = +\infty$

As in case (3), I match labor force participation moments conditional on health status, $health \in \{good, bad\}$, and sex. Therefore the moment condition is:

$$\mathbb{E}\left[LFPR_{hit} - L\bar{F}PR_i\left(health, t, \tau; \theta_0, \chi_0\right) \mid t, \tau, health_{it} = health\right] = 0 .$$

I then convert this relationship to an unconditional moment condition:

$$\mathbb{E}\left[LFPR_{hit} - L\bar{F}PR_i\left(health, t, \tau; \theta_0, \chi_0\right) \mid t, \tau\right] \times \mathbf{1}\left\{health_{it} = health\right\} = 0 . \tag{D.5}$$

D.2   Method of Simulated Moments

Using the moment conditions discussed in the previous section, I use 408 moment conditions to over-identify the 48 preference parameters, denoted by θ. Let $m(\bullet)$ represents the moment condition based on observed life-cycle histories LC_i for individual i in household h, and let θ_0 represent the true value of the preference parameters θ, from the data generating process, χ_0. Note that the life cycle histories, LC_i, comprises all observables, including endogenous outcomes, exogenous or potentially endogenous state variables, X_t, and instrumental variables. Given the vector of moment conditions such that

$$\mathbb{E}\left[m\left(LC_i, \theta_0; \chi_0\right)\right] = 0 ,$$

then the generalized method of moments (GMM) estimator, $\hat{\theta}_{gmm}$ minimizes:

$$Q_n(\theta) = \left[\frac{1}{n}\sum_{i=1}^{N} m\left(LC_i, \theta_0; \chi_0\right)\right]' \mathbf{W}_n \left[\frac{1}{n}\sum_{i=1}^{N} m\left(LC_i, \theta_0; \chi_0\right)\right] ,$$

where $\mathbf{W}_n$ is the symmetric positive definite weighting matrix that does not depend on θ. Now if there is no closed-form solution for $m\left(LC_i, \theta; \chi_0\right)$ such that:

$$m\left(LC_i, \theta; \chi_0\right) = \int k\left(LC_i, u_i, \theta; \chi_0\right) g(u_i) du_i$$

then $m\left(LC_i, \theta; \chi_0\right)$ can be replaced by $\hat{m}\left(LC_i, u_{is}, \theta; \chi_0\right)$, an unbiased simulator, and u_i denotes s draws from the marginal density $g(u_i)$. The method of simulated moments (MSM) estimator $\hat{\theta}_{msm}$ instead minimizes:

$$\hat{Q}_n(\theta) = \left[\frac{1}{n}\sum_{i=1}^{n} \hat{m}\left(LC_i, u_{is}, \theta; \chi_0\right)\right]' \hat{\mathbf{W}}_n \left[\frac{1}{n}\sum_{i=1}^{n} \hat{m}\left(LC_i, u_{is}, \theta; \chi_0\right)\right] \tag{D.6}$$

where $\hat{m}(LC_i, u_{is}, \theta; \chi_0)$ is defined by the moment conditions in (D.1)-(D.5) above, and $\hat{\mathbf{W}}_n$ is the optimal weighting matrix from the simulated data. Following Gourieroux and Monfort (1996), as $n \to \infty$ and for a fixed number of simulations s, $\hat{\theta}_{msm}$ is both consistent and asymptotically normally distributed:

$$\sqrt{n}\left(\hat{\theta}_{msm} - \theta_0\right) \xrightarrow[n \to \infty]{d} N\left(0, \hat{\Upsilon}\right),$$

where:

$$\hat{\Upsilon} = \left(D'WD\right)^{-1} D'WSWD \left(D'WD\right)^{-1} \tag{D.7}$$

such that $D = \partial m(\cdot)/\partial \theta' \big|_{\theta=\theta_0}$ and $W = plim_{n \to \infty} \hat{W}$, which is estimated by:

$$\hat{W} = \left\{ \tilde{V}\left(m\left(LC_i, \theta; \chi_0\right)\right) + \frac{1}{s} \tilde{V}\left(\hat{m}\left(LC_i, u_{is}, \theta; \chi_0\right)\right) \right\}^{-1}$$

where $\tilde{V}(\cdot)$ is the estimated variance with respect to a larger simulation sample and S is the variance-covariance matrix of the simulated sample. Thus the first term represents the moment condition from the data with respect to the larger simulated sample, and the second term represents the moment condition with respect to the smaller simulation sample from which the estimates are selected. Note that the optimal choice of W, corresponds to $W = S^{-1}$, simplifying the asymptotic variance-covariance matrix to

$$\hat{\Upsilon} = \left(D'\hat{W}D\right)^{-1}$$

In practice, I use only the diagonal terms of $\tilde{V}\left(m\left(LC_i, \theta; \chi_0\right)\right)$ when calculating $\hat{W}$ in order to minimize (D.6). This is to ensure invertibility (non-singularity) and because S may be biased in small samples. When I calculate the standard errors of the preference parameter vector $\hat{\theta}_{msm}$ and test the moment conditions (i.e. over-identified restrictions of the model) against the zero restrictions implied by the model, I use equation (D.7) as the approximate variance-covariance matrix, $\hat{\Upsilon}$.

When calculating, $D = \partial m(\cdot)/\partial \theta' \big|_{\theta=\theta_0}$, most calculations are done by taking the straightforward numerical derivative using a two-sided approach with a 1 percent variation in the underlying parameter. However, the first two moment conditions, since they are based on asset quantiles, require additional simplification. Recall that equation (D.1) was written as

$$\mathbb{E}\left[A_{ht} - \bar{A}_j(t, \theta_0; \chi_0) \mid t\right] \times \mathbf{1}\left\{\hat{Q}_{v_{j-1}}(t; \theta_0, \chi_0) \leq A_{ht} \leq \hat{Q}_{v_j}(t; \theta_0, \chi_0)\right\} = 0.$$

This equation can be rewritten as

$$\int_{Q_{v_{j-1}}(A_{ht},t)}^{Q_{v_j}(A_{ht},t)} \left\{ \mathbb{E}\left[A_{ht} \mid t\right] - \bar{A}_j(t,\theta_0;\chi_0) \right\} \times f\left(A_{ht} \mid t\right) dA_{ht} = 0.$$

Applying Liebnitz's rule, the first-order condition becomes,

$$\begin{aligned}
D &= -Pr\left[\hat{Q}_{v_{j-1}}\left(t;\theta_0,\chi_0\right) \leq A_{ht} \leq \hat{Q}_{v_j}\left(t;\theta_0,\chi_0\right) \mid t\right] \times \frac{\partial \bar{A}_j(t,\theta_0;\chi_0)}{\partial \theta'} \\
&\quad + \left\{ \mathbb{E}\left[\hat{Q}_{v_j}\left(t;\theta_0,\chi_0\right) \mid t\right] - \bar{A}_j(t,\theta_0;\chi_0) \right\} \times f\left(\hat{Q}_{v_j}\left(t;\theta_0,\chi_0\right) \mid t\right) \times \frac{\partial \hat{Q}_{v_j}\left(t;\theta_0,\chi_0\right)}{\partial \theta'} \\
&\quad - \left\{ \mathbb{E}\left[\hat{Q}_{v_{j-1}}\left(t;\theta_0,\chi_0\right) \mid t\right] - \bar{A}_j(t,\theta_0;\chi_0) \right\} \times f\left(\hat{Q}_{v_{j-1}}\left(t;\theta_0,\chi_0\right) \mid t\right) \times \frac{\partial \hat{Q}_{v_{j-1}}\left(t;\theta_0,\chi_0\right)}{\partial \theta'}.
\end{aligned}$$

Similarly, recall equation (D.2):

$$\mathbb{E}\left[h - \bar{h}_j\left(\tau,t;\theta_0,\chi_0\right) \mid t,\tau\right] \times \mathbf{1}\left\{\hat{Q}_{v_{j-1}}\left(t;\theta_0,\chi_0\right) \leq A_{ht} \leq \hat{Q}_{v_j}\left(t;\theta_0,\chi_0\right)\right\} = 0$$

It can be rewritten as

$$\int_{Q_{v_{j-1}}(A_{ht},t)}^{Q_{v_j}(A_{ht},t)} \left\{ \mathbb{E}\left[h \mid A_{ht},t,\tau\right] - \bar{h}_j\left(\tau,t;\theta_0,\chi_0\right) \right\} \times f\left(A_{ht} \mid t\right) dA_{ht} = 0,$$

where the first order condition becomes,

$$\begin{aligned}
D &= -Pr\left[\hat{Q}_{v_{j-1}}\left(t;\theta_0,\chi_0\right) \leq A_{ht} \leq \hat{Q}_{v_j}\left(t;\theta_0,\chi_0\right) \mid t,\tau\right] \times \frac{\partial \bar{h}_j\left(\tau,t;\theta_0,\chi_0\right)}{\partial \theta'} \\
&\quad + \left\{ \mathbb{E}\left[h \mid \hat{Q}_{v_j}\left(t;\theta_0,\chi_0\right),t,\tau\right] - \bar{h}_j\left(\tau,t;\theta_0,\chi_0\right) \right\} \\
&\qquad \times f\left(\hat{Q}_{v_j}\left(t;\theta_0,\chi_0\right) \mid t,\tau\right) \times \frac{\partial \hat{Q}_{v_j}\left(t;\theta_0,\chi_0\right)}{\partial \theta'} \\
&\quad - \left\{ \mathbb{E}\left[h \mid \hat{Q}_{v_{j-1}}\left(t;\theta_0,\chi_0\right),t,\tau\right] - \bar{h}_j\left(\tau,t;\theta_0,\chi_0\right) \right\} \\
&\qquad \times f\left(\hat{Q}_{v_{j-1}}\left(t;\theta_0,\chi_0\right) \mid t,\tau\right) \times \frac{\partial \hat{Q}_{v_{j-1}}\left(t;\theta_0,\chi_0\right)}{\partial \theta'}.
\end{aligned}$$

E Data and Sample Selection

This appendix provides greater detail on the data used in estimating the model described in §3.

E.1 Data

I use the original cohort of the Health and Retirement Study (HRS), which was born between 1931 and 1941, and has 12,652 respondents and 7,704 households in the main analysis. However when calculating the transition probabilities for health and mortality, as well as medical expenses, I also use the Asset and Health Dynamics among the Oldest Old (AHEAD) cohort from the HRS, which consists of non-institutionalized individuals born before 1923.

I use the RAND HRS cross-wave supplement (version L) as the initial data set. I then import Social Security earnings history from a separate file where I have calculated, conditional on the assumptions specified in §F, each individual's AIME and PIA as of 1992 as well as each individual's defined benefit levels for every possible age of retirement between 1992 and 2010. Using the combined data set, I use the RAND tenure variable to determine the number of jobs, including baseline job, that are observed between 1992 and 2010.

I define an individual to have retiree health insurance if they report having health insurance coverage that persists after retirement or have access to VA or CHAMPUS benefits (retired or active duty U.S. military benefits). An individual who has health insurance but does not meet these criteria is considered to have tied health insurance. If an individual has medicaid, private health insurance, or another type of means-tested health insurance, I treat them as having no health insurance, since these individuals are more likely to resemble to pool of individuals with no health insurance. I create a household health insurance variable by assuming that if one individual is eligible for retiree health insurance then everyone is. If no one in the household has retiree health insurance, but at least one individual has tied health insurance, then the household acts as if it has tied health insurance. Finally, no member of the household has health insurance, then the household is treated as having no health insurance.

Since the HRS is conducted at two year intervals, I use the reported labor force status in the RAND HRS supplement for labor force participation in years that correspond to survey waves, and then use information regarding last job and data on Social Security earnings history to fill in labor force participation between survey waves. To be participating in the labor force, an individual must report being employed full-time, employed part-time, unemployed, or partially retired. Additionally, the individual must work more than 300 hours per year. If an individual continues in the same job, then I assume that the hours in non-survey years are the same as the previous survey year. I use information on when a person ended his or her last job to deduce between-wave labor force participation and job changes. I only use Social Security data, when

an individual has changed jobs and cannot use the surrounding waves' information regarding the employment (i.e. the inter-wave job was very short, or the individual was not surveyed in adjacent waves). When I do use the Social Security data, I assume the individual is participating in the labor force if they have a positive earnings history.

I consider an individual to be working full-time if he or she reports working full-time and if he or she reports working in excess of 1600 hours per year. An individual is considered working part-time if he or she reports being employed part-time or reports being partially retired with between 300 and 1600 annual hours of work. If I am relying on Social Security reports to determine the individual's work status, then I assume that 4 quarters of coverage corresponds to full-time work and between 1 and 3 quarters corresponds to part-time work. Using Social Security's earnings records is an imperfect measure since the burden for reaching 4 quarters is low, but this is rarely used since most people's work histories can be achieved based on respondent's reports of when they stopped working at his or her last job.

As described in the section on health, individuals provide a self-reported health status to the interviewer on a scale of *excellent, very good, good, fair,* and *poor*. I reduce these self-reports to a binary measure *good* ∈ {*excellent, very good, good*} or *bad* ∈ {*fair, poor*}.

I use the RAND HRS measures of household assets. To create my measure of household assets, I sum the value of the household's primary residence, and the net value of other real estate, businesses, vehicles, stocks, mutual funds, other investment trusts, checking accounts, certificates of deposits, savings accounts, government savings bonds, treasury bills, bonds, bond funds, and any other reported savings, and subtract debt from the household primary residence's mortgage, any other debts based on the primary residence, and and remaining non-residence based debt.

E.2   Sample Selection

The original HRS sample has 7,704 households, which includes 5,813 households with at least one male. Of the male households, 1 is eliminated because the birth year of the respondent is unknown [5,812], 968 are not married [4,844] and 260 are eliminated for missing spousal information in the first wave [4,584]. I keep households that (1) are married in wave 1 and not missing spousal information [4,584], (2) are not missing information on their labor force participation in 1992 [4,575], (3) have never applied for Social Security disability benefits [3,300], (4) are without missing pension [2,628] or Social Security information [2,197], (5) have a spousal age difference of less than 10 years [1,943], (6) are not missing information on either household member's baseline earnings [1,899], and, for computational tractability, (7) households with no more than one defined benefit pension [1,729]. Additionally, I drop annual observations if employment or health status of either household member is not reported, and if health insurance status cannot be determined when the household is less than age 65 (Medicare age) [1,728].

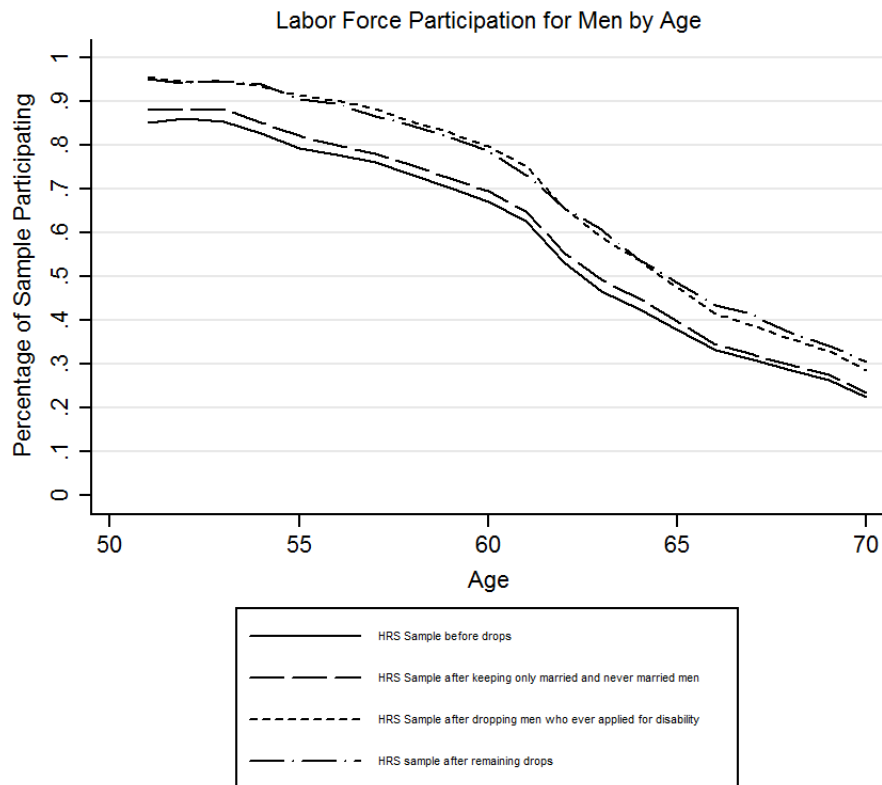**Labor Force Participation for Men by Age**

Figure 11: Sample Selection by Labor Force Participation of Men

After this sample selection, I am left with 1,728 married households. I use the Social Security Administrative data for earnings and labor force participation histories and respondent reports for periods not covered by the Social Security data. Doing so yields an average of 14.95 annual observations per household (out of a maximum possible of 20), providing a long history of observations. Figure 11 shows how my sample selection effects the average rate of male labor force participation. The omission of divorced, separated, and widowed households increases labor force participation slightly, but eliminating those household that ever apply for Social Security disability benefits increases labor force participation at all ages by approximately 10%. This result is not surprising since individuals who credibly apply for disability will likely have a reduced ability to participate in the labor force.

Table 8 provides sample statistics for the entire sample, while table 1 in the main text provide the sample statistics for the estimation sample. Finally, table 9 provides the sample statistics for the model validation sample.

Table 8: Sample Statistics from the selected HRS Sample.

		Men	Women
Age	Mean	57.81	54.89
	Median	57.75	55.08
	Standard Dev.	4.72	4.69
Earnings	Mean	$31,697.86	$12,614.91
	Median	$24,400	$8,320
	Standard Dev.	$35,786	$18,689.36
AIME	Mean	$2,187.74	$703.8
	Median	$2,321.5	$489
	Standard Dev.	$971.85	$692.8
Predicted Annual Pension Benefit	Mean	$14,829.28	$6,696.37
	Median	$7,195.43	$2,611.16
	Standard Dev.	$27,029.28	$12,061.7
% with Current Pension Benefit		25.78	24.34
% Working		78.36	59.43
% Working Full-time		85.16	62.41
% in Bad Health		10.76	9.84
% White		89.93	90.1
Average Years of Education		12.65	12.48

		Household
Assets	Mean	$324,744.2
	Median	$163,272.5
	Standard Dev.	$564,908.1
% with Retiree Health Insurance		64.06
% with Tied Health Insurance		18.34
% with No Health Insurance		17.54
Preference Type (Work, Spousal)	Out	12.04
	Low, Low	23.26
	High, Low	20.54
	Low, High	21.59
	High, High	22.45
Fraction of Women Eligible for Spousal Benefit	Overall	57.18
	1st Asset Quantile	56.77
	2nd Asset Quantile	58.68
	3rd Asset Quantile	56.08
Number of Households		1,728

Notes: Sample consists of only those households with one member between the ages of 51 and 61 in 1992. Individual income is conditional on participating in the labor force in 1992. Predicted Annual Pension Benefit is defined benefit pensions that are vested and is conditional on having a pension. The percentage with current pension is conditional on participating in the labor force in 1992. The percentage working full-time is conditional on participating in the labor force in 1992.

Table 9: Sample Statistics from the portion of the HRS Sample used for Model Validation.

		Men	Women
Age	Mean	54.78	51.75
	Median	54.17	52.08
	Standard Dev.	3.62	3.29
Earnings	Mean	$35,131.12	$13,734.66
	Median	$28,000	$9,800
	Standard Dev.	$37,067.81	$16,785.95
AIME	Mean	$2,355.99	$733.6
	Median	$2,504	$524
	Standard Dev.	$1,015.35	$714.67
Predicted Annual Pension Benefit	Mean	$12,133	$6,631.55
	Median	$5,267	$2,225
	Standard Dev.	$17,493.59	$12,610.25
% with Current Pension Benefit		27.74	23.19
% Working		88.06	66.59
% Working Full-time		88.56	63.65
% in Bad Health		8.98	9.53
% White		90.69	91.24
Average Years of Education		12.81	12.55

		Household
Assets	Mean	$305,626.9
	Median	$141,600
	Standard Dev.	$532,144.8
% with Retiree Health Insurance		59.36
% with Tied Health Insurance		22.45
% with No Health Insurance		18.18
Preference Type (Work, Spousal)	Out	5.48
	Low, Low	24.42
	High, Low	24.21
	Low, High	22.02
	High, High	23.88
Fraction of Women Eligible for Spousal Benefit	Overall	57.5
	1st Asset Quantile	59.34
	2nd Asset Quantile	57.57
	3rd Asset Quantile	55.59
Number of Households		913

Notes: Sample consists of only those households with one member born between 1937 and 1941. Individual income is conditional on participating in the labor force in 1992. Predicted Annual Pension Benefit is defined benefit pensions that are vested and is conditional on having a pension. The percentage with current pension is conditional on participating in the labor force in 1992. The percentage working full-time is conditional on participating in the labor force in 1992.

F Pensions

F.1 Defined Benefit Plans

DB plans provide a guaranteed payment to an employee who is vested. An employee typically becomes vested after 5 or 10 years of service, at which point they will be eligible for a pension benefit based on years of service. Many pension plans define the workers annual benefit ($db_{i,t}$) as:

$$db_{i,t} = (\text{Years of Service}) \times (\text{PoFS}) \times (\text{AFS})$$

where PoFS is the percent of final salary, usually between 1.5% and 2.5%, and AFS is the average final salary, usually the best three or five years of service. PoFS may follow a bend point system based on years of service (e.g. 2.2% for the first 20 years of service and 2% thereafter). Note that to accommodate more gradual retirement, most plans take the best average annual salaries over a worker's lifetime. Depending on the plan, these best years may be required to be consecutive. Most plans offer an early retirement option, usually at ages 50, 55, 58, 60, or 62 assuming the employee is vested. Individuals taking early retirement may have their annual benefit reduced, but this reduction can vary widely by plan. For example, the California State Teacher Retirement System reduces monthly benefits by 50% of the PoFS for each month before age 60 and then keeps it at this rate for the same number of months after age 60.[35] Alternatively, Michigan's teacher pension system permanently reduces monthly benefits by an annualized amount of 6% for each year before age 60.

Once an individual reaches the full-retirement age, usually age 60, 62 or 65, some plans may offer delayed retirement benefits, such as a higher PoFS, but many offer no benefit beyond increased years of service increments. Other employers may offer a longevity bonus to a monthly benefit (e.g. an extra $300 for employees with at least 3 years of service). Re-employment at the same place of employment after claiming a pension plan is discouraged by most plans through benefit reduction or elimination. Alternatively, some employers in an effort to retain older workers have implemented deferred retirement option plans (often called DROPs) which permit a worker to claim his or her benefit, but this benefit is placed in an interest bearing account payable upon retirement.

Those employees who are not vested can receive a refund with interest on the amount that they personally contributed to the pension plan. Some non-vested plans allow for some payment of the employer portion if the employee has greater than 5 years of service.

In the model presented in §3, if an individual has too few years of service to qualify for an

---

[35]For example, consider Jane who is eligible for a $2000 monthly benefit if she retires at her full retirement age in June. If Jane claims in April and she will receive a $1000 monthly benefit from April to August of that year and then she will receive $2000 per month thereafter.

annual benefit, then the vested benefit level is treated as a lump sum benefit when the individual leaves the baseline job. Additionally, if the individual's plan is like the California plan above (lower benefits early, higher benefits later) or has a Social Security topper (higher benefits before age 62 and lower benefits after) then the long term rate (i.e. what the benefit level is 10 years after claiming) is treated as the monthly benefit level and the individual has to pay a lump sum payment at claiming that makes up for the difference. This is done for computational feasibility.

Some individuals have access to multiple defined benefit pension plans. I assume that the individual cannot claim until the early eligibility date of his or her largest DB plan. If smaller plans are not yet eligible, then I still assume that the individual receives the same annual benefit that they would receive in the long-term, but he or she makes lump sum payment upon claiming to cover those additional benefits.

Employees eligible for a benefit can generally elect to have a survivor benefit that is 0-100% of their benefit amount, where benefits are reduced according to the actuarially fair rate of adjustment (i.e. the pension provider will consider the possible survivor's gender and relative age). Most plans include a survivor option should the employee die prior to retirement that pays a fixed benefit at death (similar to life insurance) and may pay a monthly benefit that makes an assumption about what the employee would have done if he or she had survived and chosen a plan. For example, in the California State Teacher Retirement System, the survivor receives a benefit based on a 50% beneficiary option, so that the survivor would be eligible to receive 50% of the employee's benefit, which would be reduced based on the survivor's gender and relative claim age.

DB plans work much like Social Security, often providing disability insurance to the employee and life insurance benefits to spouses and children. Due to data limitations, and similar to what I do for Social Security, I ignore these benefits here as most couples in the HRS do not have children living with them. For survivor benefits, I will assume that individuals have claimed a 0% beneficiary option to simplify the analysis. Otherwise, benefits from DB plans will be defined as in the summary plan descriptions provided to HRS.[36]

F.2 Defined Contribution Plans

DC plans do not provide a guaranteed payment to an employee upon retirement. Many plans require an employee to be vested, usually 3 to 5 years, before he or she is eligible to retain employer contributions. Employers generally match employee contributions, up to some maximum level, such as a $1 match for every $1 contributed or $1 match for every $2 contributed. Most of these plans are administered by private entities, and provide the employee with a wide range of invest-

---

[36]This assumption can be quite strong since it is possible that the different household types, described in §5.1.4, could vary in their choice of benefit plan. However, this reflects an income effect, and does not induce further notches in the budget constraint, so I view this as a reasonable assumption to make in light of the computational difficulties that this would otherwise entail.

ment possibilities. These plans generally do not act as a form of insurance for the employee, so employees have to separately subscribe to disability or life insurance plans. Any surviving beneficiaries receive access to the DC plan's account balance.

Taxation of defined contribution plans is based on the taxable amount, which is generally considered the amount that has not previously been taxed. In most cases, this is comprised of the deferred wage, taxed based on your income tax bracket in the year the individual receives the annuity and any gains in those contributions over the lifetime. There are only two major notches in a household budget constraint based on DC plans: the 10% tax penalty for withdrawals before age $59\frac{1}{2}$ and the required withdrawals from investment retirement account in the year after an individual turns $70\frac{1}{2}$. I do not expect these constraints to be binding for the majority of the HRS sample. One of the major reasons an agent would want to withdraw from a DC plan before $59\frac{1}{2}$ would be due to a medical expense shock, which would be exempted under Internal Revenue Service rules.

I will treat defined contribution plans as additional post-income tax assets, therefore these plans will be subject only to a personal income tax on any growth and I will assume that they grow at the rate of return, r. This is a strong assumption because standard 401(k) contributions by the worker and all contributions made by an employer are generally are not taxed as income until disbursement. I omit this detail for computational feasibility because a more accurate model of these assets would require both knowledge on HRS's part of which assets are and are not pre-tax (which HRS does not know) and an additional state variable in the estimation of the model that tracks pre-tax assets.

F.2.1 Defined Contribution Imputations

The Health and Retirement Study (HRS) has released imputations for DC plan wealth through the sixth wave (2002). To fill in DC plan wealth for 2002 to 2010, I impute the DC plan wealth following a procedure similar to how RAND imputes income and wealth levels, and compare my imputations with the earlier imputations done by the HRS staff for the overlapping years (2000,2002 or waves 5 and 6).[37]

The HRS collects information for up to 4 pensions each interview wave. If the respondent reports having either a DC or a combination plan and is missing information on the plan's balance, I impute a balance amount. Individual's who did not know their balance amount were asked a series of unfolding brackets to help approximate the balance (i.e. if you do not know your pension, is it greater or less than $20,000). Unlike RAND's imputation procedure, I do not impute ownership. Conditional on reported ownership of a DC plan, I impute the bracket if none is given, and then conditional on bracket, I impute an account balance.

---

[37]See *Imputations for Pension-Related Variables, Final, Version 1.0 (June 2005)* by the Health and Retirement Study for a description of the HRS's imputation process for waves 1-6.

I impute brackets for individuals who report a DC or combination plan, but do not provide a complete range. I begin by estimating an ordered logit model of the DC balance bracket on the sample of individuals who report complete brackets but do not report a balance. The covariates include dummies for if there is greater than 50% chance of leaving a bequest of more than $10,000, greater than 50% chance of leaving a bequest of more than $100,000, high school diploma or higher, college degree or higher, whether the respondent self reports excellent or very good health, whether the respondent self reports poor or fair health, whether the respondent works in a professional occupation, self-employed, married, spouse-age missing, and non-white, as well as continuous measures of tenure, own-age, own-age squared, spouse-age, and spouse-age squared. All covariates with the exception of the bequest arguments are interacted with the individual's gender. Second, I use the fitted model to predict the probability of being in each of the five brackets, and then use these probabilities to generate a cumulative distribution. Third, I draw a random number from a uniform distribution, and compare the random number to the cumulative distribution in order to assign each of the individuals with missing bracket information to a bracket.

Finally, I impute account balances for all individuals who report a DC or combination plan, but do not provide a specific balance amount. I begin by estimating a standard regression on log account balances using the same covariates as used in the bracket imputation in addition to the dummies for each individual's respective balance level. Second, I use the fitted model to predict DC account balances for all individuals who report a DC or combination plan. Using a modified hot-deck approach, I sort the data by the imputed account balances and then assign account balances to missing observations based on a weighted average of the nearest-neighbors.

This imputation procedure produced similar results to the HRS imputation procedure used for the first six waves. Waves 5 and 6 were estimated for both samples. In the 6th wave, the imputation procedure produced 558 additional balances, bringing the total observed to 1,569. The mean [standard deviation] of log account balances before imputation was 10.19 [1.84] and after this imputation procedure it was 9.87 [2.00]. HRS's imputation procedure produced a mean [standard deviation] of 9.79 [1.99]. In the 5th wave, the imputation procedure produced 524 additional balances, bringing the total observed to 1,882. The mean [standard deviation] of log account balances before imputation was 10.05 [1.90] and after this imputation procedure it was 9.84 [1.91]. HRS's imputation procedure produced a mean [standard deviation] of 10.07 [1.89]. Since both imputation methods produce similar results, I use the HRS's imputations for the first 4 waves (1992-1998), and the aforementioned imputation method for the remaining waves (2000-2010).

F.2.2 Tax Treatment of IRAs and tax-deferred accounts

At baseline, 1992, I observe the respondent's report of how much money the household has saved in defined contribution plans, such as 401(k), 403(b), and IRA accounts. Standard accounts like

these are usually composed of pre-tax earnings, meaning the individual has not paid income tax on this money. Therefore, when the money is disbursed, it would be taxed as unearned income (i.e. it will not be subject to payroll taxes, but it will be subject to income tax). The model, as currently estimated, does not include a distinction between pre- and post-tax savings because of the computational burden associated with estimating these separately. Since defined contribution plans are treated as post-tax savings in the model, I must make an assumption about how much of the account at baseline should be reduced to reflect the future payout of income taxes.

I assume that the money is disbursed based on the expected joint life-expectancy from the 1994 IRS joint life expectancy table.[38] I account for the couple's estimated Social Security benefits and defined benefit if both claimed their Social Security at age 62 and DB pension benefits at age of first eligibility, and assume the DC disbursements start at age 70. Put simply, an individual will draw on his or her DC account starting at age 70, and will take the minimal disbursements required by the IRS. Therefore, the respective amount that they will be taxed will be based on their annual income comprised of Social Security benefits, DB benefits, and DC disbursements. I account for the limited amount of Social Security income that is taxable. I set the tax attributable to the DC disbursement assuming that DC disbursements are the last dollars taxed. I then sum tax payments from the DC plan across years and subtract this from the total DC account at baseline. The account is then assumed to be comprised of post-tax dollars.[39]

---

[38]The 1994 IRS publication 590 was the earliest I could locate. Age 62 corresponds to the mean age people plan to begin collecting benefits.

[39]Note that this procedure is ad hoc: While I account for age differences within the couple, I do not account for the individual decision of when it is disbursed and whether the couple continues to work. Since the only source of income for these individuals is via annuity payments from Social Security, DB plans and minimal DC plan disbursments, and not based on earnings from work, the taxed amount should be lower than expected.

G   Earnings Profiles

The model described in §3 assumes that each individual can choose between employment in her baseline full-time job (FT-B), a non-baseline full-time job (FT-NB), a part-time non-baseline job (PT), and no job. Earnings in these employment states are assumed to be non-stochastic and known to the individual, similar to Blau and Gilleskie (2006). However, unlike Blau and Gilleskie (2006), I allow earnings to change with age in all possible employment states. This is done to reflect diminished employment prospects with age and the fewer hours worked after age 58 among those participating in the labor force. In this section, I will first specify how baseline full-time earnings are determined, and then consider how non-baseline full-time and part-time earnings are determined.

I define the baseline job as the full-time job an individual currently holds at baseline (1992). If an individual leaves there FT-B job for any other state, then he or she cannot return to the FT-B job. Annual earnings for FT-B jobs are determined from individual self-reports in the Health and Retirement Study, and grow at a constant rate, consistent with the HRS pension calculator. The HRS pension calculator uses information collected from employer reported "summary plan descriptions" in combination with the worker's reported annual earnings and user-specified assumptions regarding nominal wage growth, inflation, and real interest rates to predict the worker's annual benefit levels by respective quit dates. Consequently, the earnings model must reflect the same assumptions used in the pension calculator to ensure that the correct benefit levels are predicted. The assumptions used in the pension calculator are a real interest rate of 4%, inflation of 2%, and nominal wage growth of 0%. This is consistent with the realized negative real wage growth rate of approximately 2%, following baseline, among individuals with pension plans in the sample specified in §4.

The situation is more complicated for non-baseline earnings. Approximately 56.6% of men and 32.0% of women are in a FT-B job at baseline. From the men (women) who have a FT-B job at baseline, 18.1% (24.9%) will transition to a PT job from the FT-B job, and 16.4% (15.5%) will transition to a FT-NB job from the FT-B job. Of the men (women) transitioning from FT-B to a FT-NB job, 31.9% (33.3%) will receive earnings increases after the move. Median annual earnings for men at FT-NB jobs rise until about age 57 and then decline, as seen in figure 12. This is despite median annual hours falling prior to age 57 and then remaining relatively constant for FT-NB jobs (as in figure 13). Alternatively, the story for women in FT-NB jobs is that annual earnings decline after 54 and then becomes noisy for ages 60+, despite annual hours remaining largely unchanged. Finally, part-time earnings decline as hours decline for both sexes.

Non-baseline jobs represent an alternative employment option for individuals at baseline and each subsequent period (up until the maximum working age of 70). Therefore, it is important

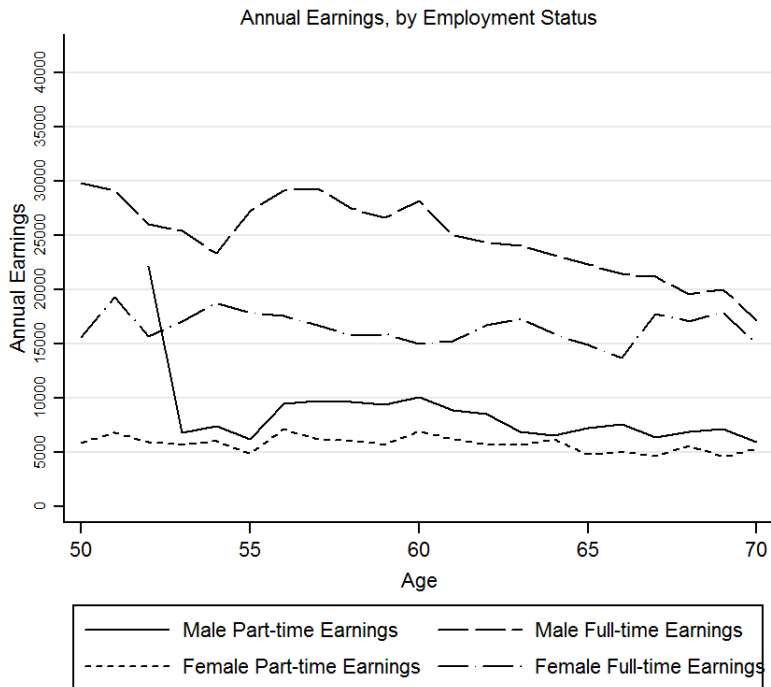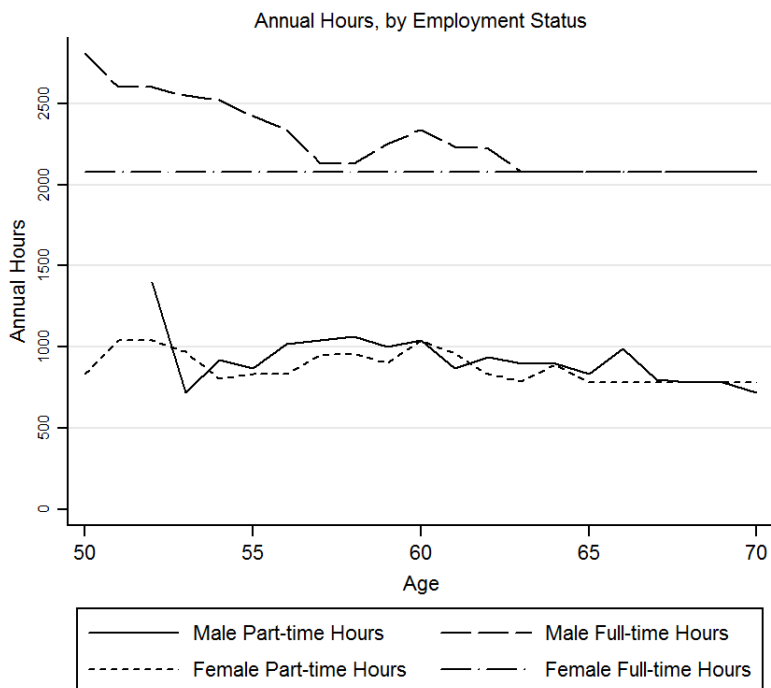## Figure 12: Median Earnings - Non-baseline jobs

Annual Earnings, by Employment Status

Legend:
- Male Part-time Earnings
- Male Full-time Earnings
- Female Part-time Earnings
- Female Full-time Earnings

Figure 13: Median Hours - Non-baseline jobs

Annual Hours, by Employment Status

Legend:
- Male Part-time Hours
- Male Full-time Hours
- Female Part-time Hours
- Female Full-time Hours

to assign a feasible wage that a worker might believe is available to her outside of her baseline job (if she is working), or if she was to return to the workforce (if she was not working at baseline).

I estimate individual log earnings profiles (separately by sex and employment status), $\ln w_{it}$, for jobs that begin after baseline - the first sampling wave of the HRS in 1992. Baseline in my full sample corresponds to an average age of 57.8 for men and 54.9 for women. These jobs represent alternatives to the individual's baseline job, which most individuals have held for a long time. The independent variables, x_{it}, include a quartic in age and a quadratic in tenure (tenure is only included for FT-NB jobs). At this late age, I model wages as being primarily determined by an individual i's time invariant ability, c_i^j, if he or she is working in job $j \in \{\text{FT-NB, PT}\}$:

$$\ln w_{it}^j = x_{it}^{'} \beta^j + c_i^j + \varepsilon_{it}^j , \tag{G.1}$$

where ε_{it}^j is a the model error term such that $\mathbb{E}\left[\varepsilon_{it}^j | j, x_{i1}...x_{iT_i}, c_i^j\right] = 0$, and T_i corresponds the last observed period for individual i . The model can then be used to predict the time invariant fixed effect, $\hat{c}_i^j = \overline{\ln w_i^j} - \bar{x}_i \hat{\beta}$ where $\overline{\ln w_i} = \sum_t (\ln w_{i,t}/T_i)$

When (G.1) is estimated, a value of $\hat{c}_i^j$ can be calculated for all individuals with at least two periods where non-baseline jobs are observed. The β^j terms in equation (G.1) are identified by variation within individuals over time.

Some individuals will not have a predicted fixed effect, $\hat{c}_i^j$. Specifically, individuals who (i) never work in another job after quitting his or her baseline job, and (ii) individuals who never work. In order to predict a fixed effect for these individuals, I regress

$$\hat{c}_i^j = \theta_0^j + \theta_1^j educ_i + \theta_2^j AIME1992_i + \theta_3^j EarningsBaseline_i + \eta_i^j \tag{G.2}$$

on the same individuals used in estimating equation (G.1), and then use (G.2) to predict $\hat{c}_i^j$ for those missing individuals due to (i). I do the same thing for individuals who never work, but exclude baseline earnings.

Predicted earnings profiles for individual i at each age t in job j are made by substituting the respective values into equation (G.1). Predicted profiles for the mean worker are included in figure 14.

I do not estimate a combined model in (G.1), because this model specifically prevents the change in the quality of the match, $\Delta \varepsilon_{it}$, from being correlated with change in employment status, which rules out most types of endogenous job search. This is particularly problematic in my setting, where I observe workers occasionally getting higher wages on part-time jobs relative to full-time jobs. In fact, since 23% of individuals who have both a PT and FT-NB jobs after baseline have higher part-time wages, it is very likely that changes in observed employment status may be driven by positive shocks during job search.

Figure 14: Annual Earnings by Employment Status

Annual Earnings, by Employment Status

Legend:
- Male Part-time Earnings
- Male Full-time Earnings
- Female Part-time Earnings
- Female Full-time Earnings
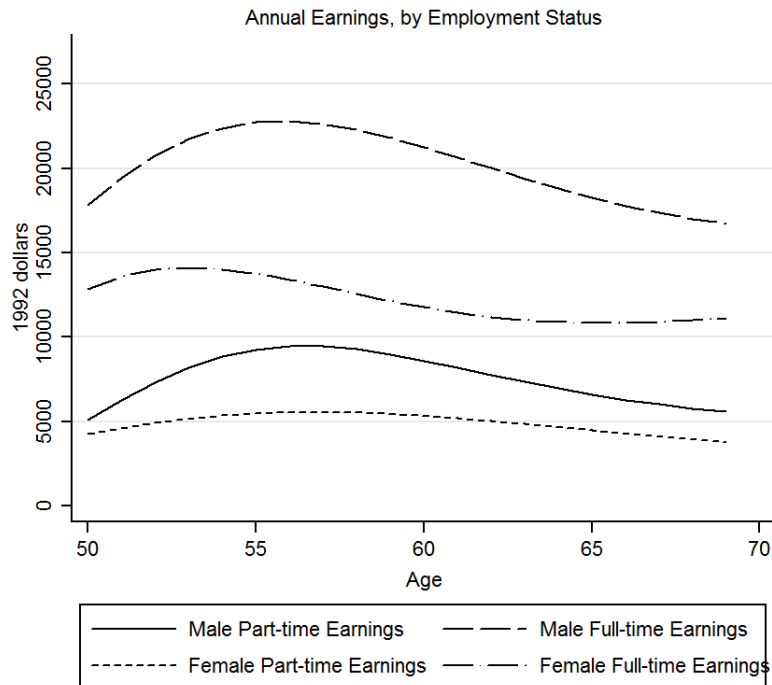
H Health and Mortality Transitions

Health and mortality transitions are estimated using logit model based on a cubic in age, and lagged health status.

Figure 15 shows the 1 year transition probabilities from good to bad health and bad to bad health, for men and women. Men are more likely to move into and stay in bad health (relative to women) as they age. The probability of the average man (woman) remaining in bad health steadily increases from around 72% (72%) at age 50 to 98% (96%) at age 100. Likewise, the probability of the average man (woman) transitioning from good health to bad health increases from 7% (7%) at age 50, to 50% (40%) at age 100.

Figure 16a shows the probability of death for men conditional on health status. As a point of reference, I include information from the Social Security actuarial tables for the 1933 birth cohort. The figure indicates that at younger ages, my model under-predicts the conditional population mortality rate, which is to be expected since the sample used is going to be more likely to have worked and includes younger cohorts. At older ages, the model over-predicts the mortality rate, which is also expected since members of the AHEAD cohort, comprised of birth cohorts before 1924, are used identify mortality rates at these ages. Additionally, figure 16b shows the comparable result for women.

I Medical Expense Distribution

Each period, the household faces a medical expense shock based on its health status. As discussed in §5.1.3, I use a transitory shock from a distribution that is based on the the original HRS sample.

The HRS collects data on self-reported out of pocket medical expenditure $(M_{i,t})$, which is imputed by the Labor and Population Program at the RAND Institute on Aging. In estimating the medical expense distribution, I include members from the Asset and Health Dynamics among the Oldest Old (AHEAD) cohort from the HRS. This sample consists of individuals born in 1923 or before. The combined sample is used to identify the distribution of medical expenses into old age.

I estimate the distribution of medical expense separately for ages above and below age 65, by regressing the logarithm of out-of-pocket medical expenses on a quadratic in age conditional on health insurance, labor force participation, and health status, which represent states of the structural model. Age 65 is chosen as a break-point since most individuals qualify for Medicare at this age and it becomes the primary insurer of the population above 65. As a result, the expense distribution can be expected to differ across groups on either side of age 65.

Previous work has used estimates of total medical expenses, and has generally used another data source for total medical expenditure because it is not observed by the HRS. I compare the distribution of $M_{i,t}$ to total medical expenditure found by Blau and Gilleskie (2006), who use

Figure 15: Probability of being in bad health
Conditional on previous health status
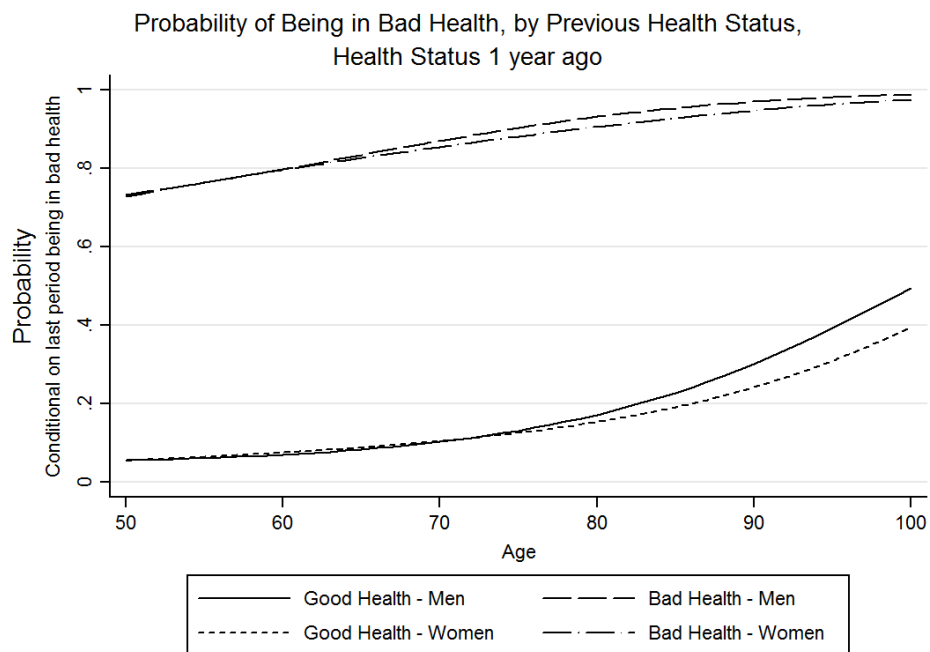
Figure 16: Probability of Death by Sex
Conditional on previous health status
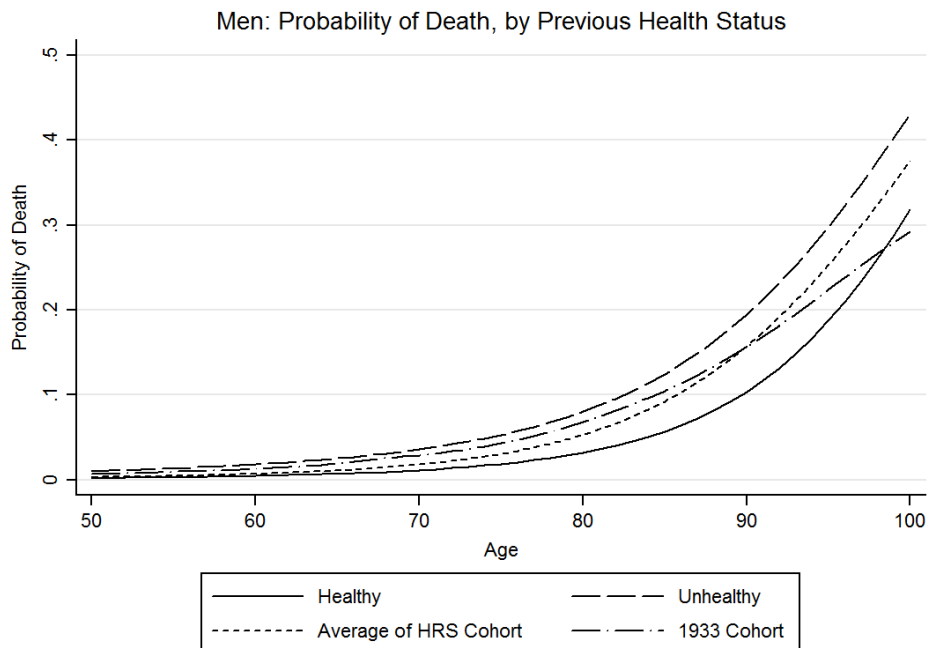
(a) Men
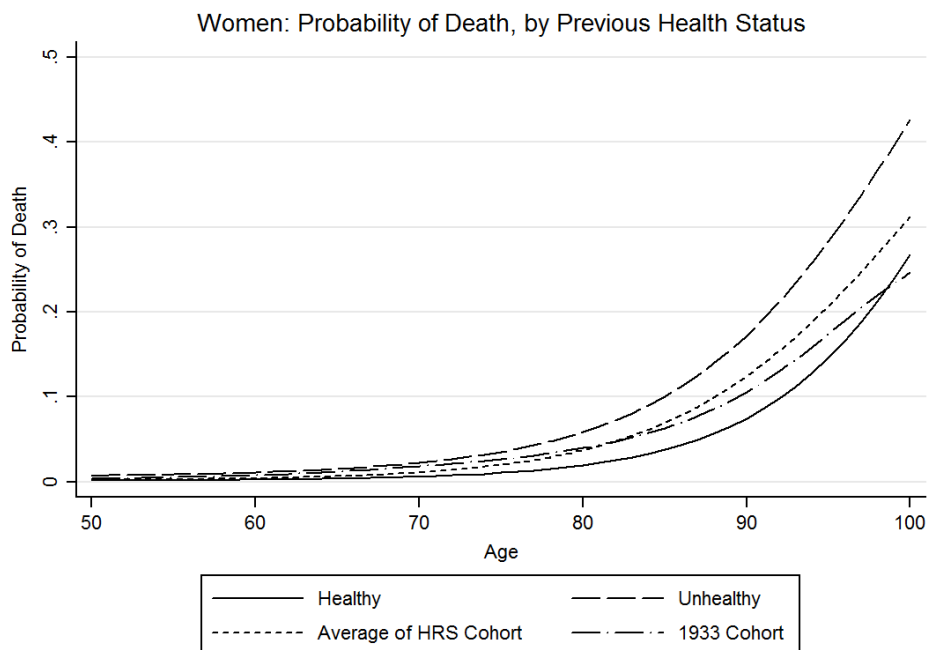
(b) Women

an external survey - the 1987 National Medical Expenditure Survey. I observe that my medical expenditure estimates are generally lower at every level, particularly they are much lower at higher levels of medical expenditures. This is to be expected since they were attempting to estimate total medical expense, and health insurance limits catastrophic medical expenses.

Past literature that has included medical expense uncertainty has usually been focused on how health insurance alters retirement behavior. Due to computational limitations, I am unable to include a persistent process for medical expenses. Persistence in medical expenditures does exist indirectly through persistence in health status. I expect that this will lead to underestimating the household's lifetime medical expense risk.

J Preference Types

As described in §5.1.4, households can vary based on characteristics that will be reflected in their preference for consumption versus leisure, but are not otherwise captured by the typical state variables. For this reason I include a preference index, as in Keane and Wolpin (1997), van der Klaauw and Wolpin (2008), and French and Jones (2011), to capture heterogeneity in preferences for consumption, own-leisure, spousal leisure, time, and household decision-making.

I construct my preference index by regressing each individual's labor force participation on a quartic in age, household health status, assets, earnings, health insurance status, the individual's AIME, defined benefit flow (if eligible), marital status, and a full set of interactions of these terms. Furthermore, I include in this regression three variables pertaining to the individual's preference for work:

1. Even if I didn't need the money, I would probably keep on working. (Agree or disagree)

2. When you think about the time when you and your husband or wife will retire, are you looking forward to it, are you uneasy about it, or what?

3. On a scale of 1 to 10, how much do you enjoy your job?

and, I include four more variables the pertain to the individual's preference for his or her spouse:

1. Generally speaking, would you say that the time you spend together with your husband or wife is extremely enjoyable, very enjoyable, somewhat enjoyable, or not too enjoyable?

2. When it comes to making major family decisions, who has the final say – you or your husband or wife?

3. Some couples like to spend their free time doing things together, while others like to do different things in their free time. What about you and your husband or wife? (together, separate, or sometimes together and sometimes separate)

4. I am going to read you a list of things that some people say are good about retirement. For each one, please tell me if, for you, they are very important, moderately important, somewhat important, or not important at all. Having more time with husband or wife.

For each of the above questions, I create a binary variable for each, either lumping answers such as agree and strongly agree together, or partitioning it by the median answer. I estimate the above regression separately for men and women. For each individual, the work preference index is the sum of the work preference coefficients multiplied by their respective independent variables, and similarly for the spouse preference index. The household's work or spouse preference index is

simply the equally weighted sum for each household member's respective preference indices. The household preference indices are then converted into binary measures by partitioning them at each measures' median.

I observe that the work preference index is positively correlated with marriage, earnings, assets, AIME, defined-benefit pension flows, and negatively correlated with health. The spouse preference index is positively correlated with assets and health, but negatively correlated with earnings and AIME. An "out' preference index is created for households who were not asked the work questions in the first period because they were not working. As noted in table 1, the initial distribution consists of 17.4% of the "out" preference type, and then a relatively even distribution between the four other preference types.

K   Additional Figures depicting Moment Matching

Figure 17: Asset Quantiles (by thirds) by Male Age
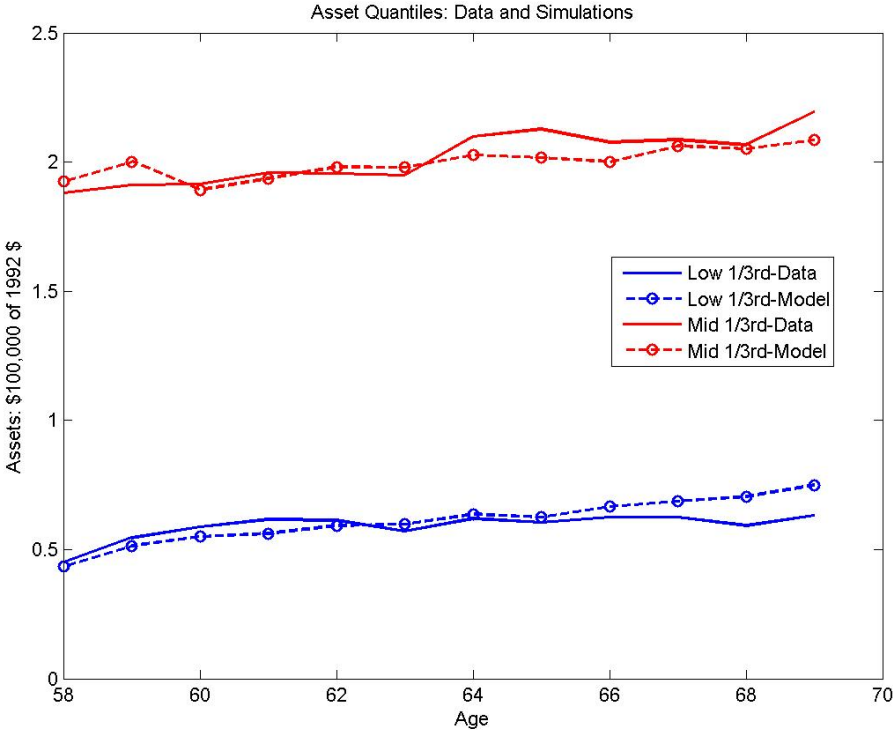
Figure 18: Asset Quantile Shares by Preference Type

(a) Type 0 - the Out Type
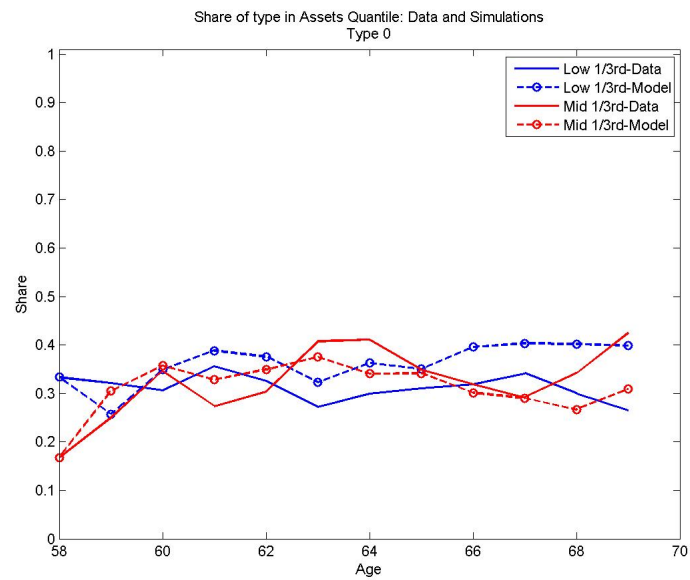
Share of type in Assets Quantile: Data and Simulations
Type 0

(b) Type 2 - Low Preference for Own Leisure, Low Preference for Spousal Leisure

Share of type in Assets Quantile: Data and Simulations
Type 2

Share of type in Assets Quantile: Data and Simulations
Type 3

(c) Type 3 - High Preference for Own Leisure, High Preference for Spousal Leisure

Share of type in Assets Quantile: Data and Simulations
Type 4
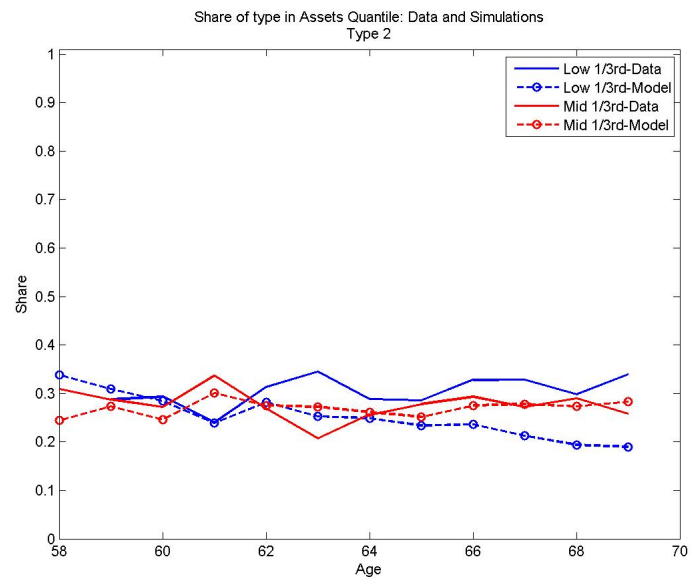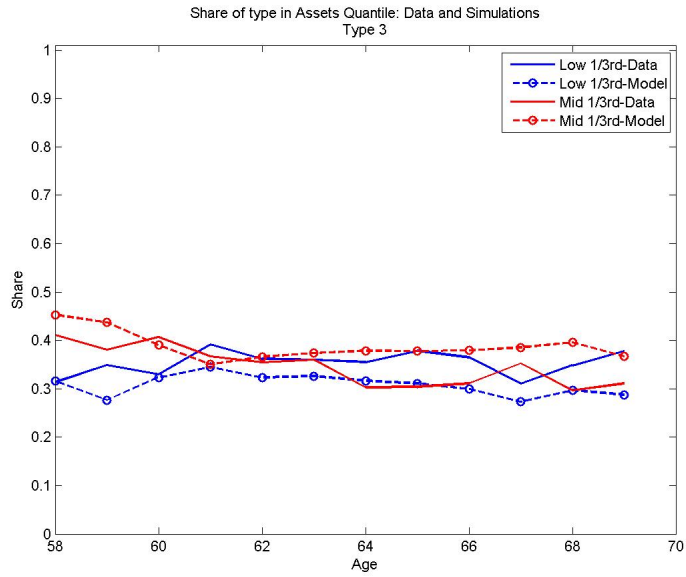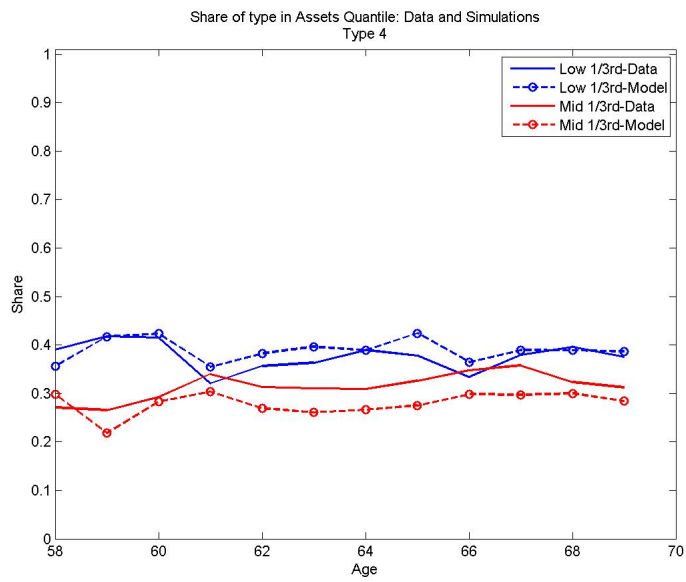
(d) Type 4 - Low Preference for Own Leisure, High Preference for Spousal Leisure

Figure 19: Men Labor Force Participation by Preference Type

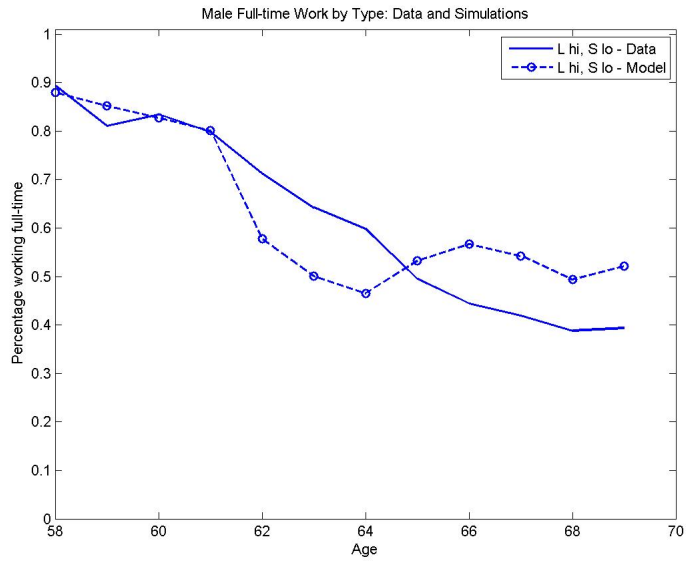(a) Type 2 - Low Preference for Own Leisure, Low Preference for Spousal
Leisure

Male Participation Rates by Type: Data and Simulations

(b) Type 3 - High Preference for Own Leisure, High Preference for Spousal
Leisure

Male Participation Rates by Type: Data and Simulations

Male Participation Rates by Type: Data and Simulations
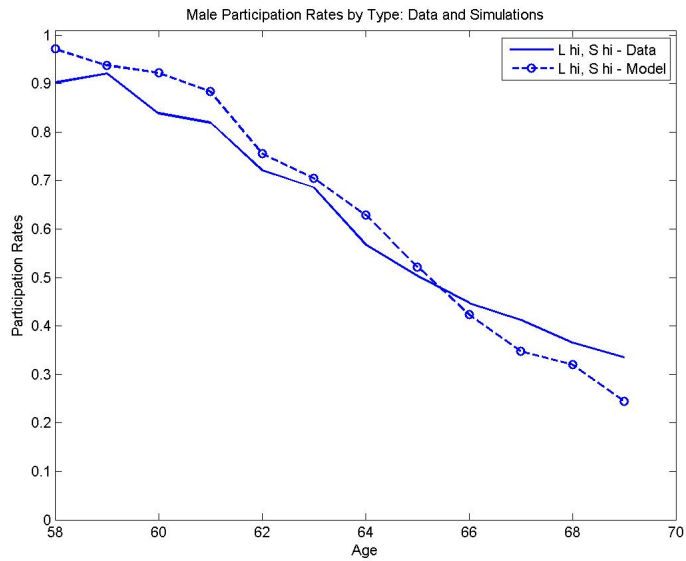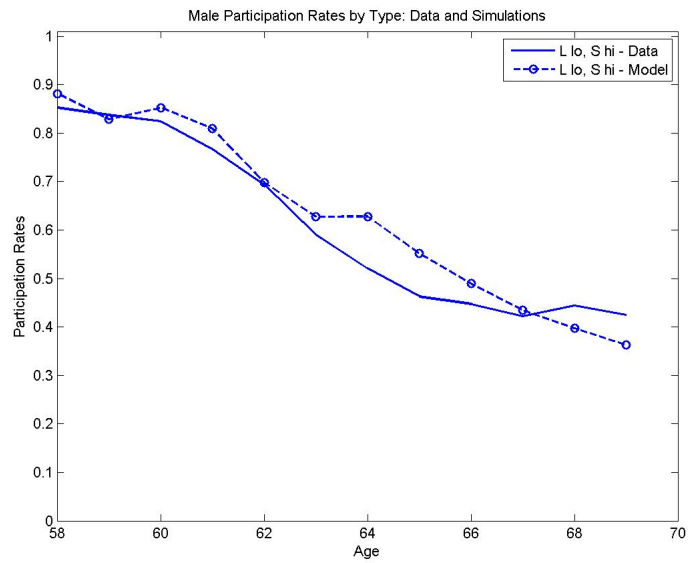
(c) Type 4 - Low Preference for Own Leisure, High Preference for Spousal Leisure

Figure 20: Women Labor Force Participation by Preference Type

(a) Type 2 - Low Preference for Own Leisure, Low Preference for Spousal Leisure

Female Participation Rates by Type: Data and Simulations
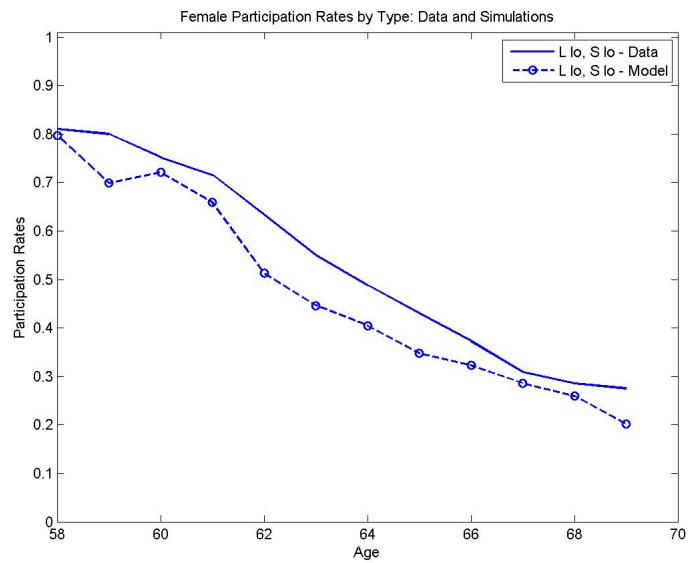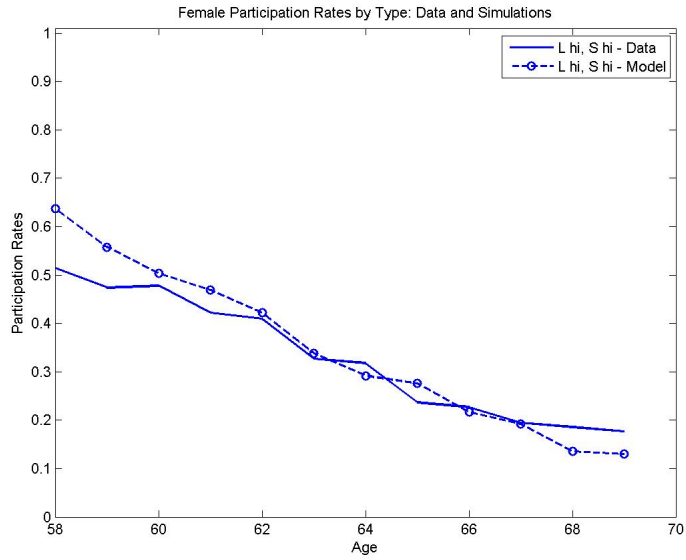
(b) Type 3 - High Preference for Own Leisure, High Preference for Spousal Leisure
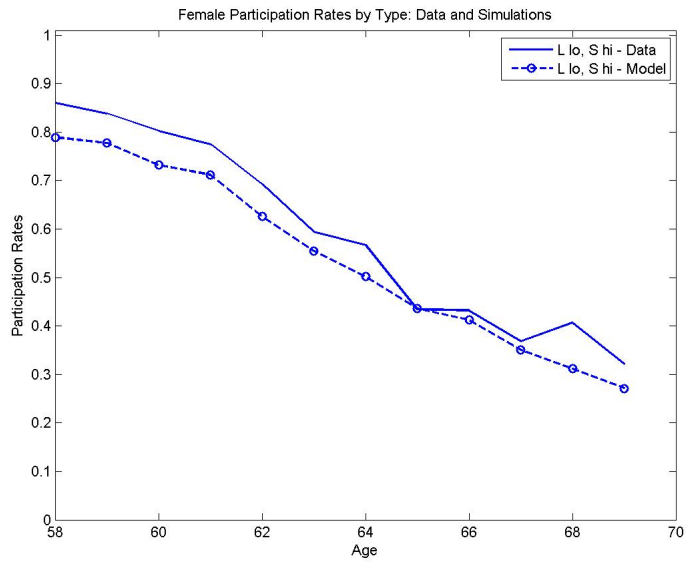
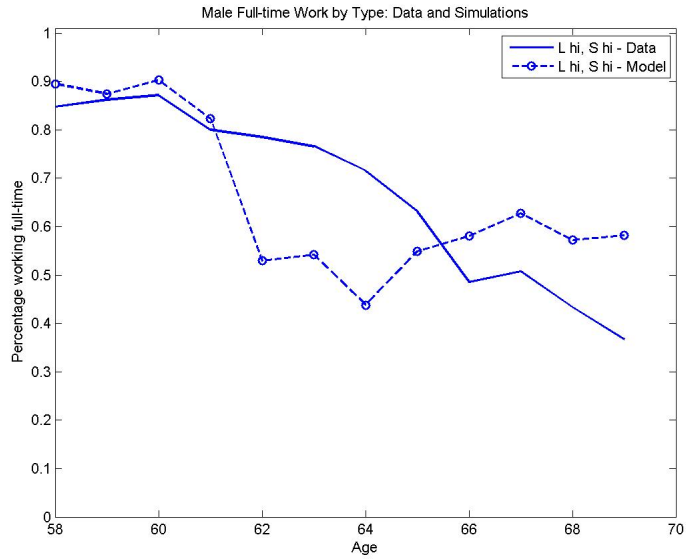
(c) Type 4 - Low Preference for Own Leisure, High Preference for Spousal Leisure

Figure 21: Men Full-time work by Preference Type

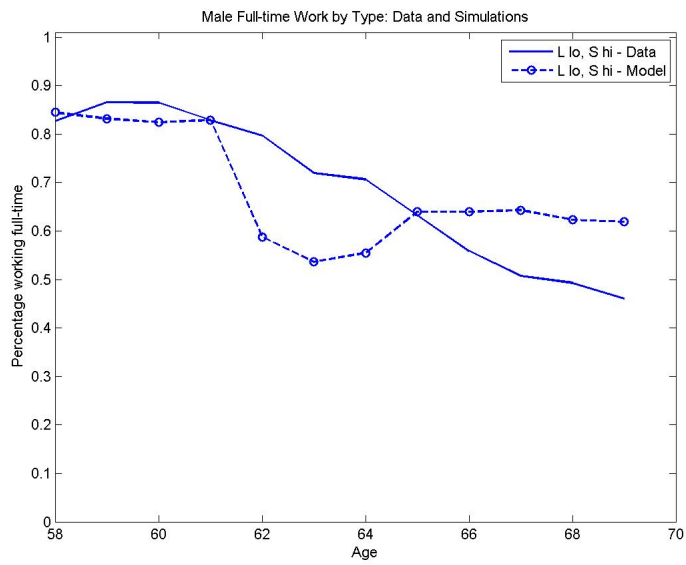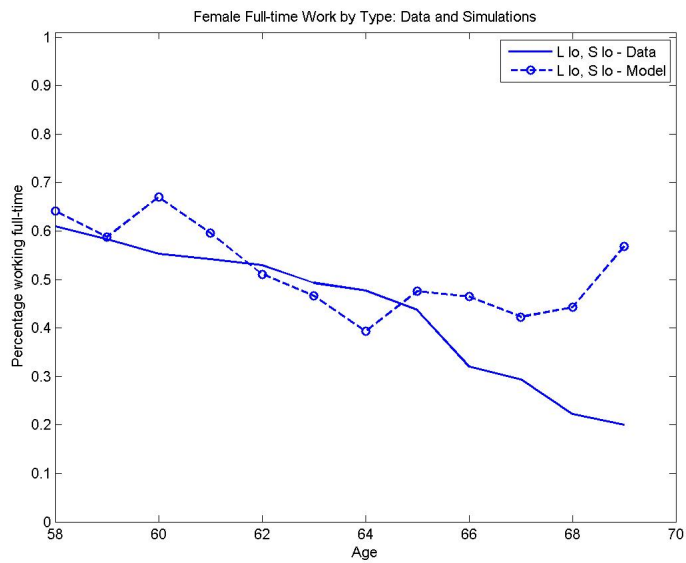(a) Type 1 - High Preference for Own Leisure, Low Preference for Spousal Leisure

Male Full-time Work by Type: Data and Simulations

(b) Type 2 - Low Preference for Own Leisure, Low Preference for Spousal Leisure

Male Full-time Work by Type: Data and Simulations
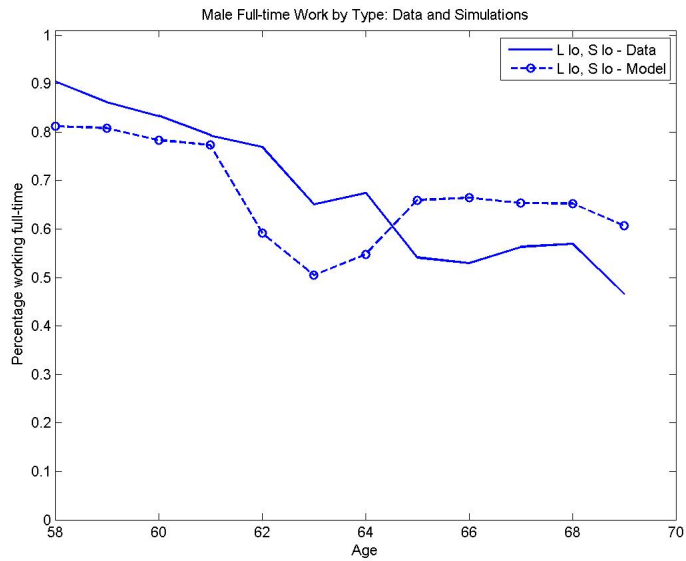
(c) Type 3 - High Preference for Own Leisure, High Preference for Spousal Leisure

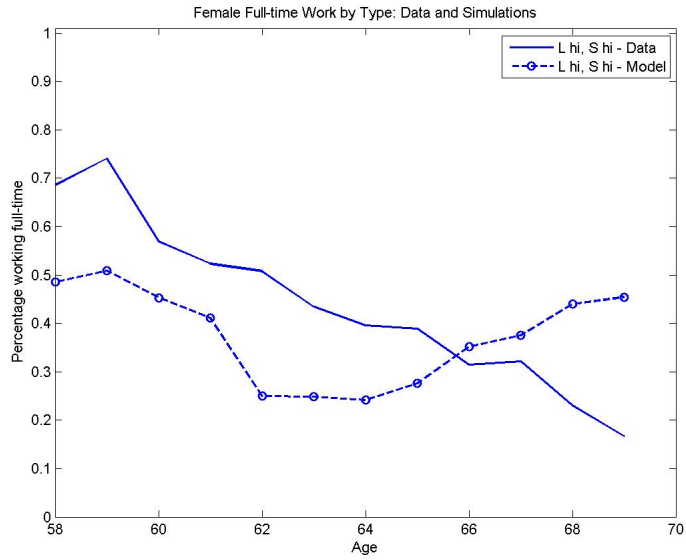(d) Type 4 - Low Preference for Own Leisure, High Preference for Spousal Leisure

Figure 22: Women Full-time work by Preference Type

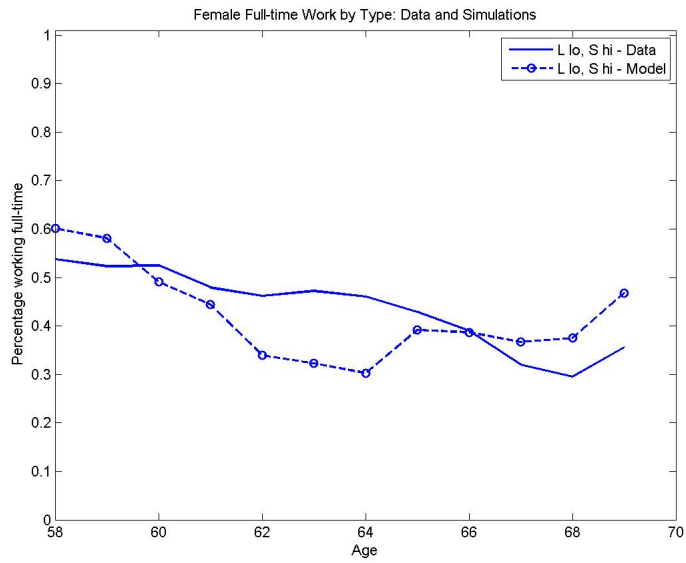(a) Type 1 - Low Preference for Own Leisure, Low Preference for Spousal Leisure

Female Full-time Work by Type: Data and Simulations

(b) Type 2 - Low Preference for Own Leisure, Low Preference for Spousal Leisure

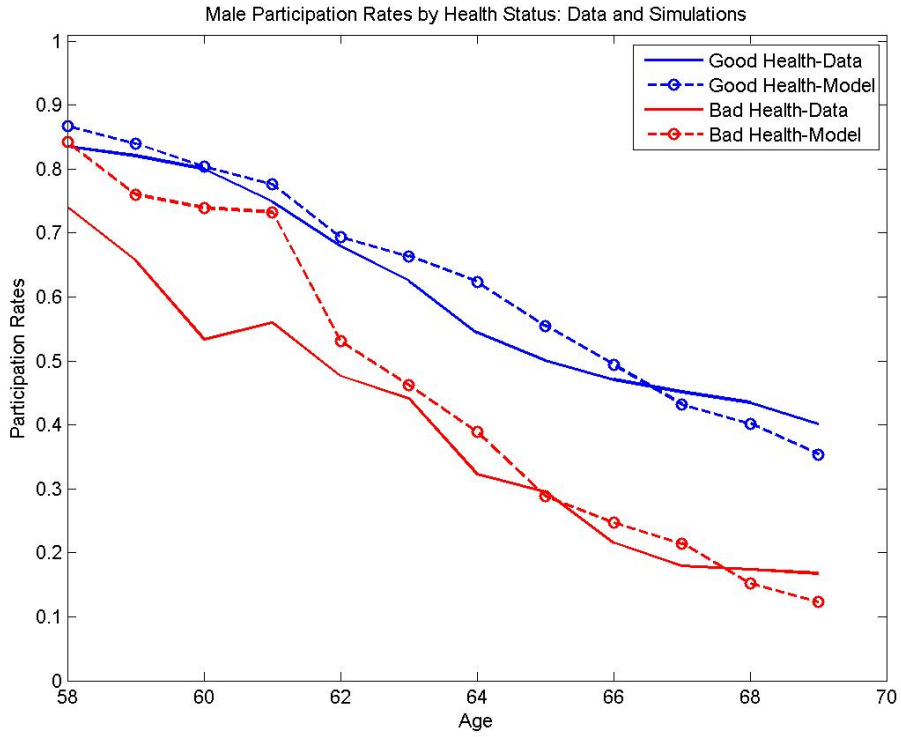Female Full-time Work by Type: Data and Simulations

(c) Type 3 - High Preference for Own Leisure, High Preference for Spousal Leisure

(d) Type 4 - Low Preference for Own Leisure, High Preference for Spousal Leisure

Figure 23: Participation Rate by Health Status

(a) Men

Male Participation Rates by Health Status: Data and Simulations

(b) Women

Female Participation Rates by Health Status: Data and Simulations