Essays on Dynamic Structural Models for Employment and Organization

by

David M. K. Knapp

A dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy (Economics) in the University of Michigan 2014

Doctoral Committee:

Professor Charles C. Brown, Chair
Professor John P. Laitner
Professor Brian P. McCall
Assistant Professor Xun Wu

Dedication

To my Mother and Father

Acknowledgment

I thank Charles Brown, John Laitner, Brian McCall, Brian Wu, Michael Nolte, Janet Keller, Helena Stolyarova, Nancy Herlocher, Molly Saunders-Scott, Reid Dorsey-Palmateer, Yu Zhou and Evan Starr for their thoughtful comments and assistance. Support for this project has been provided by the Boston College Center for Retirement Research (project BC12-D6) pursuant to a grant from the U.S. Social Security Administration, as well as Rackham Graduate School and the Michigan Institute for Teaching and Research in Economics. Any errors, opinions, or results herein are those of the author and are not to be construed as representing the opinions or policy of the Social Security Administration, the Center for Retirement Research, or the University of Michigan.

Table of Contents

List of Figures

List of Tables

List of Appendices

Chapter 1:
The Effect of Social Security Auxiliary Spouse and Survivor's Benefits on the Household Retirement Decision

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 experi-

[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.

ments 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.

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 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, 1986a; 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 (2000a, 2004a) provide a framework for household decision-making that accounts for interdependence of

[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.

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

[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.

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 1.1 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

[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.

Figure 1.1: Growth Rates in Annual Benefit Payment of Defined Benefit Pensions, by Age (Multiply vertical axis by 100 for percent growth rates)

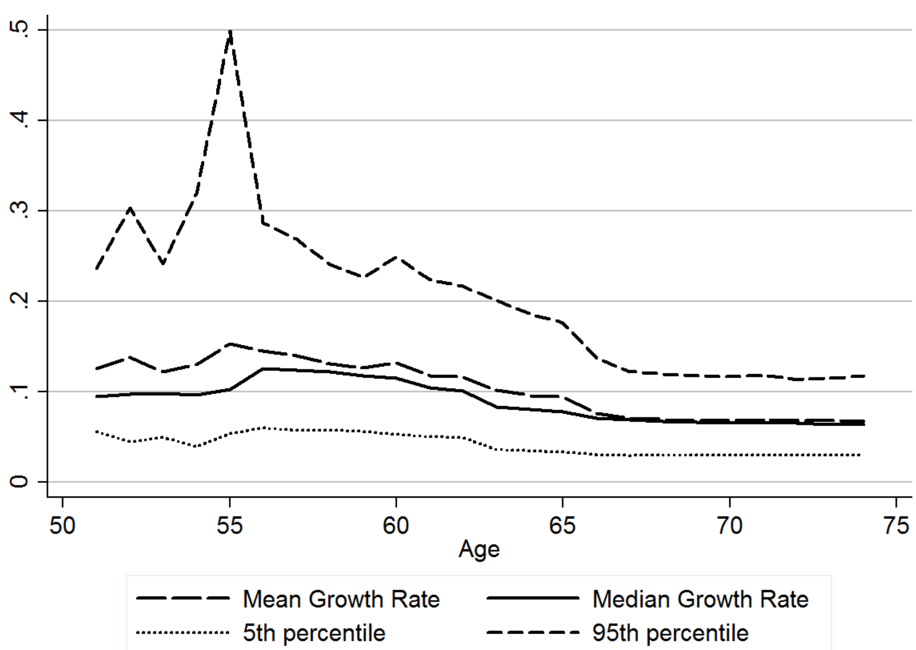

Phelan, 1997; van der Klaauw and Wolpin, 2008). By omitting work histories and pension plan details, these papers focus on a portion of the population that has lower incomes and for which Social Security benefits represent a very important part of retirement wealth. These household will be more likely 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 early claiming rates for Social Security than are predicted by a typical life-cycle model (Hurd, Smith, and Zissimopoulos, 2002; Coile, Diamond, Gruber, and Jousten, 2002; Sass, Sun, and Webb, 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 10 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:

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

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.

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 2.1 and 2.2, 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

[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.

her own benefit level, whichever is greater.[7] The household's budget constraint in period 1, assuming no preexisting assets, is

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

The budget constraint in period 2 is

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

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 2.1 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 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

[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:

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

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

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

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

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 2.1). 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 2.1: Example of a Household Budget Constraint

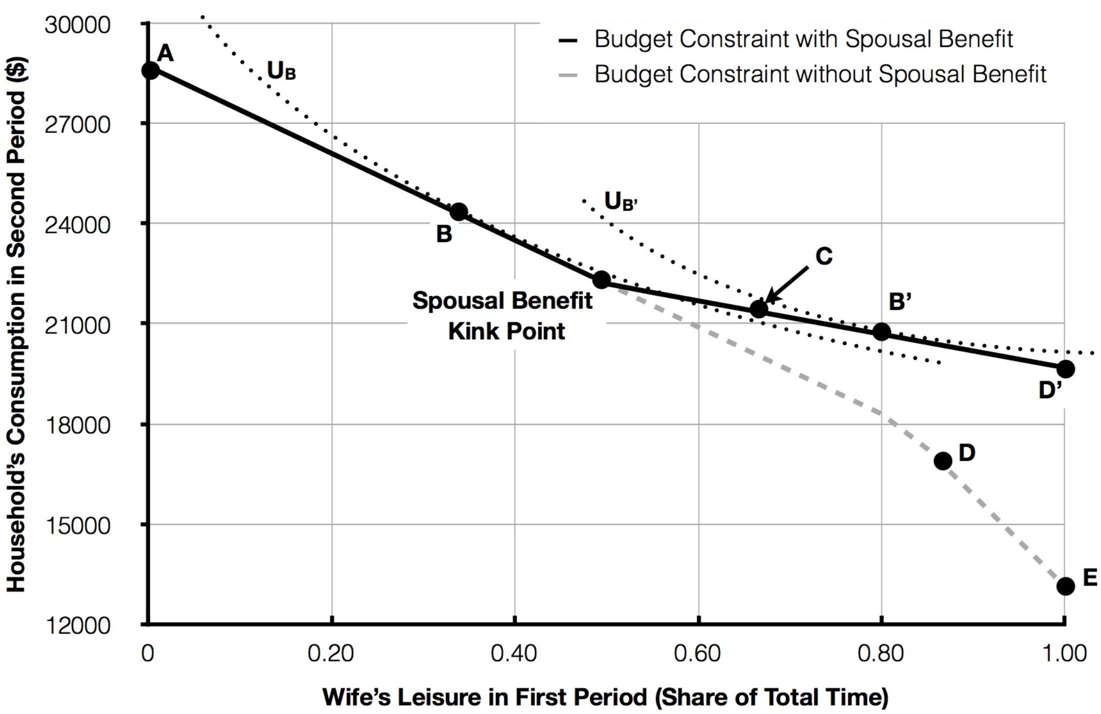

Notes: The baseline budget constraint (solid line) assumes the husband earns $40,000 and saves none of his income, and the wife's potential income was $25,000 if she worked full time and the wife saves 20% of her income.

figure 2.1.

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 2.1.

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 2.1).

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 2.1. In this case wife's work behavior is unaltered by the existence of the spouse benefit, but consumption increases from E to D'.

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 2.2a, 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 2.2 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 2.2.

Figure 2.2: Household budget constraint relative to Husband's Leisure Choice

(a) Dominant Income Effect

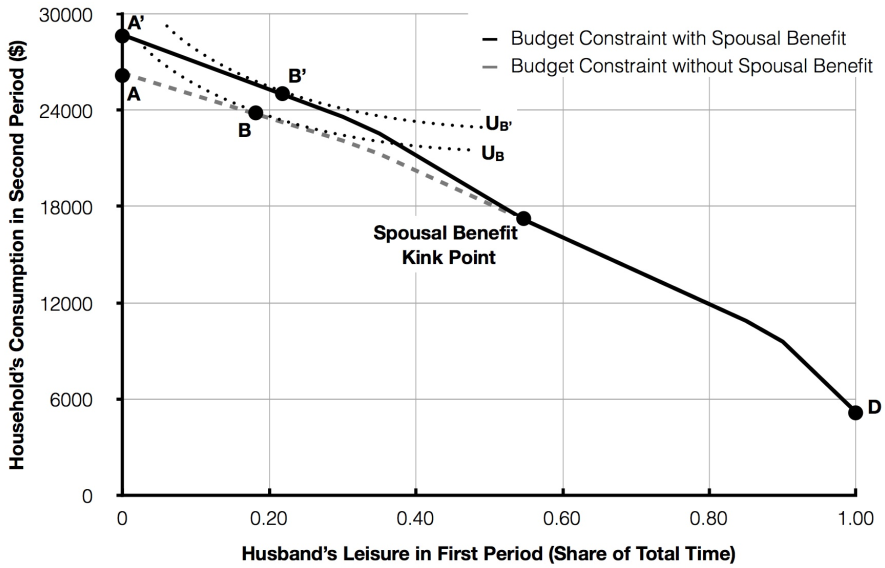

(b) Dominant Substitution Effect

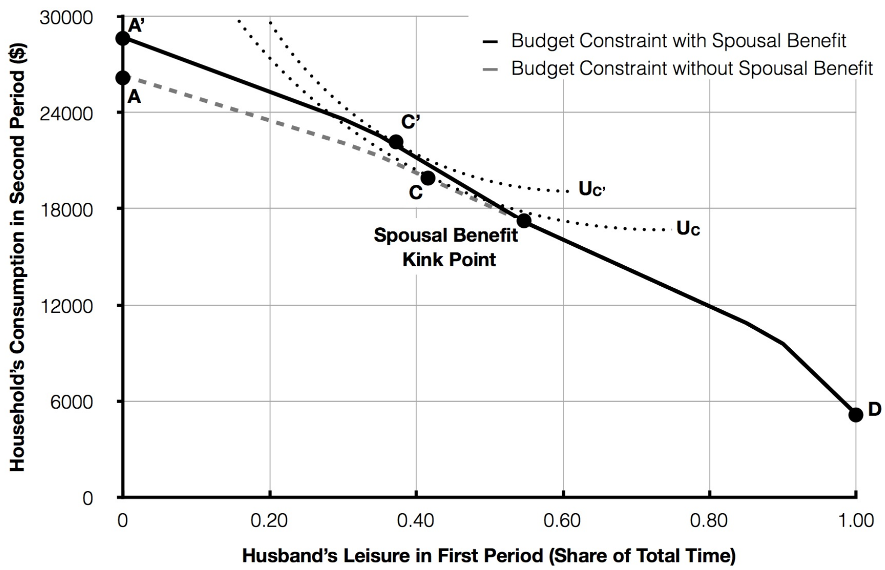

Notes: The baseline budget constraint (solid line) assumes the wife earns $5,000 and saves 20% of her income, and the husband's potential income was $60,000 if he worked full time and the husband saves 20% of his income.

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 2.2a.

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 2.2b.

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.

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 2.2b).

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

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

will have 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

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

the form

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

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

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

$D_{H,t}$ 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 (2004a)). This individual fixed effect for higher value of retirement to an individual, $\varepsilon_i \sim N\left(0, \sigma_{\varepsilon_i}\right)$, 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, 2005a). 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, French, and Jones (2010), households where both members are deceased value 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.

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

Finally, households are borrowing constrained based on their flow of income. Following past work (e.g. Hubbard, Skinner, and Zeldes (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]

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

$$\begin{aligned} V_t\left(X_t\right) \;=\; \max_{C_{h,t},L_{h,t},B_{h,t}} \; & \Big\{U\left(C_{h,t}, L_{h,t}\right) + \delta\left(1 - s^H_{t+1}\right)\left(1 - s^W_{t+1}\right) b(A_{h,t+1}) \\ & + \delta\left(1 - s^H_{t+1}\right) s^W_{t+1} \mathbb{E}\left[V_{t+1}\left(X_{t+1} \mid X_t, t, C_{h,t}, B_{h,t}, L_{h,t}, \text{ wife survives}\right)\right] \\ & + \delta s^H_{t+1}\left(1 - s^W_{t+1}\right) \mathbb{E}\left[V_{t+1}\left(X_{t+1} \mid X_t, t, C_{h,t}, B_{h,t}, L_{h,t}, \text{ husband survives}\right)\right] \\ & + \delta s^H_{t+1} s^W_{t+1} \mathbb{E}\left[V_{t+1}\left(X_{t+1} \mid X_t, t, C_{h,t}, B_{h,t}, L_{h,t}, \text{ both survive}\right)\right]\Big\} \quad , \end{aligned} \tag{3.9}$$

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

[13] C_{min} will depend on whether the household is single (i.e. widowed) or married. As mentioned in §10.1, 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.

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.

4 Data

4.1 Health and Retirement Study

The model in §10 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

[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, Michaud, Smith, and Soest, 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 4.1: 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.

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 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 4.1 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.

[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.

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 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

[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)

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 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%.

[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).

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 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

[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.

[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.

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 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 §10 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

[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

remain consistent with 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^i_{t+1}, 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

sample used in the model's estimation.

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 4.1, 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 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

[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%].

[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.

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.

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.

[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.

Table 5.1 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 5.1 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 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, 2005a), 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 5.1 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 $\gamma_{i,\tau}$,

[28] There is also a separate literature using behavioral questions from the HRS to determine the CRRA coefficient. Barsky, Juster, Kimball, and Shapiro (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 5.1: Preference Parameter Estimates

Parameters based on type

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

Parameters common to all types

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

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 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 (2004a), 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, Skinner, and Zeldes (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, French, and Jones (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 5.1 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 (1986a)), 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 5.1 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 5.1. As a result, the 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 5.2 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 K.2, 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 5.3, 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 K.3 and K.4, 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.

The remaining moments on full-time work and health's influence on participation are

[29] Figure K.1 in Appendix K provides a closer view of the first two thirds to verify that they do indeed match well.

Figure 5.1: Asset Quantiles (by thirds) by Male Age

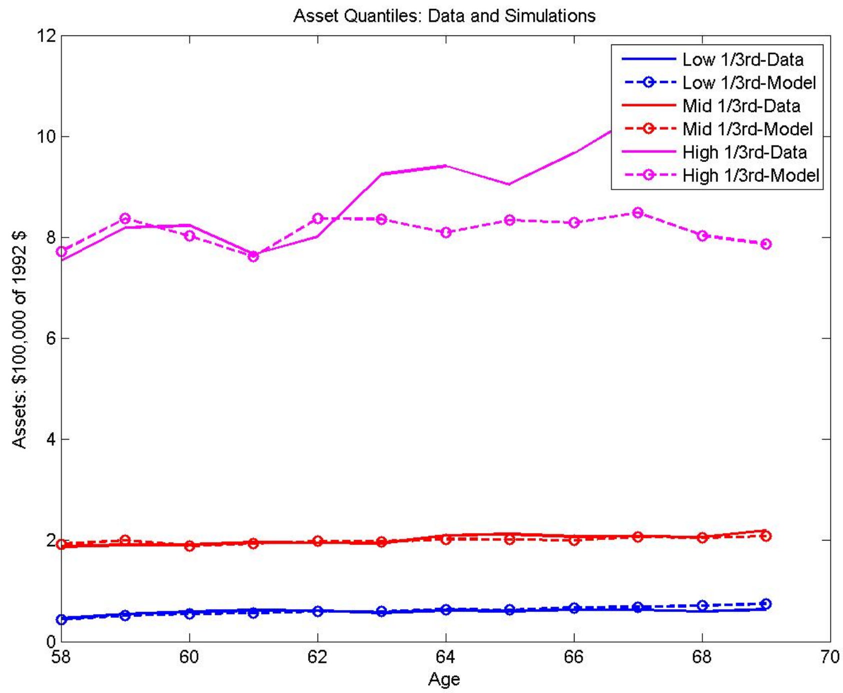

captured in figures K.5, K.6, and K.7 in Appendix K. The model has the most difficulty reproducing the employment composition of the workforce (i.e. part-time versus full-time), which could reflect my coarse discretization of labor supply - using 2 discrete states instead of hours. As expected, figure K.7, 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, 1986a). 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 5.4, 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 5.5 shows the benefit claiming history of the men and women in our sample.

Figure 5.2: Asset Quantile Shares by Preference Type
Type 1 (High Preference for Own Leisure, Low Preference for Spousal Leisure)

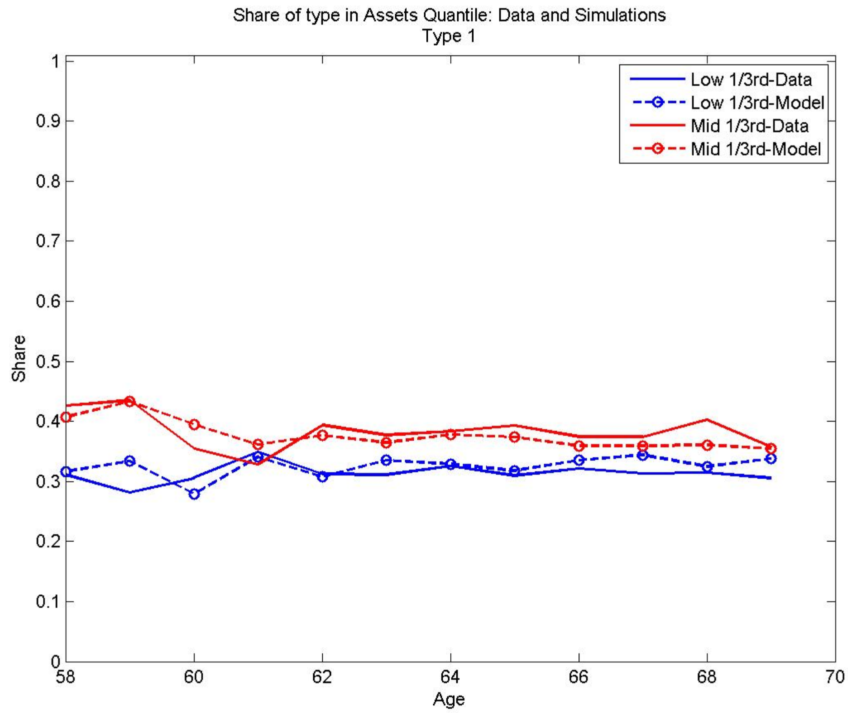

Figure 5.3: Participation by Preference Type

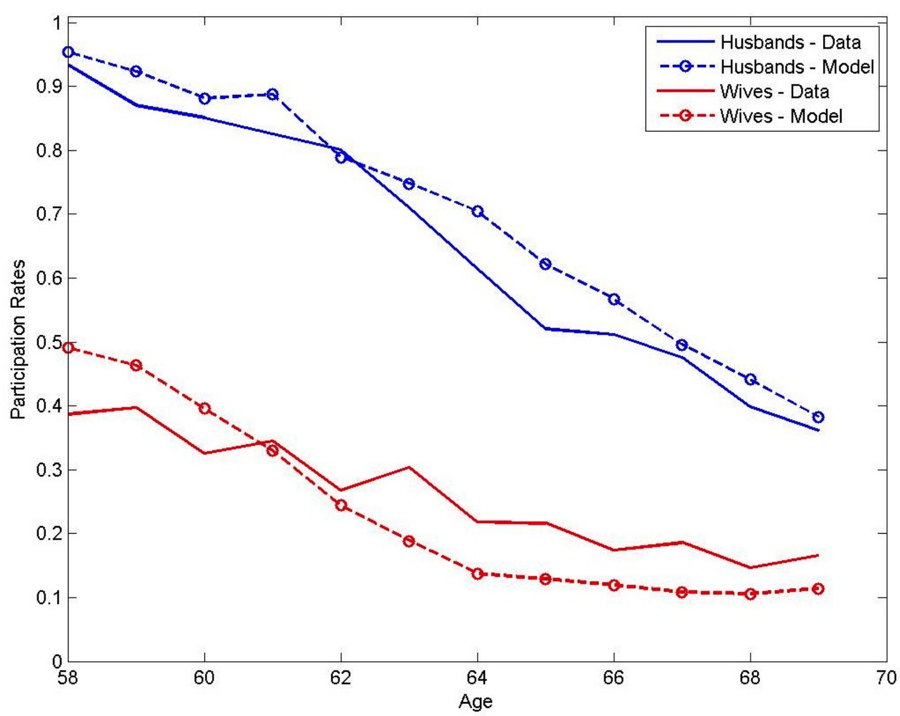

Figure 5.4: Male Labor Force Exit by Age

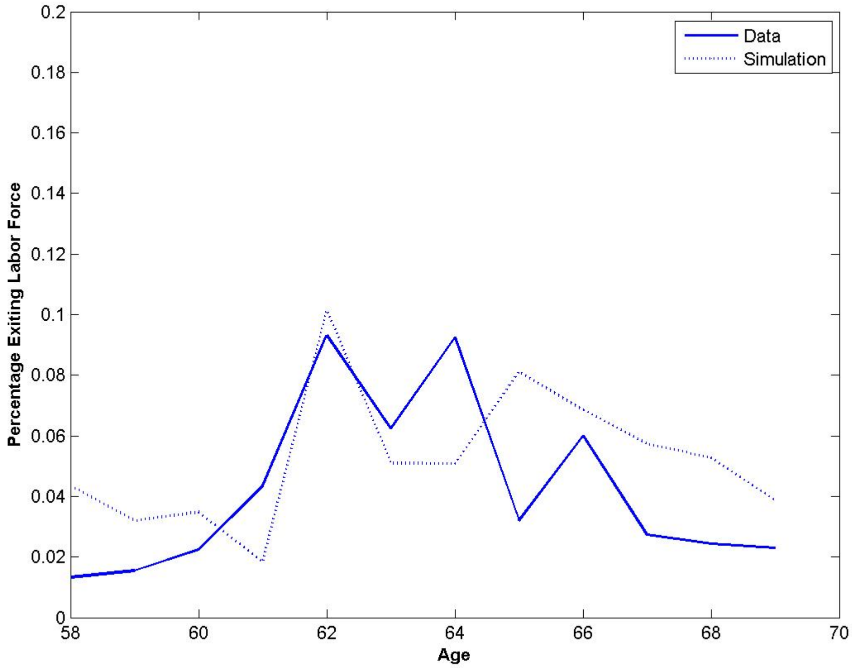

In comparing the benefit claiming rates in figure 5.5 to figure 5.3, it is clear that male participation rates remain high, despite over 50% of the male population claiming their benefit at age 62. The model correctly predicts that a substantial amount of the population claims at 62, and that women are 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, Diamond, Gruber, and Jousten, 2002; Shoven and Slavov, 2013; Sass, Sun, and Webb, 2013). It also correctly captures that very few individuals in my sample actually delay claiming beyond age 65.

Figure 5.5 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

[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

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 5.6 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 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 5.2a 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

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 5.5: Social Security Claiming Rates

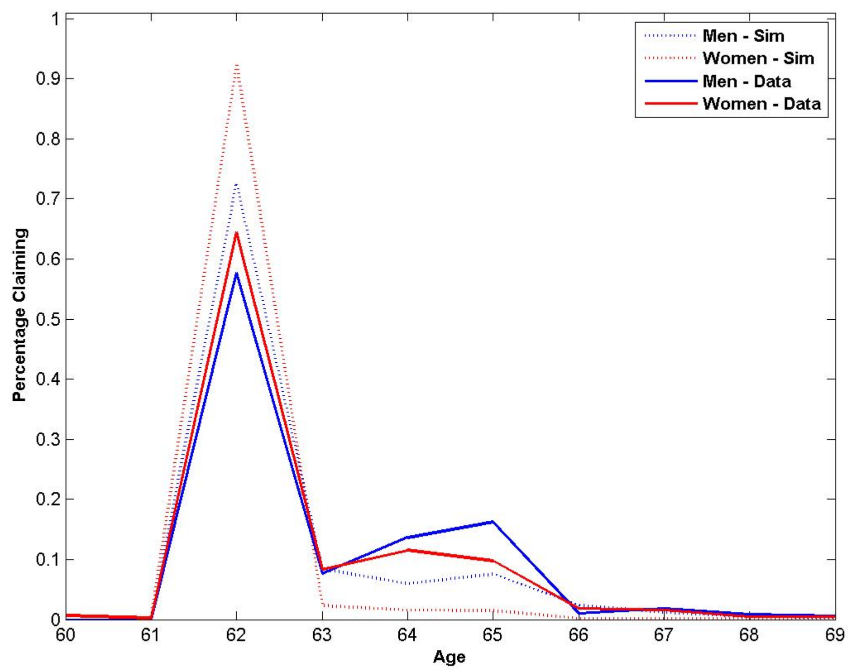

Figure 5.6: Joint Labor Force Exit for Dual Career Households

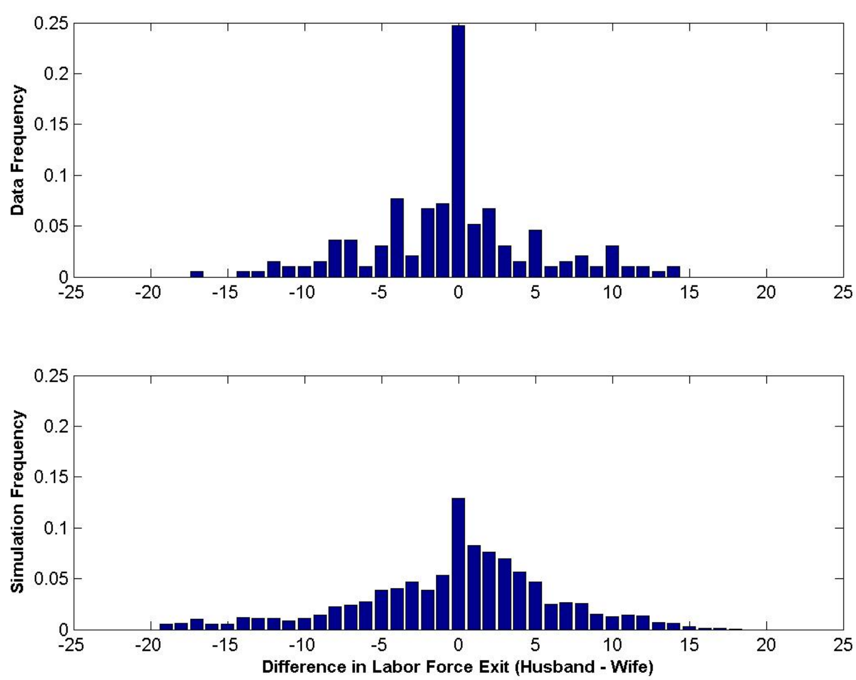

Table 5.2: Comparison of model fit in-sample versus out-of-sample

(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.

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 5.3b 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.

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 6.1a and 6.2a, I report the change in labor force participation from ages 58-69 for men and women. Figures 6.1a and 6.1b 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 2.2).

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 6.1b and 6.2b). 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

Table 6.1: 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 6.2: Policy Experiment Results - Women

(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	0.5943	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 by 100%	by 50%	Reduce Sp. & Surv. Benefits by 100%	by 50%	Increase SS 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 6.3: Policy Experiment Results - Change in Benefits paid over lifetime

Sex	Baseline	Reduce Spouse's Benefits by 100%	by 50%	Reduce Sp. & Surv. Benefits by 100%	by 50%	Increase SS 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.

the simulated claim and work histories, I can estimate the direct reduction in benefits paid through lower benefit 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 6.3 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 6.3, 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.

Figure 6.1: Labor Supply Response to Social Security Benefit Changes

(a) Men

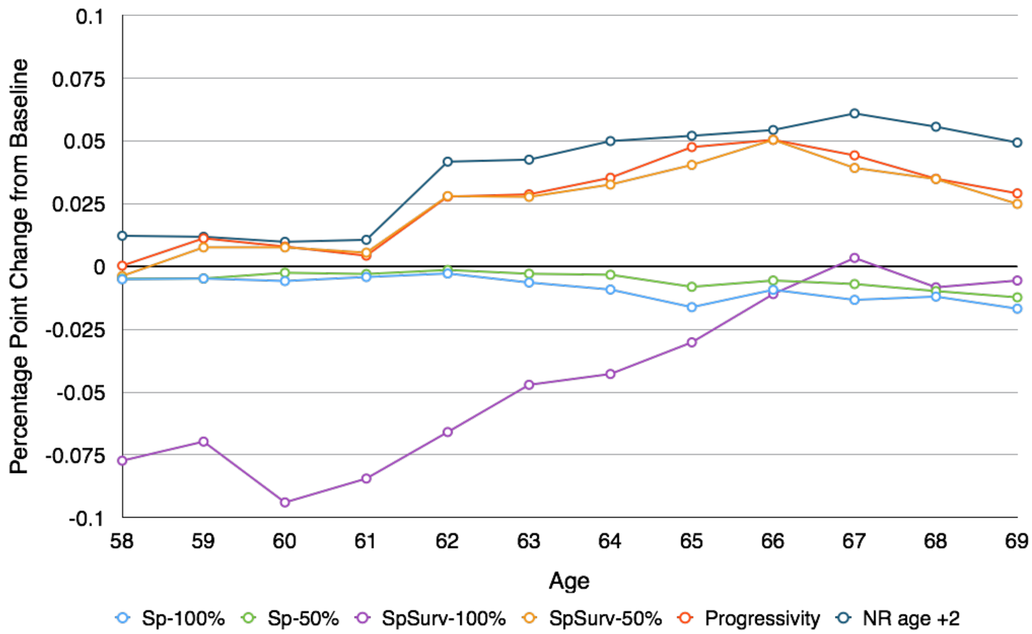

(b) Women

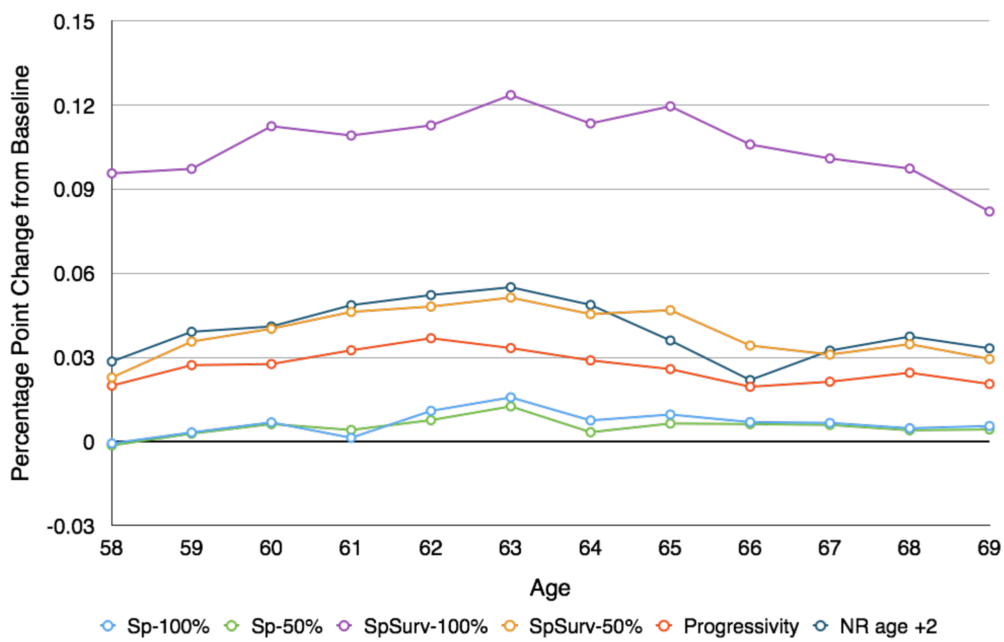

Notes: This figure presents results from counterfactual experiments of changes in Social Security's benefit structure based on the model of household savings, labor supply, and benefit claiming. The model incorporates uncertain longevity, health, and medical expenses. Sp-100% refers to eliminating the 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.

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 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 6.1a). 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 6.1a and 6.2a, 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 6.1b and 6.2b, 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 6.1a and 6.2a 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.

Chapter 2:
Old Age Divorce and Labor Supply

8 Introduction

Divorce is an extremely disruptive event, causing a redistribution of finances between household members and altering long-term objectives, such as retirement. When separation occurs near or after retirement, individuals have a limited time horizon to adjust their labor supply and savings decisions. Divorce for married women aged 55-59 has increased from 3.9 per thousand in 1980, to 4.8 per thousand in 1990, to 13.1 per thousand in 2010, yet little research has been done on divorce's effect on individual's labor supply decisions. In 2010, a married woman had a 14.2% chance of divorcing in her 50s, up from a 6.3% probability in 1990. This paper provides a descriptive analysis of late-life divorce using the Health and Retirement Study, a panel survey of households over fifty. In doing so, I ask the question: what is the immediate effect of divorce on men and women's labor supply and the long-run impact of these separations on retirement?

Divorce means that a two-person household has become two one-person households. Assets are split, not necessarily in half, and incomes are based on the individual: his or her education, past employment, and life choices made during the marriage. Households that separate late in life face a shortened time period where they can be gainfully employed and save for retirement to make up for the asset loss. Two one-person households lose the economies of scale enjoyed by a two person-household. The separated households generally require two living locations, two cellphone plans, two health insurance plans, and more. Both new homes need to be cleaned and two separate meals must be prepared. Divorce also means social and psychological upheaval. Forty percent of divorcees between age 40 and 69 express concern about "the uncertainty of their future" and 29% suffer through depression (AARP, 2004). At the same time, 41% report that the best thing about life after divorce is the "freedom/independence to do what I want to do." Motivated by a household life-cycle model of near-retirement decisions, I separate the effects of divorce into changes from loss of household assets, future spousal income, and economies of scale, as well as changes in the divorcée's willingness to work.

I find that women increase their labor force participation by up to 17.1% in the 8 years following separation. I also find little evidence that women anticipate divorce and alter their labor supply or savings behavior before the divorce occurs. Divorcing men respond by decreasing their hours worked, similar to their younger counterparts (Johnson and Skinner, 1986). When parsing divorce's effect on labor supply into contributing factors, asset loss is found to have the most significant effect on female labor supply: the average non-working woman would be 6.0 to 10.9 percentage points more likely to work per $100,000 lost in divorce in the 8 years following separation. Other effects, notably the loss of future spousal income, are found to have an insignificant impact on a woman or man's decision to return to work after a divorce, or to delay retirement.

The main contributions of this paper are twofold. First, there are few studies on the financial effects of divorce at older ages, although there are plentiful anecdotes.[31] I document a sharp increase in divorce for couples in their 50s and 60s, while also demonstrating the substantial transitional changes to households' assets, work, and spending choices. Between 1990 and 2010, divorce rates have jumped: doubling for women 50-54, quadrupling for women 65 and above. Put another way, divorce rates for women 50-54 in 2010 are similar to divorce rates of women 35-39 in 1970. Following a divorce, women retain only 35-43% of the non-housing assets of the pre-divorce household, while the average post-divorce spending of a newly single household is 78.9% of its pre-divorce level. Using the Census' poverty definition, the percentage of divorcing women in poverty rises from 6.5% before the divorce to 18.6% after the divorce.

Secondly, this paper analyzes the contributing factors to short and long-term changes in women and men's labor supply, and consequently her or his ability to consume and avoid poverty. A household life-cycle model is developed that captures the effect of divorce and widowhood on the labor supply of individuals who are near-retirement. Within the context of a life-cycle model, a negative shock to savings combined with higher relative expenses should increase the likelihood of the individual returning to work or delaying retirement. I use the 1992-2010 interview waves from the Health and Retirement Study to test the implications of the model. The change in asset levels following a divorce is the largest determinant of a women's re-entry into the labor force and whether they will delay retirement. An analysis of widowhood is included, as it offers an insightful comparison since household assets are generally not lost in widowhood.

The paper proceeds as follows. Section 9 describes changes to the frequency of old-age divorce and uses the Health and Retirement Study and its supplements to chart demographic and financial changes around the divorce. The next section introduces a household life-cycle model of labor supply and consumption that captures the transition from a two-person

[31] For an excellent collection, see *Calling it Quits: Late Life Divorce and Starting Over* by Deirdre Bair.

household to two one-person households. The model is then used to motivate the empirical specification in section 11, which is used to analyze short run labor supply responses around the time of divorce. Section 5 adapts the specification to capture how retirement is affected by divorce. Section 6 concludes.

9 Background

Old-age divorce, like widowhood, brings to mind many issues related to the well-being of individuals in old age. While there are many stories and anecdotes, there exists little direct evidence of what happens to husbands and wives after divorce. Using the Health and Retirement Study, I am able to report a number of statistics before and after the household separation, including assets, income, disability status, size of household, rental/home ownership, and poverty status. It is particularly interesting to see how these develop in the survey waves following divorce. In all cases, I also compare the divorced to widows, a notable comparison group since they are also making the transition from a two-person to one-person household.

This section is intended to provide a better understanding of financial and demographic statistics around the time of the divorce. I describe the Health and Retirement Study, focusing on the features which I use for the analysis in later sections. The advantage of the Health and Retirement Study is that it follows the same households from their introduction into the study (i.e. households where one member was between 51 and 61 in 1992, with additional cohorts between ages 51-56 added in 1998, 2004, and 2010) until present day. Despite the size of the Health and Retirement Study, the sample of affected households is still small when examining a rare and critical event like divorce. I augment the Health and Retirement Study with data from larger surveys such as the U.S. Census, the American Community Survey, and the Center for Disease Control and Prevention's Vital Statistics Report of the Divorce Registration Area. These large cross-sectional surveys can inform us regarding broad trends in divorce, but they cannot capture how an individual's situation changes after a divorce.

9.1 Data

This paper uses the Health and Retirement Study (HRS), a panel survey of over 29,000 respondents where one household member is greater than 50 years old. The HRS tracks labor force participation, health, and other key demographics variables. I reduce the sample to individuals who are at-risk for divorce and widowhood between the ages of 50 and 70 for the main analysis. This means that I exclude households that have never been married and households that divorced or were widowed prior to 50 and then never remarried.

I measure time of separation based on reports of the month and year the individual was divorced or widowed. In the case of divorce, I will examine waves prior to the official divorce event and observe whether the individual reports being separated or divorced from his or her spouse in previous waves.[32] If this is true, then I move the time of the separation earlier, since the time of separation should reflect when the individual reevaluates his or her expectations of future income and make adjustments to labor supply. Since separations do not occur at the survey interview date, there will be heterogeneity in the time between separation and when the individual is interviewed. This may lead to unclear results in the period immediately following the separation if the individual has not had time to react, and in the period immediately before separation if the effect is at all anticipated.

Table 9.1 presents summary statistics for the sample. Divorced women are more likely than both widowed women and the general female population to have participated in the labor force within the past five years, and have worked more years in their lifetimes. They are also more likely to be white, have fewer children, and are closer in age to their husbands.

Widowed women are generally from longer marriages, less educated, are much younger than their spouse, have more children, and claim their Social Security benefits sooner. At the time that their husband dies, widows are much more likely to be out of the workforce, and will be less likely to remarry when he dies. Divorced and widowed women are more likely than continuously married women to apply for disability, and their wages are lower on average than the continuously married female population.

Divorced men are more likely than continuously married and widowed men to be white, claim disability, have lower wages, and delay claiming their Social Security benefits. They also are older on average relative to their ex-spouse. Widowed men are closest in age to their ex-spouse, have the fewest years of education, claim Social Security the earliest, and are much less likely to have participated in the labor force in the year prior to their first interview by the HRS.

Death is an event that is more likely to happen later in a couple's life, while divorce is an event that is more likely to happen earlier in a couple's life. In appendix M, I standardize the same statistics by age. In order to do this, that data are weighted based on the age distribution of divorcées. The general trends mentioned above are robust to age standardization.

In estimating the empirical specification in §11.1, I include a binary measure of bad health, based on the respondent's self-reported health at each survey wave. An individual

[32] Appendix L provides greater detail on how this is done.

Table 9.1: Summary Statistics from HRS Sample

	Women			Men		
	Population	Divorced	Widowed	Population	Divorced	Widowed
Age (Mean)	53.83	51.54	55.03	57.37	54.14	57.34
Before Separation		56.24	60.04		58.07	60.98
White (%)	0.83	0.86	0.72	0.82	0.87	0.84
Education in Years	12.47	12.60	11.84	12.42	12.66	10.88
Average years worked	21.79	23.34	23.17	35.16	32.51	33.51
Before Separation		26.05	25.26		34.16	36.50
Marriage Length (Mean)	28.15	19.61	28.73	28.11	17.66	27.45
Before Separation		20.33	31.96		17.98	31.09
Marriage Length (Median)	30.40	21.40	32.45	30.40	18.00	30.50
Before Separation		20.40	34.90		16.90	34.40
Remarried after 3 waves (%)		0.32	0.16		0.47	0.43
Age difference with spouse	3.50	2.58	4.92	-3.40	-5.42	-2.98
Number of Children	2.80	2.69	3.21	2.89	2.67	2.92
Worked last survey wave	0.73	0.83	0.75	0.81	0.90	0.74
Before Separation		0.62	0.43		0.71	0.52
Worked in last five years, but not last survey wave	0.10	0.10	0.14	0.10	0.06	0.18
Before Separation		0.17	0.22		0.15	0.21
Not worked in last five years	0.16	0.07	0.11	0.09	0.04	0.07
Before Separation		0.21	0.35		0.14	0.26
Hourly Wage	25.30	18.27	16.47	37.77	26.13	31.53
Before Separation		19.92	14.35		26.65	19.39
Ever Applied Disability (%)	0.14	0.26	0.25	0.16	0.25	0.24
Age begin Social Security	62.42	62.08	61.80	63.23	63.35	62.39
Sample Size	8409	235	239	8475	237	102

Notes: Author's calculations, Data from Health and Retirement Study. Separations occurring between age 50 and 70. Unless otherwise indicated, statistics refer to the first wave.

is considered in bad health if he or she reports being in poor or fair health. Measures of spousal labor force participation and full-time status are based on the spouse's responses. For asset measures, I use imputed values of individual's defined contribution plans and RAND's HRS measures of household assets excluding primary residence. To create my measure of household assets, I sum the net value of the household's other residences, the net value of businesses, and the value of vehicles, stocks, mutual funds, other investment trusts, checking accounts, certificates of deposits, savings accounts, government savings bonds, treasury bills, bonds, bond funds, defined contribution plans, and any other reported savings, and subtract other debt.

9.2 Frequency of Divorce

Divorce is becoming more common at older ages. Figure 9.1 presents the divorce rate for married women by age group. It combines data from the CDC's Vital Health Statistics for 1970, 1980, and 1990, with data from the 2010 American Community Survey (ACS). We observe that divorce at all ages over 50 has increased in the past 40 years. Divorce rates for women between age 50 and 54 have more than doubled since 1990 from 8.2 divorces per thousand married women to 17.4 divorces. Divorces for women 65 and older have nearly quadrupled in the last 20 years, rising from 1.4 divorces per thousand married women to 5.5 divorces. An alternative perspective is that the 2010 divorce rate for women in their late 50s is greater than the 1970 divorce rate for women in their early 40s.

This is in contrast to a decline in the aggregate divorce rate for married couples since 1980 (see Stevenson and Wolfers, 2007, their figure 1). One possible reason for the surge in old-age divorce is that as women are becoming more educated and participate in the labor force, they delay marriage and consequently the possibility of divorce. The marriage rate for women at age 30 has fallen from a peak of 85.3% of the population in 1960 to 51.8% in 2010 (see figure 9.2). Alternatively, there could be cross-cohort differences in people's expectations for marriage and willingness to divorce. The baby boom generation (i.e. those who were age 16-34 in 1980, and 26-44 in 1990) divorced more frequently when they were young, and now that they are older (i.e. age 46-64 in 2010), the willingness to divorce appears to have continued into old age.

Most large data sources (i.e. U.S. Census, etc.) have historically kept track of the number of people whose current marital status is divorced. This number fails to capture people who have since remarried, or when the divorce occurred. For example, a divorced woman at 55 could be a woman who divorced at 25 and never remarried, or could be a woman who divorced just before being interviewed at age 55. Nonetheless, this measure does provide

Figure 9.1: Women: Frequency of divorce by age group

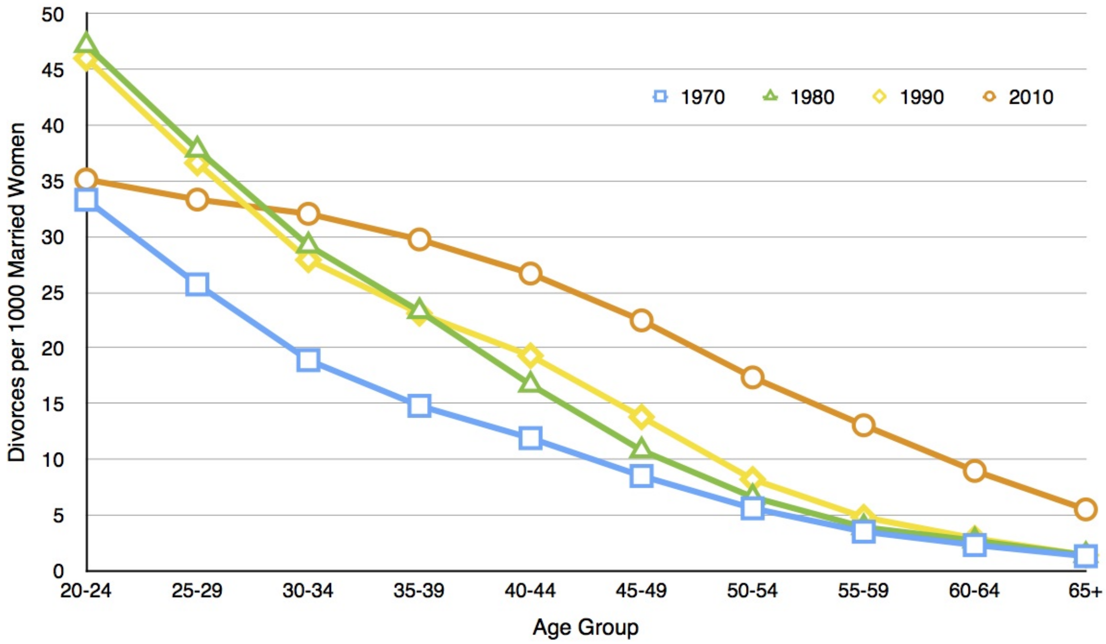

Source: 1970-90 from Clarke (1995), 2010 from author's calculations using the American Community Survey. 2000 is missing because data was not collected by any source (see Kennedy and Ruggles, 2013 for additional information on missing divorce measures).

---

an understanding of the growing number of divorcées. Figure 9.3 shows that the percentage of aged 55 women, whose current marital status is divorced, has risen from 11.0% in 1980 to 22.7% in 2010. The trend in increasing rates of women whose current marital status is divorced reflects a cohort effect: being divorced and staying single is more frequent in the baby-boom generation. Meanwhile, widowhood rates at age 55 have fallen from 11.5% to 5.5%, and female marriage rates at age 55 have fallen from 72.9% to 62.4%.

This is not the only change in marital trends. Marriages also appear to be lasting longer. Figure 9.2 shows that marital rates after age 62 have significantly increased since 1940 (between 1880 and 1940 they were constant). For example, the percentage of women who are currently married at age 75 has jumped from 33.1% in 1980 to 45.4% in 2010. The decrease in widowhood at age 75, from 54.9% in 1980 to 36.2% in 2010, more than accounts for this change. The increase in old-age married women is driven primarily by the large decrease in widowhood.

Men have experienced similar trends to women before retirement age. At younger ages, the percentage of men married has decreased (at age 30: 81.7% in 1960, 40.5% in 2010). At near-retirement ages, fewer men are likely to be married. Men aged 55, whose current marital status is divorced, have increased from 5.3% in 1980 to 19.5% in 2010.

Figure 9.2: Women - Percentage whose current marital status is married (by age)

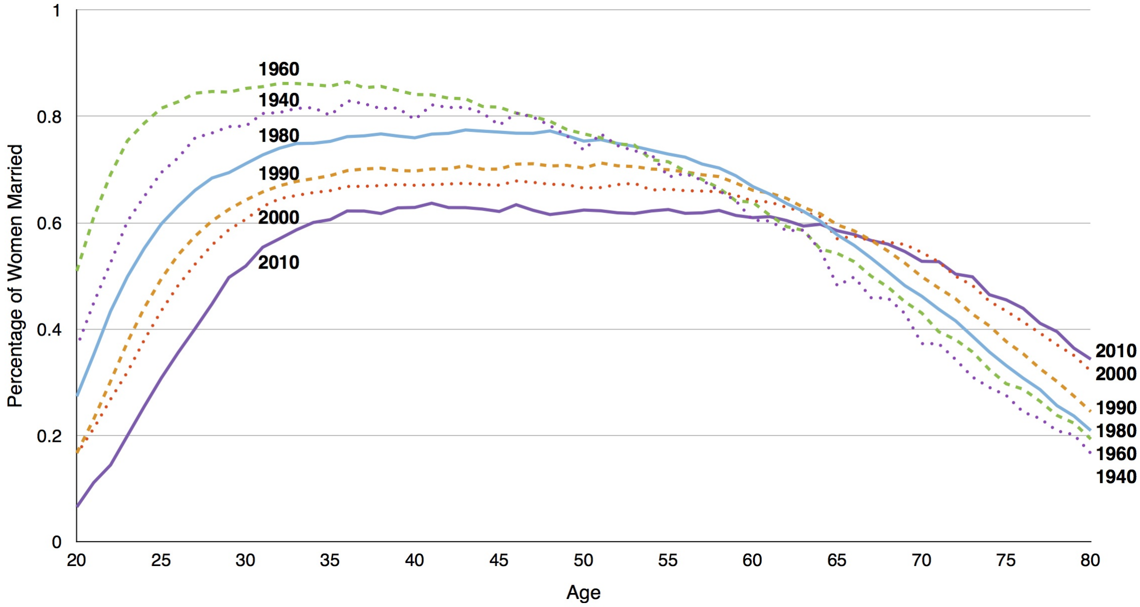

Source: Author's calculations for 1940-2000 using U.S. Census (1% samples), 2010 using the American Community Survey.

Figure 9.3: Women - Percentage whose current marital status is divorced (by age)

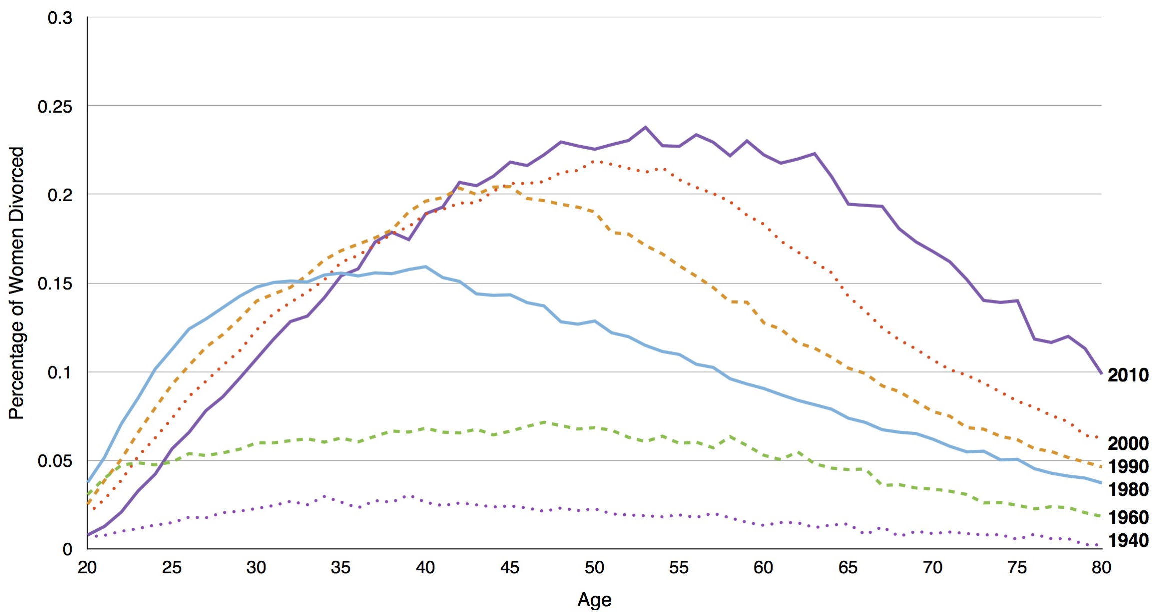

Source: Author's calculations for 1940-2000 using U.S. Census (1% samples), 2010 using the American Community Survey.

We previously observed that more women in younger generations are married after age 62. This effect is exaggerated for men. Male marital rates in 2010 stagnate between ages 38 and 49, but continue to rise after age 50, peaking at age 73 when 75.2% of men are currently married. Men's greater propensity to remarry after divorce, combined with longer lifetimes, mean that a greater percentage of retirement-age men are ending up married. Similar to figures 9.2 and 9.3, figures P.1-P.4 in the appendix chart the changes in widowhood for women and men by age and cohort, and male marriage and divorce rates by age and cohort.

If future generations continue to delay marriage until later in life, and people's willingness to divorce remains unchanged, then we should expect that divorce near and after retirement will remain common in the future. These aggregate studies, while helpful for recognizing generational trends, do not allow us to observe the household around the time of the divorce. The HRS's panel of households has captured people approaching retirement, including their experiences with divorce. The HRS continues to follow both household members after divorce. In the next section, we will observe how an individual's observable characteristics change from before to after the household's separation.

9.3 Household changes before and after separation

Transitioning from a two to one-person household results in significant changes to a person's finances as well as their future expectations. In this section, I use the fact that we can follow households in the HRS before and after the separation (whether it is through divorce or widowing) to explore changes to assets, rental rates, labor force participation, disability benefit claiming, and the person's self-reported expectation of working at ages 62 and 65. The summary statistics before and after a household separation are presented in table 9.2. All assets are reported in 2010 dollars.

Assets. Divorcing women experience a mean and median asset shock. The mean asset shock is $161,000, leaving them with, on average, only $119,000. As is true of the general population, the asset distribution is highly skewed. Median non-housing assets for the general married female population are only $94,000, while the mean is $244,000, implying that a few households have substantial net assets. Divorcing women usually end up with between 35-43% of the before separation non-housing, net assets. Figure 9.4 shows the change in non-housing net assets around separation for both men and women.[33] Divorcing individuals' asset levels do not return to their pre-divorce state. Divorced women's asset levels remain generally unchanged, while divorced men begin to accumulate assets.

Widowed women experience little median asset loss because the survivor generally retains

[33] Recall that women are generally younger than men at separation, implying that the asset profile of women pre-divorce should appear higher.

Table 9.2: Before and After Separation - Summary Statistics from HRS Sample

		Women			Men		
		Population	Divorced	Widowed	Population	Divorced	Widowed
	Number in Household	2.89	2.95	2.95	2.89	2.99	3.12
	Before Separation		2.70	2.56		2.76	2.72
	After Separation		1.82	1.95		1.70	1.99
Excluding Housing	Assets ($100,000 - Mean)	2.44	1.97	1.55	2.59	2.47	1.15
	Before Separation		2.80	1.31		2.66	1.53
	After Separation		1.19	1.39		1.35	1.45
	Assets ($100,000 - Median)	0.94	0.59	0.46	0.99	0.72	0.31
	Before Separation		0.91	0.34		0.93	0.25
	After Separation		0.32	0.30		0.40	0.29
Including Housing	Assets ($100,000 - Mean)	3.50	2.85	2.29	3.61	3.47	1.70
	Before Separation		3.71	2.23		3.60	2.18
	After Separation		1.82	2.31		2.05	2.42
	Assets ($100,000 - Median)	2.07	1.41	1.12	2.12	1.51	1.10
	Before Separation		1.85	1.14		1.74	1.13
	After Separation		0.79	1.09		0.90	1.33
	Renting (%)	0.10	0.11	0.16	0.11	0.13	0.21
	Before Separation		0.08	0.17		0.13	0.19
	After Separation		0.38	0.20		0.41	0.17
	1st wv. After Separation		0.35	0.21		0.33	0.20
	Working (%)	0.60	0.71	0.56	0.73	0.85	0.66
	Before Separation		0.62	0.41		0.71	0.51
	After Separation		0.64	0.40		0.63	0.46
	1st wv. After Separation		0.60	0.34		0.59	0.33
	Applied for Disability (%)	0.02	0.06	0.04	0.06	0.07	0.10
	Before Separation		0.07	0.12		0.12	0.16
	Ever Apply	0.14	0.26	0.25	0.16	0.25	0.24
	Prb. Work at 62 (Self-rpt)	36.88	48.07	45.19	52.29	56.52	57.10
	Before Separation		49.03	38.28		56.57	45.00
	After Separation		52.58	33.42		54.33	33.03
	Prb. Work at 65 (Self-rpt)	19.43	25.03	21.78	31.26	31.71	32.34
	Before Separation		23.24	19.46		34.14	26.18
	After Separation		31.48	16.03		36.34	15.11
	Sample Size	8409	235	239	8475	237	102

Notes: Author's calculations, Data from Health and Retirement Study. Separations occurring between age 50 and 70. Unless otherwise indicated, statistics refer to the first wave. Separations occur between waves. The interview wave response before the separation is marked as "Before Separation", and similarly for the response "After Separation". "1st wv. after separation" refers to the interview response in the first full wave since the separation occurred (minimum of 2 years). All assets are reported in 2010 dollars.

Figure 9.4: Non-housing, net assets by timing of separation

(a) Assets in waves before and after separation

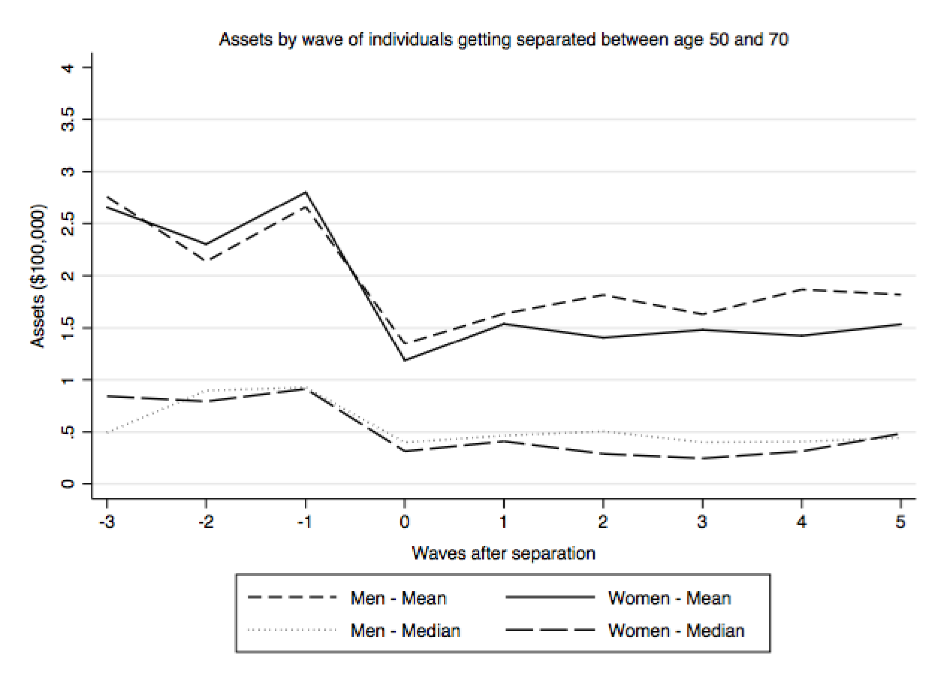

(b) Assets in waves before and after widowing event

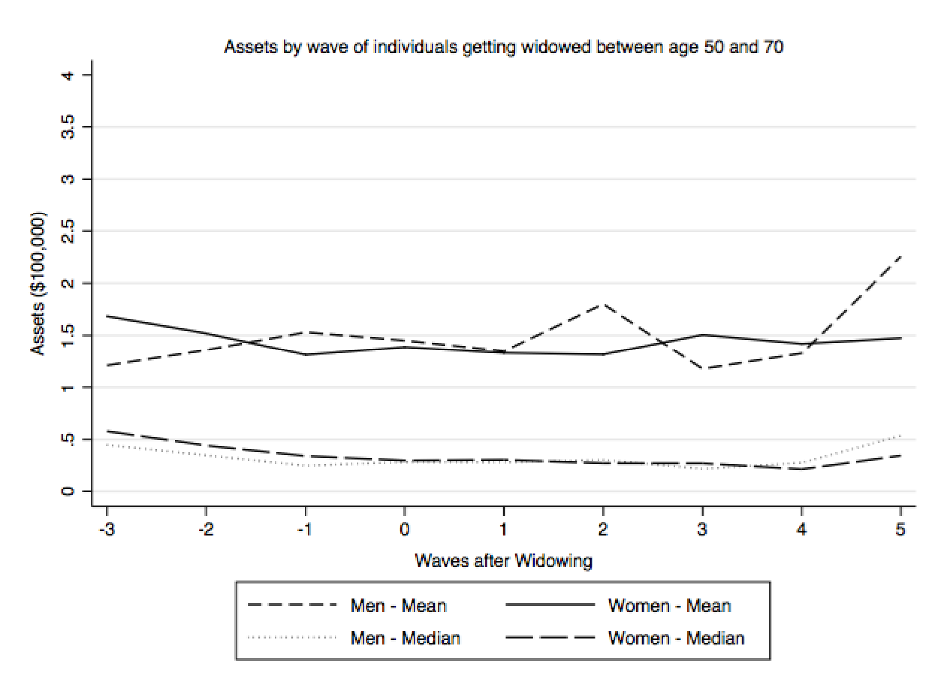

Notes: Author's calculations. Data from Health and Retirement Study. Assets are reported in 2010 dollars. Separations/losses occur between waves. Period 0 refers to the first interview after the separation/loss. Period -1 refers to the last interview before the separation/loss. Total spending includes durable and nondurable household spending.

the household's assets. We do observe a mean asset gain for widows in the wave following the widowing event, but assets experience little subsequent growth on average. Assets change little for men before and after the widowing event.

Household size does not alter any of these relationships. In appendix N, I examine how changes in household size affect the change in assets around the separation/loss. Conditioning on households that transition from a two person household to a one person houshold, none of the previously described relationships are altered in an important way.

Renting/Owning. Households that divorce or widow between age 50 and 70 are more likely than the rest of the population to rent their home. Prior to the separation many will have become home owners, but divorce leads to a 30 percentage point jump in rental rates for women and 28 percentage point jump in rental rates for men. Rental rates remain high relative to their peers in the waves following the divorce or widowing event. Rental rates remain relatively constant for widows and widowers.

Labor force participation. Households that divorce between age 50 and 70 are more likely than the rest of the population to be working, while widows are less likely to be working during their first HRS interview. Prior to the separation (leading to a divorce), female labor force participation declines from 71% to 62%. After the divorce, female labor force participation increases to 64%, before beginning to decline again. There are two opposing forces here: the ongoing retirement of individuals in this age group, and divorced women evaluating their participation in the labor force. For widows, widowers and divorced men, the decline in labor force participation continues after the separation or loss. Standardizing by age, this relationship is preserved for men. Women, however, are more likely to work after divorce or widowhood. This could imply that the generational shifts in women's labor force participation may also affect how they respond to life events. In §11.1, I will consider how other changes correlated with divorce, such as asset loss and the change in future spousal income will alter a person's labor force participation.

Disability. The HRS records when an individual applies for disability benefits either through SSDI or SSI. A notable statistical regularity is that divorced individuals are more likely to eventually apply for disability in the years following separation, while widowed individuals are much more likely to have applied for disability prior to the separation.

Retirement. The HRS asks an individual his or her expected probability of working at ages 62 and 65. Women in this sample are more likely than the rest of the population to expect to work at age 62 and 65. Divorcing women increase their belief they will work at age 65 by 8.3%, while their expected likelihood of working at 62 increases by only 3.6%. This is an interesting effect, implying that perhaps it is a certain group of divorcées whose retirement is delayed as a result of the divorce. I will consider this in my analysis in §11.1.

Table 9.3: Total household spending (in $1000s) relative to separation/loss

(a) Women

Waves relative to separation	Divorcées Total		Divorcées Nondurables		N	Widows Total		Widows Nondurables		N
-3	55.05		54.85		6	38.38		38.04		81
-2	48.76		48.36		12	38.01		37.81		149
-1	46.40		46.20		33	38.13		37.86		216
0	36.60		36.19		33	37.30		37.04		216
1	24.86	**	24.61	**	16	33.53		33.37		128
2	20.49	**	20.41	**	9	36.39		36.16		71
3	17.15	***	16.57	***	5	27.91	***	27.60	***	21

(b) Men

Waves relative to separation	Divorcées Total		Divorcées Nondurables		N	Widowers Total		Widowers Nondurables		N
-3	50.04		49.86		10	44.70		44.70		22
-2	42.20		41.97		21	33.21		33.21		45
-1	44.42		44.09		47	34.97		34.97		71
0	35.32	*	34.91	*	47	34.60		34.60		72
1	25.91	***	25.76	***	21	28.28		28.28		33
2	16.98	***	16.96	***	8	26.96	*	26.96	*	14
3	21.47	***	21.20	***	4	26.96		26.96		5

Notes: Spending in thousands of nominal U.S. dollars. Data from RAND's Consumption and Activities Mail Survey, a supplement to Health and Retirement Study. Separations occurring between 2001 and 2009 (waves 5-9 of CAMS). Data weighted by person weight in wave prior to separation/loss. Separations occur between waves. Period 0 refers to the first interview after the separation/loss. Period -1 refers to the last interview before the separation/loss. Total spending includes durable and nondurable household spending. Nondurable household spending includes mortgage payments, home/renter's insurance, property tax, rent, electricity, water, hear, phone/cable/internet, auto financing charges, auto insurance, health insurance, house/yard supplies, home maintenance, grocery, dining out, clothing, gasoline, vehicle servicing, drugs, health services, medical supplies, vacations, tickets, hobbies, contributions, and gifts. Durable spending includes refrigerator, washer/dryer, dishwasher, television, and computers. Vehicle purchases are excluded. The asterisks refer to a Wald-test of whether the wave's spendings is equivalent to spending in the wave prior to separation. Significance reported at the 1% (***), 5%, (**) or 10% (*) level, for waves 0 through 3.

Consumption/Spending. Starting in 2001, the HRS conducted a supplementary mail survey between interview waves called the Consumption and Activities Mail Survey (CAMS). The survey captures annual spending by households, and the RAND Corporation produces a data product that cleans this data and sorts it into aggregate, durable and nondurable consumption. CAMS also tracks marital status, so I use changes in martial status between CAMS interview waves as a means of tracking divorce and widowing events.[34]

Table 9.3 reports total and nondurable spending levels in the waves surrounding the separation/loss. Total household spending for women decreases from $46,400 in the interview wave prior to the household separating to $36,600 in the first interview after the separation: a decrease of 21.1%. In the waves following the separation, spending declines further and becomes significantly different from the wave prior to the separation. The sample sizes, however, are small and so we should be concerned about sample composition biasing spending levels. Similarly, men experience a 20.5% decrease in spending around the time of divorce. Both women and men's spending continues to decrease in the waves following the separation. This may imply that it takes several years for spending to stabilize. This could be in response to it taking time to sell the couple's house, or it may take time for an individual to learn which expenses are important to him or her.

Widows and widowers spending declines, but very little, in the years following the spouse's death. This may reflect that the widow/widower stays in the couple's home and does not alter his or her routine expenses significantly.

The CAMS data are informative, but the results should be interpreted cautiously. The sample sizes are generally small and, although the same sample is maintained before and after the divorce, the waves surrounding them could be made up of very different types of households.

Poverty. The RAND Corporation produces a supplement to the HRS that calculates whether or not a household is in poverty based upon the definitions used by the U.S. Census Bureau. The two major components of a person's poverty status are income and household composition.[35] The supplement only calculates poverty status for interview waves from 2002-2010. In 2010, 4.9% of married women in their 50s, 22.3% of divorced women and 28.0% of widowed women were in poverty. Table 9.4 shows women's poverty rates before and after

[34] The CAMS survey is completed by a single household member for the household. There are not enough instances where both members are followed after divorce to provide a direct comparison of consumption of one two-person household versus two one-person households.

[35] Income includes before-tax income from: earnings, unemployment, workers' compensation, Social Security, SSI, public assistance, veterans benefits, pension and retirement income, interest, dividends, rents, royalties, income from estates and trusts, educational assistance, alimony, child support, assistance from outside the household, other sources, income of all resident family members. Income does not include: non-cash benefits (e.g., food stamps), capital gains and losses.

Table 9.4: Women: Poverty rates before and after separation/loss

Interview waves relative to event	Divorce In Poverty (%)	Divorce N	Widowhood In Poverty (%)	Widowhood N
-2	.042	96	.052	76
-1	.065	112	.035	93
0	.186	117	.176	99
1	.211	106	.081	115
2	.234	91	.101	124
3	.292	95	.079	116
4	.210	97	.109	150

Notes: Author's calculations, Data from RAND's supplement to Health and Retirement Study, Version M. Separations occurring between age 50 and 70. Separations occur between waves. Period 0 refers to the first interview after the separation/loss. Period -1 refers to the last interview before the separation/loss. Includes women whose separation/loss occurred between 50 and 70, where the separation/loss was observed within the HRS survey, and whom did not remarry.

the divorce or widowing event for women whose separation/loss occurred between 50 and 70, where the divorce/loss was observed within the HRS survey, and who did not remarry. For divorcing women, poverty rates jump from 6.5% to 18.6% after the separation. While the jump is not significant, the difference between two waves before divorce and the two waves after divorce is jointly significant at the 1% level. The poverty rates for divorcing women remain persistently high for the waves following the separation. Widows experience a statistically significant jump in the poverty rate at the time of the widowing event, but the increase in the poverty rate does not persist into future periods. Widowhood poverty rates are lower in table 9.4 because widows entering widowhood in this age range are going to be somewhat wealthier than the aggregate of widows entering in this age range and widows who were widowed at younger ages (i.e. what is captured in the cross-sectional statistic of 28.0% widowhood poverty in 2010).

9.4 Impact of divorce and widowhood on retirement benefits

When a household divorces, or one member dies, the change may affect its employer-provided retirement benefits or the household's Social Security benefit. Employees are generally eligible for two types of retirement benefits: defined benefit (DB) and defined contribution (DC). Some employees are eligible for a combination plan that consists of both. Appendix O provides a summary of the impact of marital separation/loss on different retirement benefit types. In this section, I describe the retirement benefits that are correlated with divorce or widowhood and may impact an individual's labor supply decision.

DB plans, typically referred to as pensions, are common in public sector jobs and provide the employee with a lifetime income. The employee who is eligible for the pension is considered a plan member, while his or her spouse would be considered a nonmember. Members and nonmembers are only entitled to a lifetime benefit in the form of an annuity if the member is vested, meaning they have worked for a minimum number of years at the employer. In the event of a divorce, the time rule formula is a common provision for separating assets. It can occur before or after the plan's member has first received benefits. Benefits to the nonmember spouse are calculated according to the following formula:

$$\text{Nonmember Benefit} = 50\% \times \frac{\begin{array}{c}\text{Total service credits earned}\\ \text{while married and employed}\end{array}}{\begin{array}{c}\text{Total service credits}\\ \text{at retirement}\end{array}} \times \text{Retirement Benefit}$$

The member's benefit is the total benefit less the nonmember's benefit. The time rule formula need not be the 50% specified above. The time rule provision is important for households with pensions, because this is where many of these households keep their retirement savings. Shorter marriages mean a smaller future income loss for the member, but a larger future income loss for the nonmember.

Pension funds do not have to immediately provide funds to nonmembers following a divorce. Funds will often require the member to first retire before the nonmember ex-spouse can begin drawing benefits. This means an ex-spouse's pension fund is not a guaranteed source of liquidity for the non-member after the divorce.

Social Security provides a lifetime benefit to spouses. While alive, a spouse is entitled to the better of her own benefit or 50% of her spouse's benefit. Social Security also provides a divorced spouse's benefit, that is nearly equivalent to the spouse's benefit. One important difference is that if an ex-spouse has been divorced less than 2 years, she must wait for the high earner to claim his benefit before she can collect the divorced spouse's benefit (in addition to be age qualifying herself). After two years, as long as she is age qualifying, she can claim the divorced spouse's benefit, regardless of if he has claimed his benefit.

Social Security also provides a lifetime benefit to survivors. A surviving spouse is entitled to the better of her own benefit or 100% of her spouse's benefit. For example, a single income household entitled to \$1200 per month before widowhood, is still entitled to \$1200 per month after widowhood. Alternatively, if a household is comprised of two equivalent earners, each entitled to \$600 per month (\$1200 per month total), then widowhood results in Social Security income falling to \$600 per month for the household after one person dies. We, therefore, would expect that households closer in lifetime earnings might be more affected by the loss of their spouse, because the combined household's income is not protected in the event of survivorship.

Issues related to DC plans, non-vested pensions, and survivor benefits from pensions, while important for the individual, are less important for our identification of the labor supply effects from divorce and widowhood in later sections. DC plans are generally split at divorce, at which point the nonmember's account is rolled over to an individual retirement account (i.e. IRA). In the case of widowhood, spouses are the default beneficiary. Pension plans, since the passage of Employee Retirement Income Security Act, are required to provide a survivorship option.

9.5 Summary

Divorced individuals, particularly women, experience significant asset and spending shocks around the time of their divorce. They become more likely to rent, and claim disability. It takes these individuals years to rebuild their assets, and for some, they expect to work longer.

The impact of separation near retirement is substantial, pushing many individuals, especially women, into poverty (Weir and Willis, 2000). While official poverty statistics are widely criticized (e.g. Citro and Michael, 1995), even alternative measures of well-being indicate significant poverty rates among widowed and divorced individuals relative to their married counterparts (Butrica, Murphy, and Zedlewski, 2007). While widowhood poverty may have been a large issue in previous decades, divorcée poverty has the potential to become the relatively larger social issue.

The numbers above suggest that near-retirement divorce is becoming more common and that it has major financial consequences. Little is known, however, about how women and men respond to the financial consequences of divorce near retirement. Johnson and Skinner (1986) find that divorce leads women to increase labor force participation by 20 percentage points, but, their sample contains much younger women, with an average age of 32.8. Compared to the women in their study, women divorcing near retirement are a very different population, both because their marriage are longer, 20.6 years on average, and because they may face greater difficulty returning to the labor force in old-age. In the next section, for a household near retirement, I extend a household life-cycle model of consumption and labor supply to incorporate the impact of household separation on labor force participation. The model does not include a disability application decision or a choice for renting versus owning, but it is able to capture an individual's saving, spending, work and retirement decisions. Therefore, it is able to capture most of the important short and long-term decisions an individual must make following a separation or loss.

10 Model

In this section I present a household model of labor supply, and describe how I theoretically account for the transition from a two-person household to a one-person household. This is important for understanding my choice of empirical specification in the next section.

I model the household decision process prior to divorce as a finite-lived household that maximizes joint utility from consumption and each individual's leisure. Divorce is assumed to be a sudden and unexpected change from a household of two individuals to two households each with a single individual.[36]

I present a unitary model of the household that places fixed weights on each individual's leisure, which differs from more recent economic models that incorporate divorce. Most economic models of divorce focus on household formation and its interaction with labor supply decisions (e.g. Fernandez and Wong, 2013). In this case, a collective model is useful because it reflects the reallocation of decision weighting within a marriage if conditions change, such as income potential. Older households, which have been married for a long time (i.e. 28.2 years on average), have established within-household responsibilities. In most cases, these couples have raised children together, bought a home, and survived economic ups and downs. By choosing a unitary model, I assume that the within-household weighting is fixed while the household remains married.[37]

Moreover, the model presented in the next section assumes intrahousehold separable utility. Note that this is a strong assumption, abstracting away from leisure complementarities at the household level. In examining variations to the empirical specification, I will relax this assumption.

10.1 Married Household Preferences

Married households maximize the present value of their discounted lifetime utility by choosing their consumption, labor participation and whether or not to claim benefits. The house-

[36]This could be modeled as an increased risk of entering the one person state. This has been done in other work with younger individuals. Since I am looking at an older sample where marriages are in place longer, I model it as a surprise.

[37]The model could be written as a collective model, as in Browning, Chiappori, and Lewbel (2013). In their paper, Browning, Chiappori, and Lewbel (2013) assume that the household is pareto efficient, implying that by the second welfare theorem, that individuals provide intra-household transfers to maximize utility. To achieve this, the weight the household places on each members utility must depend on a sharing rule that would be function of prices, income, and distribution factors (i.e. factors with no direct impact on preferences, technology, or the budget constraint, but may influence the decision process).

hold's instantaneous utility function is given by:

$$u_t\left(C_t, L_{H,t}, L_{W,t}\right) = \frac{C_t^{1-\alpha}}{1-\alpha} + \frac{D_{H,t} L_{H,t}^{1-\gamma_H}}{1-\gamma_H} + \frac{D_{W,t} L_{W,t}^{1-\gamma_W}}{1-\gamma_W}, \qquad \alpha > 0 \tag{10.1}$$

where the parameter α captures the household's constant relative risk aversion, and $D_{m,t}$ represents household member m's relative preference between household consumption and his or her own leisure in year t. All consumption is modeled as shared within the household to match the fact that most household surveys, included the Health and Retirement Study used in this paper, do not track individual consumption.

Each household member's leisure, $L_{m,t}$, is defined as:

$$L_{m,t} = L - N_{m,t} \tag{10.2}$$

where $N_{m,t}$ represents the household member's labor supply and L represents the constant endowment of leisure in each period.

Each household member's relative preference between household consumption and his or her own leisure, $D_{m,t}$. It depends on time varying state variables such the male or female's age and health. In the case of the man ($m = H$), it takes the form:

$$D_{H,t} = exp\left(\beta_H + \beta_{H,age} age_{H,t} + \beta_{H,health} health_{H,t} + \varepsilon_H\right) \tag{10.3}$$

Analogously, the woman's marginal rate of substitution ($m = W$) is determined by:

$$D_{W,t} = exp\left(\beta_W + \beta_{W,age} age_{W,t} + \beta_{W,health} health_{W,t} + \varepsilon_W\right) \tag{10.4}$$

The constant term, β_m, represents a scaling factor for the impact of an individual's leisure on household utility. I assume when the household is created its members negotiate the weight of each individual's contribution to total utility from his or her leisure, and it is considered fixed for the duration of the household.

Each household member faces uncertainty over their continued survival and health. The survival probability and the health transition probabilities depend on age, sex, and previous health status.

Lastly, ε_m represents an individual fixed effect for higher value of retirement to an individual as in Gustman and Steinmeier (2004b). It is through ε_m that a couple's preference for consumption versus leisure are allowed to differ across individuals and households in ways that are not otherwise captured by the economic model. For example, if a woman loves work, this would mean she would have a low draw for ε_W.

10.2 Single Household Preferences

A household comprised of a single member maximizes its present value of discounted lifetime utility by choosing consumption, labor participation and whether or not to claim benefits. The single household's instantaneous utility function is given by:

$$u_t\left(C_t, L_{m,t}\right) = \frac{\left(s \cdot C_t\right)^{1-\alpha}}{1-\alpha} + \frac{\tilde{D}_{m,t} L_{m,t}^{1-\gamma_m}}{1-\gamma_m} \qquad \alpha < 1 \tag{10.5}$$

where the parameter s captures the economies of scale associated with a two person household. If $s = 1$, then all goods are public (perfect economies of scale), but if some goods are private, then $s > 1$. With the exception of s and $\tilde{D}_{m,t}$, the other terms are assumed to be invariant to the composition of the household.

Similar to the married household, each individual's relative preference between household consumption and his or her own leisure, $\tilde{D}_{m,t}$:

$$\tilde{D}_{m,t} = exp\left(\tilde{\beta}_m + \beta_{m,age} age_{m,t} + \beta_{m,health} health_{m,t} + \tilde{\varepsilon}_m\right)$$

The notation is chosen purposefully: I assume marriage status does not impact age and health effects, but does result in a re-weighting of leisure versus consumption $\left(\tilde{\beta}_m\right)$, and that the person's unobserved preference for leisure also changes ($\tilde{\varepsilon}_m$). This allows the model to reflect that the tradeoff between consumption and leisure changes discontinuously at divorce.

Each single household member faces uncertainty over their continued survival and health. The survival probability and the health transition probabilities depend on age, sex, and previous health status.

10.3 Budget Set

The household is able to accumulate assets, A_t, over its lifetime subject to the following equation:

$$A_{t+1} = A_t + Y_t - C_t - M_t \tag{10.6}$$

where C_t is per period household consumption, Y_t is per period income, and M_t is per period medical expenses.[38]

[38] Log medical expenses are based in part on an individual's health, and are assumed to be normally distributed.

A household's per period income can come from a number of sources: household member's wage income, $\omega_m(N_{m,t})$, interest income, rA_t, defined benefit and contribution pension income, $db_{m,t}$ and $dc_{m,t}$ respectively, alimony for divorced households, $alim_{m,t}$, and a household social security benefit, ssb_t, where all of these sources of income are subject to tax, tx.

$$Y_t = Y\left(rA_t + ssb_t + \sum_{m\in\{M,W\}} \left(\omega_m(N_{m,t}) + alim_{m,t} + db_{m,t} + r\cdot dc_{m,t}\right), tx\right) \tag{10.7}$$

Note that I assume the only additional income generated by work is the wage. Alternatively, as in Knapp (2013), one could imagine additional work impacting retirement benefits, but for simplicity that will not be considered in this paper.

10.4 Individual Labor Supply

In this section, I derive the individual labor supply curves for married and single households. I then consider the immediate impact of the divorce transition on the individual's labor supply as well as the longer term impacts on retirement and benefit claiming.

10.4.1 Before Separation

A married individual will optimally choose to work if her marginal utility from leisure is less than or equal to her return from working (sum of marginal wage and retirement benefits). First order conditions with respect to labor supply from the household problem imply:

$$e^{\left(\beta_m+\beta_{m,age}age_{m,t}+\beta_{m,health}health_{m,t}+\varepsilon_m\right)}\left(L-N_{m,t}\right)^{-\gamma_m} \quad\leq\quad \lambda_t\cdot\frac{\partial Y}{\partial N_{m,t}} \tag{10.8}$$

where λ_t is the Lagrange multiplier and represents the household's marginal utility of wealth. Assuming constraint (10.8) binds, it can be rewritten for a household member's leisure in period t, $(L-N_{m,t})$, as :

$$\begin{aligned}\ln\left(L-N_{m,t}\right) &= \frac{1}{\gamma_m}\left\{\beta_m+\beta_{m,age}age_{m,t}+\beta_{m,health}health_{m,t}-\ln\left(\frac{\partial Y}{\partial N_{m,t}}\right)-\ln\left(\lambda_t\right)+\varepsilon_m\right\}\\ &= \psi_m\beta_m+\psi_m\beta_{m,age}age_{m,t}+\psi_m\beta_{m,health}health_{m,t}-\psi_m\ln\left(\frac{\partial Y}{\partial N_{m,t}}\right)\\ &\quad -\psi_m\ln\lambda_t+\psi_m\varepsilon_m\end{aligned} \tag{10.9}$$

where $\psi_m = \frac{1}{\gamma_m}$. Note that the individual's labor supply is an affine mapping of the individual's leisure demand in equation (10.9), and the labor supply preserves the expected relationships including: (1) higher wages lead to less leisure, (2) if $\beta_{m,age} > 0$ then older individuals are less inclined to work, and (3) an additional dollar of assets will lead to more leisure.

The spouse's labor supply does not directly impact the individual's labor supply because of the strong separability assumption in §10.1. Instead, it only impacts the household member's decision via lifetime wealth, and hence effects the household's marginal utility of wealth, λ_t.

10.4.2 After Separation

A single household will optimally choose to work if her marginal utility from leisure is less than or equal to her return from working (sum of marginal wage and retirement benefits). First order conditions with respect to labor supply imply:

$$e^{\left(\tilde{\beta}_m+\beta_{m,age}age_{m,t}+\beta_{m,health}health_{m,t}+\tilde{\varepsilon}_m\right)}\left(L-N_{m,t}\right)^{-\gamma_m} \quad \leq \quad \frac{\tilde{\lambda}_t}{s}\cdot\frac{\partial Y}{\partial N_{m,t}} \tag{10.10}$$

where $\tilde{\lambda}_t$ is the Lagrange multiplier and represents the marginal utility of an additional dollar of assets when single. Assuming constraint (10.10) binds, it can be rewritten for an individual's leisure in period t, $(L - N_{m,t})$, as :

$$\begin{aligned}
\ln\left(L-N_{m,t}\right) &= \frac{1}{\gamma_m}\left\{\tilde{\beta}_m+\beta_{m,age}age_{m,t}+\beta_{m,health}health_{m,t}+\tilde{\varepsilon}_m-\ln\tilde{\lambda}_t-\ln\frac{\partial Y}{\partial N_{m,t}}\right\} \\
&= \psi_m\tilde{\beta}_m+\psi_m\beta_{m,age}age_{m,t}+\psi_m\beta_{m,health}health_{m,t}-\psi_m\ln\left(\frac{\partial Y}{\partial N_{m,t}}\right) \\
&\quad -\psi_m\ln\tilde{\lambda}_t+\psi_m\ln s+\psi_m\tilde{\varepsilon}_m
\end{aligned} \tag{10.11}$$

Similar to the before separation labor supply, the single individual's labor supply preserves the same expected relationships. What separates the single household's labor supply from the married household's labor supply is that the single household may tradeoff consumption and leisure differently based on $\tilde{\beta}_m$, the man or woman's weighting of consumption versus leisure following the separation; $\tilde{\varepsilon}_m$, an idiosyncratic shock to preference for leisure that varies by gender of the household member, m, but not time, t; s, household economies of scale; and $\tilde{\lambda}_t$, the household's marginal utility of wealth.

10.5 How separations change Labor Supply

In this section, I clarify what in the model drives labor force changes after separations. First I consider changes in the budget constraint due to a separation, and second I consider changes to preferences.

The budget constraint of a newly single household is primarily altered by: (i) an immediate change in assets, and (ii) elimination of current and future earnings from the previous household member.[39] In the context of a model with perfect capital markets, where individuals could freely borrow from their ex-spouses future income, then the two effects would be equivalent. In my dynamic model, assets are liquid wealth, and future expected spousal earnings are illiquid wealth. Illiquid wealth cannot immediately be disbursed in divorce, although it can be over time with alimony. Therefore if a single household has very little liquid assets after the separation, the separation itself may spur a liquidity crisis (i.e. the individual can't afford to consume their previous level of consumption), in turn forcing the individual to increase work or take government income assistance.

The preferences of a newly single household is primarily altered by: (iii) household economies of scale, s, and (iv) changes in preference for consumption relative to own-leisure after the separation, $\left(\tilde{\beta}_m - \beta_m\right) + (\tilde{\varepsilon}_m - \varepsilon_m)$. If the economies of scale from marriage are large (i.e. s is close to 1) then the utility equivalent amount of consumption for a given level of leisure that a household needs with two people is equal to the amount it needs as a one person household. This would imply that simply living as a one person household would entail little decrease in consumption costs, forcing an individual to work more in order to finance consumption. Alternatively, the newly single household might trade off consumption for leisure differently than they would have during marriage. Depending on the direction, the change in relative preference for leisure could increase or decrease the magnitude of the change preference for consumption.

[39] Assets do not necessarily have to decrease. Assets could increase if life insurance payments exceed funeral and other end-of-life expenses.

11 Estimation

This section translates the model (§10) into a tractable fixed-effects framework that can be used to understand how a separating individual alters his or her labor supply in response to the separation. I then conduct robustness checks to examine how sensitive the point estimates are to alternative specifications.

11.1 Empirical Specification

In this section I present the specification to be estimated based on the theoretical model in §10. I use the leisure demand functions of married and single individuals, based on equation (10.9) before divorce and equation (10.11) after divorce.[40] I suppress the gender of the household member, m, to simplify the notation.

As discussed in §10.5, a divorce or death's impact on the marginal utility of wealth can persist after the separation takes place and, to the degree that it is anticipated, may have an effect prior to the separation. This effect cannot immediately be separated from changes in the household's preference for leisure that might result from the household transitioning from two persons to one person. I follow a non-parametric approach by including a dummy variable, $SEP_{i,t}^{k}$, that can take a value of one if the separation leading to divorce takes place k periods before current year t (k may also take on negative values if year t is before the separation). Similarly if the household is widowed, $WID_{i,t}^{k}$. This approach is used in case the household cannot respond immediately to the separation or if it is anticipated. For example, if a divorce is anticipated, we might expect to see an increase in labor supply prior to the household separation. Additionally, if the separated household is unable to immediately find

[40]The first order condition with respect to assets of the household problem in §10, yields:

$$\lambda_t = \delta(1+r_t)\mathbb{E}_t\left[\lambda_{t+1}\right]$$

where δ is the household's discount rate, and r_t is the known rate of return on assets in period t. Uncertainty over the future marginal utility of wealth arises from uncertain changes in survival, health, and medical expenses. If, as in Blundell and MaCurdy (1999), I assume that $ln\lambda_t = \mathbb{E}_{t-1}\left[ln\lambda_t\right] + \eta_t$, then I can characterize the stochastic process of the Euler equation by $ln\lambda_t = ln\lambda_{t-1} - \left(ln\left[\delta(1+r_t)\right] + ln\left(\mathbb{E}_{t-1}\left[e^{\eta_t}\right]\right)\right) + \eta_t$. With repeated substitution, this relationship yields:

$$ln\lambda_t = \sum_{j=1}^{t} b_j + ln\lambda_0 + \sum_{j=1}^{t} \eta_j$$

where $b_t = -\left(ln\left[\delta(1+r_t)\right] + ln\left(\mathbb{E}_{t-1}\left[e^{\eta_t}\right]\right)\right)$ and η_t is independently distributed $N(0, \sigma_\epsilon)$. This equation suggests that the joint household sets λ_0 in the first period and updates it as uncertain and unforeseeable events are realized. Specific to the discussion here, I am interested in the adjustment following divorce or widowhood. Note also, that the first term, $\sum_{j=1}^{t} b_j$, implies a period fixed effect, while $ln\lambda_0$ implies a time invariant fixed effect for the household.

a job, or some other factors hinder an immediate change in labor supply, then the newly single household's response may be delayed.

The empirical specification is:

$$
\begin{aligned}
\ln(L - N_{i,t}) \quad = \quad & \phi_{age} age_{i,t} + \phi_{health} health_{i,t} - \psi ln\left(\frac{\partial Y}{\partial N_{i,t}}\right) \\
& + \sum_{k=-3}^{7} \tilde{\zeta}_k^{SEP} SEP_{i,t}^k + \sum_{k=-3}^{7} \tilde{\zeta}_k^{WID} WID_{i,t}^k \\
& + \rho_i + \sum_{j=3}^{10} b_j + \xi_{i,t}
\end{aligned}
\tag{11.1}
$$

where the change in income based on work, $ln\left(\frac{\partial Y}{\partial N_{i,t}}\right)$, is controlled by a quadratic in years of actual experience. Since Social Security and other pension incentives significantly alter the decision to retire at specific ages, I take a non-parametric approach and use a dummy variable for each age. The sixth term in (11.1) corresponds to individual fixed effects, and the seventh term is represented by survey wave dummies (for a total of 10 survey waves). The last term represent the transitory shock to preferences for leisure.[41]

I examine the impact of divorce on both labor force participation and hours worked. When examining participation, the dependent variable in (11.1) is replaced with a binary outcome for whether an individual is working for pay. Alternatively, when I examine hours worked, I set $L = 365 \times 24 = 8760$.

11.2 Initial Regression Estimates

I use a within or fixed effects estimator to estimate (11.1) because it exploits the advantages of the panel data set by using time variation within individuals to consistently estimate time-varying coefficients. The model is estimated separately for men and women.

Table 11.1 reports the estimates for equation (11.1) for women and men. These tables restrict the sample to only those individuals who (i) have been married during the HRS observation period, 1992-2010, (ii) never applied for disability, and (iii) are part of the at-risk population. It also does not include observations if an individual remarries following the initial separation. I observe that poor health leads to lower participation, and greater work

[41]My results will be able to examine the time-varying effects of the household separation on changes to the total fixed effect, $\rho_i = \beta_m + \varepsilon_m + \ln\lambda_0$. The subscript i on ρ_i is included to emphasize that it is an individual fixed effect since it relies on an individual's own realization of ε_i and the individual's initial marginal utility of wealth. Additionally $\xi_{i,t} = \sum_{j=1}^{t} \eta_j + \iota_{i,t}$ where η_t represents transitory shocks to the marginal utility of wealth and $\iota_{i,t}$ represents transitory shocks to leisure.

experience increases the likelihood of work for women and men, but has diminishing returns.

The results from the dummy coefficient estimates on divorce and widowhood indicate that asset loss is the primary driver of women's return to the labor force after a separation. Relative to the period before separation, women are more likely to participate in the labor force, jumping an average of 100 hours in the first wave after the separation (wave 0) and increasing to 500 hours by the 2nd full wave (4-6 years) after the separation (see figure 11.1). Relative to the wave before the divorce, the probability of participating in the labor force increases by 9.1% following the divorce and increases to 17.1% in the 3rd full wave following the separation. After the 3rd full wave, the effect of separation falls and becomes less precise.

Widows, who do not lose their assets on average, show no indication of returning to the labor force (see figure P.5 in appendix P). We observe that women's labor force participation is insignificantly higher in the waves before the husband's death. This could indicate that the wife exits her job to take on the role of caretaker. Following the husband's death, there is little evidence of the widow returning to work. Household separation, either through widowhood or divorce, causes a loss of future income flow from the ex-spouse, eliminates economies of scale, and changes the household's preferences. The key difference, however, is that asset loss mainly occurs in divorce. Since we only observe a major labor supply response following a divorce, this suggests asset loss is the main driver of women's return to the workforce.

Alternatively, consider the case of men. I observe that divorce has no significant impact on labor force participation or hours worked for men (See figure P.6 in appendix P). It does, however, appear that widowing leads to a reduction in hours worked for men (see figure 11.2), which could support either an asset loss/economies of scale argument or an end to household bargaining argument. For example, if the household has one less mouth to feed, and the widower was the primary earner, then the loss of the spouse might lead the man to work less because the lifetime assets required for the household's retirement plan are lower. Alternatively, if the man was only working to finance the wife's high demand for leisure, than the death of his wife would permit him to work less and consume more leisure himself.

In summary, I find that women respond significantly to a separation leading to a divorce, by increasing labor supply by up to 500 annual labor hours, or increasing labor force participation by 17.1 percentage points. Little effect is found for divorcing men or widowed women. Widowed men decrease their labor force participation prior to the widowing event, and this decreased level continues after the event.

These results are consistent with a household life-cycle model where the wife is less attached to the labor force and the husband works full-time. A divorcing husband does not respond immediately. If he is already employed full-time, time constraints may limit his

Table 11.1: Estimates of Leisure Demand and Labor Force Participation

		Women				Men			
		Leisure Demand		Labor Force Participation		Leisure Demand		Labor Force Participation	
		Divorces	Widows	Divorces	Widows	Divorces	Widows	Divorces	Widows
		(1w)	(2w)	(3w)	(4w)	(1m)	(2m)	(3m)	(4m)
After Separation (wave)	0th	-0.00916	-0.00914	0.0912**	0.0289	0.00851	0.0247	0.00336	-0.0123
		(0.0119)	(0.0110)	(0.0415)	(0.0366)	(0.0143)	(0.0241)	(0.0398)	(0.0639)
	1st	-0.0339***	-0.0112	0.122***	0.0252	0.00791	0.0276	-0.0122	-0.0317
		(0.0130)	(0.0121)	(0.0425)	(0.0409)	(0.0167)	(0.0253)	(0.0499)	(0.0775)
	2nd	-0.0650***	-0.0232*	0.154***	0.0589	-0.0102	0.0183	0.00284	-0.0212
		(0.0231)	(0.0125)	(0.0482)	(0.0467)	(0.0197)	(0.0244)	(0.0560)	(0.0833)
	3rd	-0.0368	-0.0128	0.171***	0.0828*	-0.0162	0.0330	-0.00353	-0.0628
		(0.0228)	(0.0119)	(0.0571)	(0.0478)	(0.0272)	(0.0314)	(0.0636)	(0.109)
	4th	-0.00507	-0.0135	0.0379	0.0607	0.00132	0.0285	0.00973	-0.0462
		(0.0199)	(0.0127)	(0.0636)	(0.0523)	(0.0270)	(0.0326)	(0.0747)	(0.113)
	5th	-0.0131	-0.0135	0.0706	0.0325	-0.00129	0.0174	-0.0404	0.0850
		(0.0191)	(0.0130)	(0.0741)	(0.0553)	(0.0230)	(0.0342)	(0.0773)	(0.117)
	6th	0.0227	-0.0244	-0.0541	0.0372	0.0185	0.0354	-0.0305	0.0335
		(0.0211)	(0.0154)	(0.0769)	(0.0570)	(0.0304)	(0.0358)	(0.0861)	(0.126)
	7th	0.0143	-0.0503**	0.0487	0.0575	0.0192	-0.00503	-0.0159	0.0547
		(0.0239)	(0.0204)	(0.0867)	(0.0646)	(0.0415)	(0.0321)	(0.0944)	(0.138)
Before Separation (wave)	1st	(control)	(control)	(control)	(control)	(control)	(control)	(control)	(control)
	2nd	0.00554	-0.00454	0.0108	-0.0145	-0.0183	-0.0295	0.0385	-0.00692
		(0.0124)	(0.00913)	(0.0397)	(0.0387)	(0.0187)	(0.0305)	(0.0483)	(0.0408)
	3rd	0.00291	-0.0256*	0.00909	0.0626	-0.0250	-0.0384	0.0285	0.0903
		(0.0191)	(0.0135)	(0.0630)	(0.0537)	(0.0276)	(0.0319)	(0.0668)	(0.102)
	4th	-0.0158	-0.0259	-0.0249	0.0770	0.0259	-0.0452	-0.0816	0.336***
		(0.0238)	(0.0210)	(0.0746)	(0.0769)	(0.0210)	(0.0362)	(0.0717)	(0.0993)
	5th	0.0307	-0.0196	-0.0336	0.150*	0.00118	-0.0300	0.0340	0.226*
		(0.0276)	(0.0148)	(0.0940)	(0.0793)	(0.0296)	(0.0309)	(0.0706)	(0.117)

	Women				Men			
	Leisure Demand		Labor Force Participation		Leisure Demand		Labor Force Participation	
	Divorces	Widows	Divorces	Widows	Divorces	Widows	Divorces	Widows
	(1w)	(2w)	(3w)	(4w)	(1m)	(2m)	(3m)	(4m)
Health (lagged)	0.00265*		-0.0161***		0.00661***		-0.0279***	
	(0.00151)		(0.00520)		(0.00197)		(0.00578)	
Actual Experience	-0.00271***		0.0164***		-0.00650***		0.0493***	
	(0.00104)		(0.00330)		(0.00214)		(0.00556)	
Actual Experience^2	8.69e-05***		-0.000232***		7.61e-05***		-0.000314***	
	(1.55e-05)		(4.50e-05)		(2.20e-05)		(5.67e-05)	
Observations	50,376		50,806		43,263		43,693	
R^2 (without F.E.)	0.137		0.150		0.241		0.227	
R^2 (with F.E.)	0.660		0.703		0.695		0.697	

Notes: Person-level fixed effect regressions. The 0th wave is the wave interview wave immediately after separation, and may not be a full two years (the time period of a normal interview wave). All regression use dummies for each interview wave and age between 51 and 75. The first and third columns report coefficients where the base year is the year prior to separation leading to divorce. Alternatively, the second and fourth columns report the same models except that the base year is the before widowhood. Standard errors, clustered at the person level, are reported in parentheses. Results for the remaining regressors are available from the author. *** p<0.01, ** p<0.05, * p<0.1.

ability to respond by working more hours. The wife, if she is not working full-time, can increase her labor hours or return to the labor force. In widowed households, the survivor has reduced long-term costs implying that accumulated savings, in combination with life insurance or old age pension income, may be enough to permit an early retirement.

This first specification does not allow for the separation of the different factors that might be driving the result. Is it the fact the wife loses assets in divorce that causes her to work longer? Or, is it driven by new found freedom to make decisions independent of pre-ordained household roles? In order to answer this question, I relax the specification used in equation (11.1) to separate out how changes in the marginal utility of wealth are amplified by the relative changes in a household's finances.

11.3 Accounting for Financial Losses

As mentioned above, the initial specification does not permit the separate identification of the labor supply effects of (i) asset loss, (ii) future spousal income loss, (iii) changes to

Figure 11.1: Effect of Divorce timing on Women's Labor Force Participation

(a) Labor Force Participation, Ages 50-70

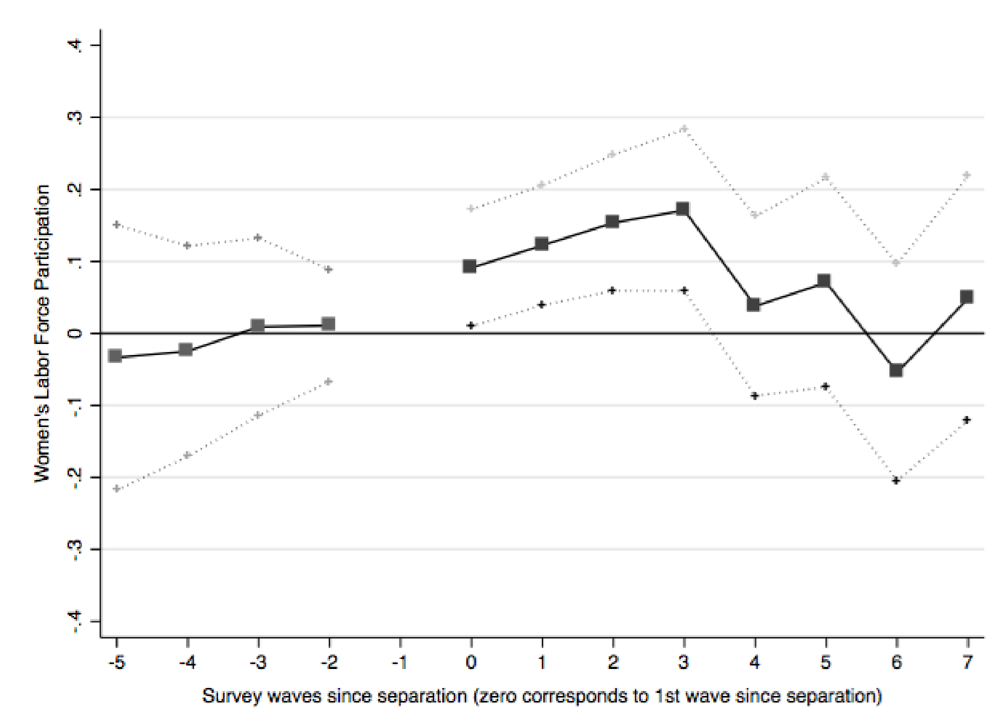

(b) Hours Worked, Ages 50-70

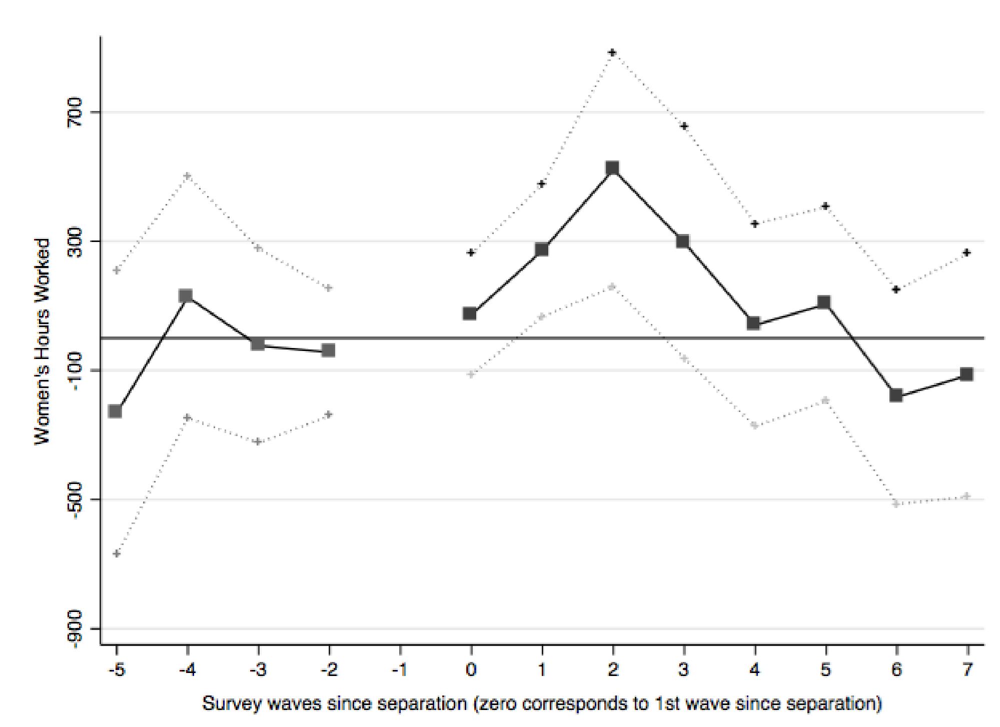

Notes: Dashed lines represent 95% confidence intervals. See table 11.1 for additional notes.

Figure 11.2: Effect of Widowhood timing on Men's Labor Force Participation

(a) Labor Force Participation, Ages 50-70

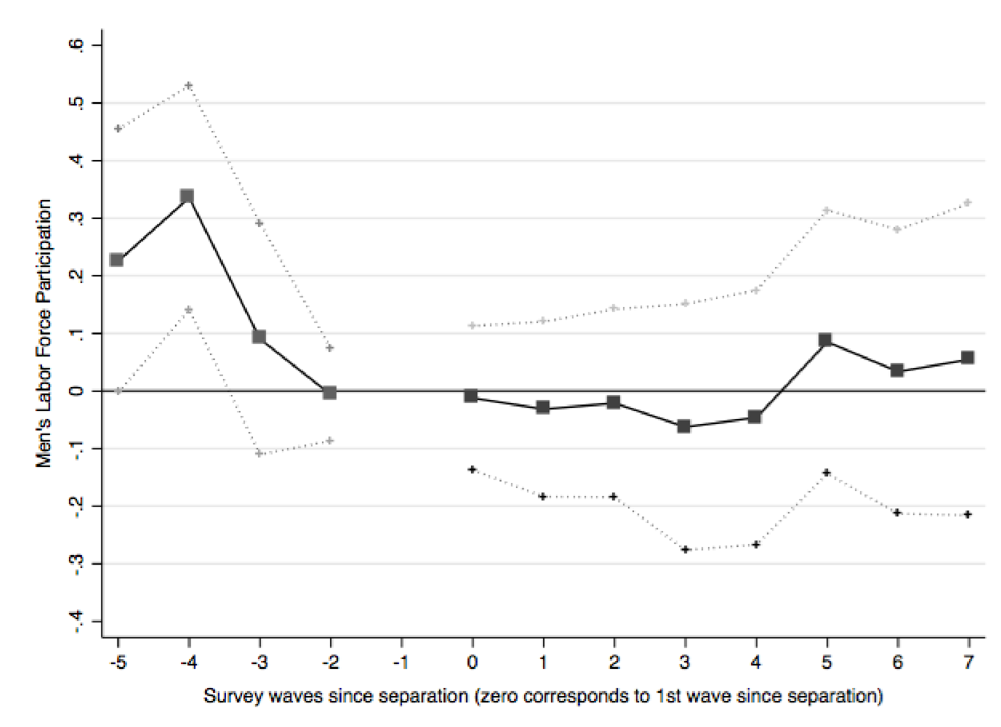

(b) Hours Worked, Ages 50-70

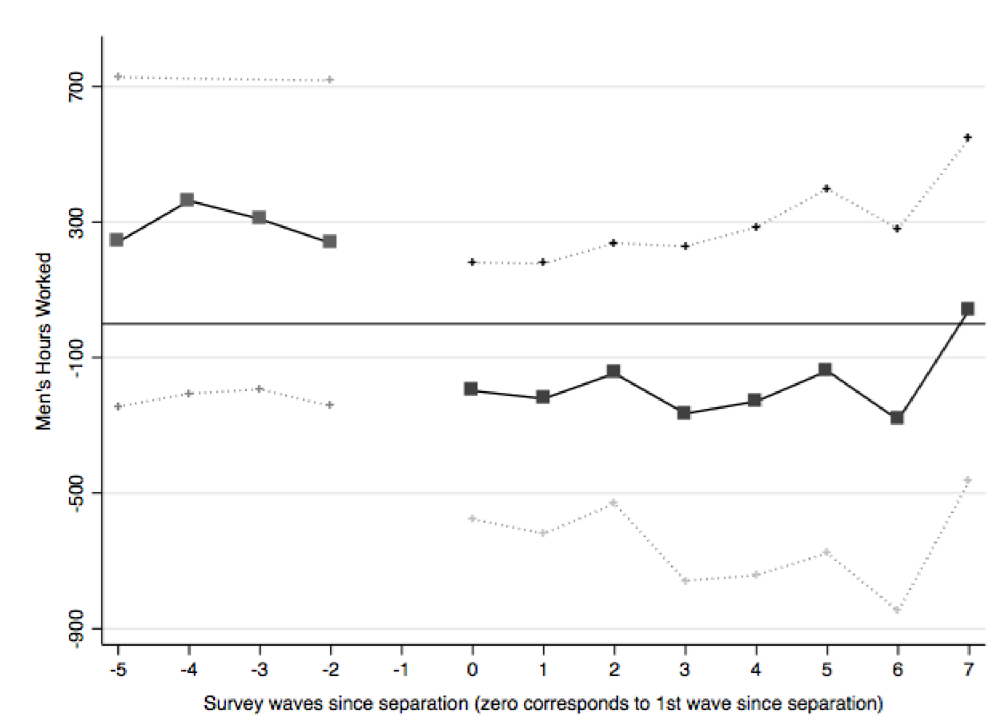

Notes: Dashed lines represent 95% confidence intervals. See table 11.1 for additional notes.

preferences, and (iv) changes to economies of scale. In the data, households vary in both whether the separation leads to large asset loss and in the pre-separation income levels of the spouse. The post-divorce asset change and the pre-divorce spousal income level can be included in the empirical specification to capture the contribution of each of these effects. If the household faces a liquidity crisis after a separation, then the household would be forced to return to the labor force. If asset loss is what is driving the majority of labor supply change from the last subsection, then the remaining effects should be rendered insignificant. A similar analysis would hold if the household was counting on the spouse's income as a primary means of survival.

Building on the specification in equation (11.1), I interact the divorce and widowhood dummy variables with asset loss (e.g. the household asset levels after the separation less the asset levels before the separation), and lost potential spouse income (e.g. the amount of future income the household would have received if the ex-spouse continued working in their baseline occupation until age 65, conditional on mortality). The divorce and widowhood dummy variables are retained, and capture the residual effect that cannot be described by asset loss or lost future spousal income. This residual effect will capture, as described in §10.5, changes in preferences and economies of scale resulting from these events.

Tables 11.2 and 11.3 present the results for wives and husbands of this new specification. For divorced women, a loss of \$100,000 in assets leads to an increase in the labor supply of 2.27-4.13% for the decade following divorce. To put this number in perspective, recall that 62% of divorced women worked before the separation. If I use the point estimate of 0.0405 in the second full wave since the separation, these results would then suggest that the average non-working women would be $(0.0405) / (1 - 0.62) = 0.107$ or 10.7% more likely to work per \$100,000 lost during the divorce. The effect of asset loss is persistent, raising labor force participation for more than a decade by 6.0-10.9% per \$100,000 lost during the divorce. The fact that only asset loss is a significant driver of an immediate change in labor force participation may indicate that divorce's wealth shock depletes buffer stock savings, forcing some women to return to work.

There is no significant effect of loss in future spousal income, however the sign is generally negative, which is what we would expect if higher spousal income meant a greater likelihood for alimony. If liquid assets are important as a buffer against immediate expenses after divorce, then in present value, \$1 lost from assets should be different from \$1 lost from future spousal earnings. A simple Wald test rejects the hypothesis that they are equivalent at the 10% level.[42] I interpret this as generally supportive of the importance of asset loss in

[42] If expected future spousal earnings are corrected for mortality probability, the simple Wald test would reject the hypothesis at the 5% level.

determining a woman's decision to return to the labor force after a divorce.

For women, the residual effect of the the time after divorce still indicates a large effect of the separation on labor supply, with the increase approaching 11.4% in the 2nd full wave following divorce. These effects are significant only at the 10% level. These residual, post-separation effects, which are uncorrelated with asset loss and the ex-spouse's income, reflect changes in the household's preferences. These results suggest that either the loss of economies of scale or leisure complementarities may impact the return to the labor force.

In the last subsection, I found that widowed women's leisure demand and labor force participation decreased in anticipation of the husband's death. When I account for asset loss and lost spousal income, you might expect that lost spousal income could be the major factor in the woman's decision to return to the workforce, since asset loss is rare in the widowing event (unless significant sums are spent on medical care leading up to the husband's death). The estimates in models 1w and 2w of table 11.2 suggest that, if this is the case, the effect is very small. Counterintuitively, the marginal effect of asset loss indicates that women are (insignificantly) more likely to leave the labor force in the period following the husband's death. Since asset loss during widowhood could be positive or negative and is generally small, it is not surprising that no effect is found.

Men's leisure demand and labor force participation are largely unaffected by a divorce, even when controlling for asset loss and spousal income loss. This largely reflects that most men are working in their 50s or are not dependent on their wife's income. Widowers, like widows, exit the labor force prior to their wife's death. The significance of these factors is reduced by the inclusion of variables capturing asset loss and spousal income loss. The estimates in model 2w of table 11.3 suggests that men are much more likely to exit the labor force if asset loss is significant. This could be the case if medical expenses are very high, and the wife's persistent disease alters the husband's long-run preference for work. Since men are largely working full-time, their response is limited. Even when accounting for other factors, divorce or death seem to have little immediate effect on men's hours worked or labor force participation. What these results do not comment on is the duration of employment. These men might choose to retire sooner if the change in marital status has reduced their marginal utility of wealth.

11.4 Robustness Checks

11.4.1 Leisure Complementarities

The specification estimated in §11.2 is based on a model where the spouse's labor supply only effects an individual through the marginal utility of wealth. This is equivalent to arguing

Table 11.2: Estimates of Leisure Demand and Labor Force Participation of Women accounting for asset loss and loss of future spousal income following separation.

		Leisure Demand (1)						Labor Force Participation (2)					
		Divorced (1d)			Widowed (1w)			Divorced (2d)			Widowed (2w)		
		Dummy	× Asset Loss	× Sp. Loss	Dummy	× Asset Loss	× Sp. Loss	Dummy	× Asset Loss	× Sp. Loss	Dummy	× Asset Loss	× Sp. Loss
After Separation (wave)	0th	0.00270	-0.00923***	0.000424	-0.00514	-5.36e-05	-0.000598*	0.0802	0.0227**	-0.00175	0.0240	-0.0103	6.42e-05
		(0.0128)	(0.00300)	(0.000414)	(0.0118)	(0.00218)	(0.000341)	(0.0532)	(0.0105)	(0.00126)	(0.0405)	(0.0115)	(0.000465)
	1st	-0.0172	-0.0130***	0.000656	-0.0101	0.00323	-0.000614	0.0634	0.0272**	-0.000775	0.0180	-0.0224*	0.000818
		(0.0142)	(0.00355)	(0.000486)	(0.0137)	(0.00237)	(0.000436)	(0.0585)	(0.0122)	(0.00143)	(0.0453)	(0.0125)	(0.00125)
	2nd	-0.0338	-0.0288*	0.00127	-0.0208	0.00822***	0.000359	0.114*	0.0405***	-0.00273	0.0426	-0.0376**	-0.000623
		(0.0300)	(0.0159)	(0.000918)	(0.0141)	(0.00304)	(0.000273)	(0.0595)	(0.0126)	(0.00179)	(0.0515)	(0.0160)	(0.00119)
	3rd	-0.0363	-0.00394	0.00131*	-0.00957	0.00533**	-0.000247	0.0879	0.0283*	-0.00114	0.0641	-0.0239*	0.000452
		(0.0315)	(0.00460)	(0.000691)	(0.0132)	(0.00248)	(0.000162)	(0.0709)	(0.0164)	(0.00215)	(0.0534)	(0.0134)	(0.000422)
	4th	0.0188	-0.00433*	-0.000963	-0.0112	0.00697**	0.000364**	-0.0517	0.0413**	0.000899	0.0380	-0.0223	0.00120***
		(0.0218)	(0.00247)	(0.000686)	(0.0145)	(0.00319)	(0.000168)	(0.0741)	(0.0171)	(0.00298)	(0.0594)	(0.0144)	(0.000411)
	5th	0.0109	-0.00567*	-0.000808	-0.0139	0.00580**	0.000503***	-0.00796	0.0280	0.00113	0.0384	-0.0293**	-0.00367***
		(0.0221)	(0.00301)	(0.000884)	(0.0150)	(0.00254)	(0.000185)	(0.0831)	(0.0223)	(0.00327)	(0.0620)	(0.0128)	(0.000492)
	6th	0.0476**	-0.0196**	0.00531**	-0.0239	0.00331	-0.0308***	-0.0787	0.0429	-0.0193***	0.0271	-0.00729	0.0306**
		(0.0232)	(0.00852)	(0.00243)	(0.0179)	(0.00250)	(0.00189)	(0.0848)	(0.0263)	(0.00555)	(0.0676)	(0.0165)	(0.0155)
	7th	0.0568**	-0.0266***	-0.00705	-0.0537**	0.00607	-0.0316***	-0.0733	0.0789***	0.0511	0.0514	-0.0220	0.0844***
		(0.0255)	(0.00646)	(0.00501)	(0.0233)	(0.00674)	(0.00700)	(0.0898)	(0.0199)	(0.0338)	(0.0739)	(0.0175)	(0.0270)

		Leisure Demand (1)						Labor Force Participation (2)					
		Divorced (1d)			Widowed (1w)			Divorced (2d)			Widowed (2w)		
		Dummy	× Asset Loss	× Sp. Loss	Dummy	× Asset Loss	× Sp. Loss	Dummy	× Asset Loss	× Sp. Loss	Dummy	× Asset Loss	× Sp. Loss
Before Separation (wave)	1st	(control)			(control)			(control)			(control)		
	2nd	9.25e-06			-0.00391			0.0260			-0.0168		
		(0.0122)			(0.00918)			(0.0387)			(0.0389)		
	3rd	-0.00791			-0.0256*			0.0365			0.0628		
		(0.0190)			(0.0135)			(0.0629)			(0.0533)		
	4th	-0.0250			-0.0250			0.00185			0.0749		
		(0.0230)			(0.0209)			(0.0775)			(0.0759)		
	5th	0.0194			-0.0185			-0.00790			0.147*		
		(0.0269)			(0.0146)			(0.0944)			(0.0780)		
Obs.		50,376						50,806					
R^2		0.140						0.152					

Notes: Model 1 represents estimates of a fixed effect regression of leisure demand on regressors as specified in the text. The difference between model 1d and 1w is what is used as the control group. In model 1d, I use the interview wave before the separation for households divorcing. Alternatively, in model 1w, I use the interview wave before the widowing event. This is done to facilitate interpretation. Standard errors, clustered at the person level, are reported in parentheses. Results for the remaining regressors are available from the author.
*** $p<0.01$, ** $p<0.05$, * $p<0.1$.
"× Asset Loss": interaction of assets lost after separation with dummy for wave before/after separation (Assets in \$100,000). ""× Sp. Loss": interaction of spousal income lost after separation with dummy for wave before/ after separation (Spousal income in \$10,000 and is multiplied by number of potential years left to work if spouse would have worked until 65).

Table 11.3: Estimates of Leisure Demand and Labor Force Participation of Men accounting for asset loss and loss of future spousal income following separation.

		Leisure Demand (1)						Labor Force Participation (2)					
		Divorced (1d)			Widowed (1w)			Divorced (2d)			Widowed (2w)		
		Dummy	× Asset Loss	× Sp. Loss	Dummy	× Asset Loss	× Sp. Loss	Dummy	× Asset Loss	× Sp. Loss	Dummy	× Asset Loss	× Sp. Loss
After Separation (wave)	0th	0.0206	0.00152	-0.000420	0.0251	0.0127***	-0.000170	-0.0128	-0.00265	0.000779	0.00961	-0.0472***	-0.00110
		(0.0173)	(0.00302)	(0.000344)	(0.0274)	(0.00410)	(0.000552)	(0.0557)	(0.0106)	(0.00111)	(0.0648)	(0.0143)	(0.00147)
	1st	0.0383*	0.000363	-0.00100***	0.0296	0.00909**	-0.000203	-0.0865	-0.00374	0.00272***	-0.0259	-0.0363**	-0.000120
		(0.0225)	(0.00273)	(0.000294)	(0.0286)	(0.00447)	(0.000632)	(0.0646)	(0.00938)	(0.00105)	(0.0882)	(0.0183)	(0.00155)
	2nd	0.00285	-0.00696**	-0.000233	0.0215	0.00852*	-0.00100	-0.0412	0.00534	0.00174	-0.0389	-0.0300	0.00229
		(0.0258)	(0.00314)	(0.000566)	(0.0284)	(0.00464)	(0.000717)	(0.0685)	(0.0104)	(0.00171)	(0.104)	(0.0190)	(0.00194)
	3rd	-0.00747	0.00517	-0.00132	0.0371	0.00549	0.00246*	-0.0235	-0.0191	0.00483**	-0.0552	-0.0210	-0.0125***
		(0.0392)	(0.00781)	(0.000971)	(0.0375)	(0.00618)	(0.00145)	(0.0749)	(0.0178)	(0.00232)	(0.125)	(0.0213)	(0.00362)
	4th	0.0349	0.0125***	-0.00306***	0.0340	0.00562	0.0195***	-0.0893	-0.0307***	0.00837***	-0.0607	-0.0198	-0.0599***
		(0.0325)	(0.00417)	(0.000873)	(0.0384)	(0.00596)	(0.00405)	(0.0861)	(0.0114)	(0.00238)	(0.132)	(0.0194)	(0.0115)
	5th	0.00460	0.00881***	-0.00255***	0.0281	0.00273		-0.00794	-0.0142	0.00560***	0.0712	-0.0123	
		(0.0274)	(0.00335)	(0.000602)	(0.0422)	(0.00732)		(0.0910)	(0.0114)	(0.00181)	(0.138)	(0.0229)	
	6th	0.0127	0.00819***	-0.00157**	0.0437	-0.00996	0.0650**	-0.00443	-0.0107	0.00517**	0.0434	-0.0118	-0.263***
		(0.0365)	(0.00272)	(0.000667)	(0.0409)	(0.0119)	(0.0253)	(0.102)	(0.0114)	(0.00234)	(0.145)	(0.0450)	(0.0746)
	7th	0.0194	0.122**	-0.00601***	0.0205	-0.00693	-0.00587***	-0.0381	-0.0963	0.0125***	-0.00948	0.0384	0.0191***
		(0.0404)	(0.0517)	(0.00192)	(0.0359)	(0.00753)	(0.00175)	(0.0977)	(0.177)	(0.00433)	(0.165)	(0.0373)	(0.00579)

		Leisure Demand (1)						Labor Force Participation (2)					
		Divorced (1d)			Widowed (1w)			Divorced (2d)			Widowed (2w)		
		Dummy	× Asset Loss	× Sp. Loss	Dummy	× Asset Loss	× Sp. Loss	Dummy	× Asset Loss	× Sp. Loss	Dummy	× Asset Loss	× Sp. Loss
Before Separation (wave)	1st	(control)			(control)			(control)			(control)		
	2nd	-0.0107			-0.0304			0.0176			-0.00124		
		(0.0187)			(0.0298)			(0.0484)			(0.0392)		
	3rd	-0.0175			-0.0444			0.00739			0.106		
		(0.0277)			(0.0302)			(0.0674)			(0.0998)		
	4th	0.0342			-0.0499			-0.104			0.348***		
		(0.0214)			(0.0372)			(0.0729)			(0.109)		
	5th	0.0103			-0.0393			0.0166			0.254**		
		(0.0294)			(0.0336)			(0.0711)			(0.122)		
Obs.		43,263						43,693					
R^2		0.242						0.229					

Notes: Model 1 represents estimates of a fixed effect regression of leisure demand on regressors as specified in the text. The difference between model 1d and 1w is what is used as the control group. In model 1d, I use the interview wave before the separation for households divorcing. Alternatively, in model 1w, I use the interview wave before the widowing event. This is done to facilitate interpretation. Standard errors, clustered at the person level, are reported in parentheses. Results for the remaining regressors are available from the author.
*** p<0.01, ** p<0.05, * p<0.1.
"× Asset Loss": interaction of assets lost after separation with dummy for wave before/after separation (Assets in \$100,000). ""× Sp. Loss": interaction of spousal income lost after separation with dummy for wave before/ after separation (Spousal income in \$10,000 and is multiplied by number of potential years left to work if spouse would have worked until 65).

that a wife only cares about her husband's leisure to the extent that it might impact the household's lifetime wealth. If individuals value joint leisure, then this assumption is no longer valid. If a couple enjoys their time together (i.e. leisure is complementary), then, assuming a spouse is still working full-time in their 50s and 60s, this may encourage the individual to continue working. Once a separation occurs, however, these complementarities are discarded, leading the individual to exit the labor force.

When the specification in §11.3 is re-estimated accounting for spousal labor force participation (full and part-time dummies) before the separation, I find that the effect on women's leisure demand is significant and negative (or positive for women's labor force participation), implying that women view their husband's leisure as complementary for their own. The effect on men is similarly significant and in the same direction, implying that men view their's wife's work as a complement to their own. Including leisure complementarities does not substantively alter any of the previous results. The estimates are included in tables P.1 and P.2.

11.4.2 Pensions and Social Security

Section 9.4 described changes in pension and Social Security benefits which could be correlated with divorce, asset loss, and loss of future spousal earnings. Defined benefit pensions are often split following a divorce by the time rule formula, which is a function of time married while employed relative to total time employed. If a couple is married for a shorter period of time, then the nonmember spouse is entitled to a smaller fraction of the pension member's monthly benefit. Pensions are often not accessible until after the pension member is retired. Using retrospective marital and employment histories from the HRS, I construct measure for whether or not the ex-spouse has retired and the respective time rule formula (conditional on having a pension).

Social Security requires that the ex-spouse must be at least the early retirement age (i.e. 62) before the individual (who also must be at least 62) may claim the divorced spouse's benefit. Furthermore, Social Security provides a survivor's benefit that is the better of the individual's own benefit and the dead spouse's own benefit. We, therefore, would expect that households closer in lifetime earnings might be more affected by the loss of their spouse, because the combined household's income is not protected in the event of survivorship.

I reestimate the model including dummies for whether the ex-spouse is retired, whether the ex-spouse is at least 62, and continuous variables for the time rule formula for divorcing households, and the relative value of survivor's benefit to the household benefit if each member claimed at 62 (or immediately if older). All factors have the expected negative effect on women's labor force participation (i.e. if a wife's ex-husband is older than 62, then she is eligible to collect the divorced spouse's benefit, and hence less likely to work). The effects, however, are generally insignificant. For men, the effects of whether the ex-spouse is retired and at least 62 are positive but insignificant. The effect of the time rule formula is positive and significant at the 10% level. This implies that the more of a pension the husband loses, the more likely he is to continue working. None of these factors significantly alter the effects of asset loss described above. The addition of these factors, however, increases the residual effects of divorce, although the effects are still insignificant.

11.4.3 Children at Home

The effect of divorce on labor supply was largest in the decade following divorce. Although the median household consists of two people before divorce, many of the household still have children at home or temporarily away at school. Whoever takes over responsibility of the children following a separation may be able to collect spousal support, but generally this amount is small. Being a single parent requires covering the costs of a child and, separate from the preferences described in §10, could explain a person's return to the labor force after a divorce.

Accounting for children at home or temporarily away at school does not significantly alter the estimates in the previous section. Children at home cause men to work longer, but have only a small, insignificant, positive effect on women's labor force participation.

12 Divorce and Delayed Retirement

The estimates in the last section show that women's immediate response to divorce is to return to the labor force, and the effect is larger based on the level of assets lost during the divorce. For many women and men, they can not return to work or increase their hours because they are already working full-time. In this case, they might respond by delaying their retirement. While our previous analysis shows the immediate response to an individual's labor supply, in this section I will consider divorce's effect on the individual's retirement decision.

12.1 Estimation Strategy

The previous analysis used a fixed effect framework because the estimating equation followed directly from the household life-cycle model of labor supply and consumption. This framework cannot be extended directly to estimate the effect of divorce on retirement decisions. The retirement likelihood is a cumulative probability of continued labor force participation until a given date. Ideally, the parameters of the model in §10 would be structurally estimated, permitting counterfactual simulations to test the retirement dates of individuals should the divorce or widowing event not have occurred. Structural estimation, however, is a time consuming process, and we instead follow a more straightforward estimation technique common in the reduced-form retirement literature (Coile and Gruber, 2007; Friedberg and Webb, 2005).

Following our previous discussion, I construct a regression framework that uses the divorce/widowhood variables in combination with asset loss, future spousal income loss, and years since separation from our previous analysis. The regression form is:

$$\begin{aligned}\mathrm{RET}_{i,t} \quad = \quad & \beta_0 + \beta_1 \mathrm{SEP}_{i,t} + \beta_1 \left(\mathrm{SEP}_{i,t} \times \mathrm{AssetLoss}_{i,t}\right) \\ & \qquad + \beta_2 \left(\mathrm{SEP}_{i,t} \times \mathrm{SpIncLoss}_{i,t}\right) + \beta_3 \left(\mathrm{SEP}_{i,t} \times \mathrm{WvsSinceSep}_{i,t}\right) \\ & + \beta_4 \mathrm{WID}_{i,t} + \beta_5 \left(\mathrm{WID}_{i,t} \times \mathrm{AssetLoss}_{i,t}\right) \\ & \qquad + \beta_6 \left(\mathrm{WID}_{i,t} \times \mathrm{SpIncLoss}_{i,t}\right) + \beta_7 \left(\mathrm{WID}_{i,t} \times \mathrm{WvsSinceWid}_{i,t}\right) \\ & + \beta_8 \mathrm{EARN}_{i,t} + \beta_9 \mathrm{ASSETS}_{i,t} + \beta_{10} \mathrm{AGE}_{i,t} + \beta_{11} \mathrm{WAVE}_t + \beta_{12} X_{i,t} + \varepsilon_{i,t}\end{aligned} \tag{12.1}$$

where $\mathrm{RET}_{i,t}$ is a dummy variable for whether or not the individual retires in wave t; $\mathrm{SEP}_{i,t}$ [$\mathrm{WID}_{i,t}$] are dummy variables for whether or not the individual has separated leading to a divorce [become widowed] from her first marriage since her first interview wave; $\mathrm{AssetLoss}_{i,t}$ is the negative change in assets following the separation/loss; $\mathrm{SpIncLoss}_{i,t}$ is the loss in future income had the ex-spouse continued to work until age 65; $\text{WvsSinceSep}_{i,t}$ [$\text{WvsSinceWid}_{i,t}$] is the number of interview waves since the separation [widowing event]; $\text{EARN}_{i,t}$ is the individual's earnings; $\text{ASSETS}_{i,t}$ are the individual's current savings, not including defined benefit pensions or Social Security; $\text{AGE}_{i,t}$ is a set of dummy variables for an individual's age from 50-70; WAVE_t is a set of dummy variables corresponding to interview wave; and $X_{i,t}$ corresponds to a set of control variables including: race, education level, health status, quadratic in actual experience, number of children, census division, 13 industry dummies, 17 occupational dummies, number of hospitalization events since last interview (i.e. 0,1,2+), dummy variables for health insurance coverage (i.e. medicare, medicaid, VA, own-employer coverage, spouse-employer coverage, other health coverage), and whether health insurance includes retiree coverage.

I choose measures for future spousal income loss and asset loss that are exogenous from post-separation/loss labor supply decisions. For spousal income loss, $\text{SpIncLoss}_{i,t}$, I use the amount the ex-spouse was making prior to separation interacted with the age difference between the ex-spouse's age before separation and age 65. If the age difference is negative, it is set to zero. This should be exogenous because it does not depend on the ex-spouse's actual labor supply decision. For asset loss, $\text{AssetLoss}_{i,t}$, I use the difference between the period before the separation to the period immediately after the separation. This means the individual has to be observed in both of these waves. I recognize that this is partially endogenous: if the individual immediately responds and begins working, then they could build up assets before the first interview wave after the separation causing the asset loss to appear smaller. Alternatively, if the individual chose not to work, the asset loss could appear larger if the existing assets are being used to finance consumption. This is only true regarding the effect on labor supply in the first interview wave after separation, and not the subsequent interview waves.

I pool observations between 1992 and 2010 of individuals who are full-time employees in the first interview wave, are at risk of divorce/widowhood between ages 50 and 70, and who never apply for disability. If an individual remarries, his or her observations are dropped from that period on. I estimate a probit of equation (12.1), with Huber-White standard errors and adjusted for person-level clustering.

12.2 Estimation Results

Tables 12.1 and 12.2 report probit estimates based on equation (12.1) for women and men. A positive coefficient indicates a higher probability of retiring which I define as transitioning from full-time work, part-time work, or unemployment into partial-retirement, full-retirement, or not in the labor force.[43] The first model is a probit regression of retirement on divorce/widowhood without the interactions in equation (12.1). The second model corresponds to equation (12.1), and includes interactions with asset loss and future spousal earnings. The third model includes an interaction between asset loss and the time since the divorce/widowhood. Each table includes the weighted and unweighted results.

Initially, controlling only for separation (leading to a divorce) and widowhood reveals separation has a small negative effect on women and men, while widowhood has a small positive effect on women and men. None of the effects are precise.

Controlling for the interaction of asset loss, lost spousal earnings, and time since separation/loss reveals that divorce, outside of these interactions, has a significant effect for women, causing their probability of retiring to decrease by 13.2-14.6%. For women, asset loss has a very small effect, and in the unweighted case the wrong sign is observed. Loss of future spousal income is positive, implying that the more the ex-spouse would have been eligible to earn (between divorce and age 65), the more likely the wife is to retire in any given period. This could be the result of wealthier individuals being more likely to provide alimony to their ex-spouses. Longer time between the divorce and the current period implies that the respondent is more likely to retire. This implies that the large effect we found for divorce begins to wear off as time passes, but it still persists even 10-12 years after the divorce. Since asset loss and loss of future spousal income have negligible effects on retirement, the major drivers to retirement decisions are the residual effects: the loss of economies of scale and new preferences for consumption relative to leisure.

The third model, which includes an interaction between asset loss and time since the marital separation, helps to clarify why the time since separation matters. Controlling for this interaction removes the positive effect of asset loss and captures most of the impact of time passing since the divorce. We know from the previous section, that asset loss can be a major reason for why women return to the labor force. If a woman returned to the labor force following the separation due a substantial loss of savings, then the amount of time that has passed will be correlated with her ability to retire. Controlling for this interaction, however, does not eliminate the persistent, negative effect of divorce on retirement.

Comparing the results for women with what we observed in the last section, we find that while asset loss may effect immediate labor supply decisions, with the passage of time, its effect on retirement is limited. The decision of when to retire is primarily driven by other factors that are related to divorce, but that are not directly correlated with asset loss at the time of the divorce. These long run consequences could be the divorced wife's new found

[43]Partial retirement is part-time work where the respondent states in a separate question that they are "retired".

Table 12.1: Probit Estimates of Retirement (Women)

	Women			Women (weighted)		
	(1)	(2)	(3)	(1w)	(2w)	(3w)
SEP	-0.171	-0.704**	-0.636**	-0.0679	-0.772**	-0.739**
	(0.139)	(0.276)	(0.281)	(0.177)	(0.304)	(0.326)
	[-0.036]	[-0.146]	[-0.132]	[-0.015]	[-0.165]	[-0.157]
×AssetLoss		0.0575*	-0.0338		0.0285	-0.0845
		(0.0315)	(0.0474)		(0.0305)	(0.0771)
		[0.012]	[-0.007]		[0.006]	[-0.018]
×SpIncLoss		0.00909	0.0122*		0.0199***	0.0248***
		(0.00661)	(0.00739)		(0.00703)	(0.00923)
×WvsSinceSep		0.137**	0.0430		0.131	0.0222
		(0.0694)	(0.0908)		(0.0814)	(0.115)
×WvsSinceSep × AssetLoss			0.119***			0.187**
			(0.0411)			(0.0734)
			[0.025]			[0.026]
×WvsSinceSep × SpIncLoss			-0.00477			-0.0114**
			(0.00328)			(0.00513)
WID	0.0474	0.0660	0.0835	-0.119	-0.0641	-0.0480
	(0.170)	(0.239)	(0.246)	(0.205)	(0.263)	(0.268)
×AssetLoss		0.0468	0.00403		0.0208	0.0219
		(0.0514)	(0.101)		(0.0650)	(0.109)
×SpIncLoss		0.00466	-0.00335		0.00753	-0.00965
		(0.0112)	(0.0130)		(0.0113)	(0.0145)
×WvsSinceWid		-0.0500	-0.0669		-0.0684	-0.0715
		(0.0907)	(0.0954)		(0.104)	(0.101)
×WvsSinceWid × AssetLoss			0.0153			-0.0281
			(0.0590)			(0.0594)
×WvsSinceWid × SpIncLoss			0.0315			0.0424*
			(0.0243)			(0.0247)
EARN	-2.60e-07	-2.91e-07	-3.10e-07	-1.70e-07	-2.55e-07	-2.88e-07
	(5.21e-07)	(5.25e-07)	(5.27e-07)	(7.62e-07)	(7.69e-07)	(7.74e-07)
ASSETS	0.00557***	0.00555***	0.00553***	0.00702**	0.00701**	0.00700**
	(0.00186)	(0.00185)	(0.00185)	(0.00279)	(0.00278)	(0.00278)
N	8,022	8,022	8,022	5,944	5,944	5,944
Pseudo R^2	0.134	0.135	0.137	0.1244	0.1265	0.1269

Notes: Dependent variable is retirement in the next period measured as a transition from employment in the current period (full-time, part-time, unemployed) to retirement in the next period (retired, partly retired, or not in labor force). Marginal effects included in square brackets for select covariates. Retirement is treated as an absorbing state, so observations after retirement are dropped from the estimation. "weighted" indicates that data is weighted using person-level weights corresponding to the specific HRS interview wave. Regressions are conducted only on individuals between ages 50 and 70 who worked full-time in their first HRS interview wave, were at-risk for a divorce/widowhood between ages 50 and 70, did not remarry, and did not ever apply for disability. Descriptions of the included covariates are in the text.

Table 12.2: Probit Estimates of Retirement (Men)

	Men			Men (weighted)		
	(1)	(2)	(3)	(1w)	(2w)	(3w)
SEP	-0.237*	-0.239	0.0516	-0.310*	-0.235	-0.0146
	(0.125)	(0.226)	(0.237)	(0.169)	(0.294)	(0.302)
	[-0.054]	[-0.054]	[0.011]	[-0.066]	[-0.050]	[-0.004]
×AssetLoss		0.0266	0.0330		0.0334	0.0387
		(0.0195)	(0.0461)		(0.0240)	(0.0444)
×SpIncLoss		-0.00605*	-0.0187*		-0.00643	-0.0153
		(0.00346)	(0.00980)		(0.00484)	(0.00936)
×WvsSinceSep		0.0672	-0.0536		0.0199	-0.0827
		(0.0955)	(0.103)		(0.0990)	(0.113)
×WvsSinceSep × AssetLoss			-0.00516			-0.00759
			(0.0240)			(0.0255)
×WvsSinceSep × SpIncLoss			0.00534			0.00482
			(0.00337)			(0.00355)
WID	0.355	0.512*	0.695**	0.475*	0.473	0.879**
	(0.271)	(0.281)	(0.270)	(0.257)	(0.376)	(0.351)
	[0.080]	[0.116]	[0.157]	[0.101]	[0.100]	[0.186]
×AssetLoss		-0.103*	0.0185		-0.0712	0.166
		(0.0547)	(0.145)		(0.0475)	(0.115)
		[-0.023]	[0.004]		[-0.150]	[0.035]
×SpIncLoss		0.000812	-0.00866		0.00256	-0.00861
		(0.00467)	(0.0132)		(0.00461)	(0.0136)
×WvsSinceWid		-0.350	-0.652**		-0.218	-0.890**
		(0.241)	(0.325)		(0.275)	(0.410)
		[0.000]	[-0.000]		[-0.000]	[-0.000]
×WvsSinceWid × AssetLoss			-0.321			-0.532***
			(0.248)			(0.195)
			[-2.7e-06]			[-3.7e-06]
×WvsSinceWid × SpIncLoss			0.00991			0.0107
			(0.00963)			(0.00947)
EARN	2.90e-08	2.86e-08	2.82e-08	4.27e-09	3.83e-09	3.31e-09
	(1.55e-07)	(1.55e-07)	(1.55e-07)	(1.24e-07)	(1.24e-07)	(1.24e-07)
ASSETS	0.00364**	0.00362**	0.00363**	0.00440**	0.00440**	0.00441**
	(0.00157)	(0.00156)	(0.00157)	(0.00180)	(0.00180)	(0.00179)
N	10,461	10,461	10,461	9,854	9,854	9,854
Pseudo R^2	0.140	0.141	0.142	0.1262	0.1266	0.1274

Notes: Dependent variable is retirement in the next period measured as a transition from employment in the current period (full-time, part-time, unemployed) to retirement in the next period (retired, partly retired, or not in labor force). Marginal effects included in square brackets for select covariates. Retirement is treated as an absorbing state, so observations after retirement are dropped from the estimation. "weighted" indicates that data is weighted using person-level weights corresponding to the specific HRS interview wave. Regressions are conducted only on individuals between ages 50 and 70 who worked full-time in their first HRS interview wave, were at-risk for a divorce/widowhood between ages 50 and 70, did not remarry, and did not ever apply for disability. Descriptions of the included covariates are in the text.

desire to work (i.e. greater preference for consumption), or it could be driven by the loss of marriage's economies of scale. Using this straightforward specification, we can not identify which of these effects are causing divorced women to delay retirement.

Widowed women who are working at baseline have a higher probability of retiring, but the coefficient estimate is imprecise. This is similar to the previous analysis which indicated that widowed women were more likely to withdraw from the labor force in advance of their husband's death.

Men are likely to delay retirement following a separation, but the point estimates are small and imprecise. Widowed men are more likely to retire following the loss of their wife. Asset loss and lost future spousal income have no effect on the retirement decision of widows or divorcées. Both men and women are more likely to retire as they have more assets.

Divorce's effect on delayed retirement is generally imprecise, but two general facts emerge. First, asset losses during divorce do not have a persistent effect on the retirement decision. Women respond to asset loss by immediately returning to the labor force, thereby mitigating its long-term impact on retirement. Second, there is a large residual effect of divorce for women and widowhood for men after controlling for the level of assets lost during divorce and the forgone wages from the ex-spouse. This implies the long-term consequences stemming from lost economies of scale or a shift in preference for consumption over leisure is encouraging the delayed retirement of households. Preference changes are likely to be heterogenous: some people will prefer more relative consumption when they are single, but some will not. On the other hand, economies of scale is uni-directional. A single household's consumption will be much larger than half of the two-person household. While there is no clear cut method of determining whether it is loss of economies of scale or changes in preferences for consumption, I suspect that the magnitude of the retirement delay could only be driven by changes in economies of scale. In future work, a structural model could be used to separate individual heterogeneity in preferences from a common shift in economies of scale.

12.3 Robustness Checks

In this subsection, I consider the effect of my retirement definition on the estimated results, as well as the same issues we considered for robustness in §11.4. Results of the robustness checks are included in appendix P.

12.3.1 Retirement Definition

The results in the last section relied on retirement in the next period being measured as a transition from employment in the current period (full-time, part-time, unemployed) to retirement in the next period (retired, partly retired, or not in labor force). This measure relies partly on the individual to define when they are retired. For example, an individual who transitions from being a full-time teacher to substitute teaching might consider himself transitioning from full-time employment to partial retirement. My retirement measure would treat this individual as retiring. Alternatively, in this section I consider a measure based on annual hours of work. If an individual transitions from working more than 300 hours per year to less than 300 hours per year, the individual is treated as retired (the individual is still required to be working full-time in his or her first interview wave).

The alternative measure of retirement still produces a large decrease in the probability of retirement from a separation, but the effect is smaller. For women, the residual effect of a separation, after accounting for asset loss and lost spousal income, is around 10.6%, compared with 13.2-14.6% using the other retirement measure. No effect is found for the interaction between asset loss and time since separation is found. The elimination of this effect may be due in part to how retirement is being defined. Using the alternative measure, which is based on annual labor hours, would include individuals who are working in bridge jobs. Bridge jobs may be more about enjoyment than they are about financing consumption. This would imply that the individual derives utility from work, or disutility from leisure. As a result, the alternative definition of retirement might not accurately capture the work/consumption tradeoff in the lifecycle model.

Tables P.3 and P.4 in appendix P demonstrate the new estimates from the alternative definition of retirement.

12.3.2 Leisure Complementarities, Pensions and Social Security

I account for married couples' leisure complementarities by including a dummy for whether the spouse is in the labor force and whether the spouse is working full-time. Accounting for leisure complementarities reveals little change to the estimates in §12.2. Couple's leisure time is complementary, although the effect is not statistically significant.

I account for the effect of pensions and Social Security by including dummies for whether the ex-spouse is retired, whether the ex-spouse is at least 62, and continuous variables for the time rule formula for divorcing households, and the relative value of survivor's benefit to the household benefit if each member claimed at 62 (or immediately if older). Accounting for the effects of Social Security and pension structure due to divorce/widowhood increases the parameter estimates for divorce. The general results discussed in the previous section, however, remain the same.

Tables P.5 and P.6 in appendix P demonstrate the new estimates from these robustness checks.

Since the retirement decision does not account for fixed effects, I follow Coile and Gruber (2007) who allow the present value of lifetime benefits to influence the retirement decision. I include the expected present value of Social Security's own-benefits, survivor's benefit, spouse's benefits, and total household benefits (i.e. capturing what one's spouse might be eligible for) over the lifetime. A higher value of one's own benefit decreases the likelihood of retiring, while greater spouse, survivor, and household benefits increase the likelihood of retirement. The coefficients for the expected lifetime benefits are insignificant, and the effects of divorce and widowhood discussed earlier remain the same.

13 Discussion and Conclusion

The Baby Boomer generation has shown a higher propensity to divorce relative to younger generations, and this trend appears to be lasting into retirement. Regardless of future trends, marital separation near retirement has short-term employment effects, forcing individuals back into the labor force, as well as long-term effects in the form of delayed retirement. I have introduced a household life-cycle model of labor supply that captures the transition from a two-person household to two one-person households. The model shows that a separation alters labor supply decisions based on four contributing factors: household asset loss, lost future spouse income, changes in fixed household preferences, and economies of scale.

I find that the labor force participation response to divorce is small for men but is as large as 17.1% for women in the ten years following a separation leading to divorce. When the effects are separated to capture the relative influence of each of the contributing factors, it is asset loss that drives the female labor supply response immediately after the separation. A non-working woman is 6.0% to 10.9% more likely to work for every $100,000 lost in non-housing assets. Other factors, such as lost future spousal income, exhibit limited influence on short-run changes to labor supply. In the long-run, divorce primarily effects retirement decisions through the loss of economies of scale, or greater preference for consumption. Women divorcing in their 50s or 60s will be 13.2-14.6% less likely to retire at any age after controlling for losses in assets and future spousal income.

These results suggest that wealth shocks may be the primary reason for women returning to the labor force following their divorce. Divorce courts currently give little consideration to forward-looking household decisions. When assets are split at divorce, the woman's earnings potential is often exaggerated, particularly if she has not been working. She might have to drawdown her remaining assets while finding a job and establishing her new home. Moreover, she is more likely to move into a lower earning job relative to her ex-spouse, reducing the likelihood that she will be able to accumulate savings to prepare for retirement.

Policymakers and courts trying to ascertain equity in divorce are interested in the welfare of post-divorce individuals. This study demonstrates that we do observe the predicted effects of a household life-cycle model of labor supply that incorporates divorce and widowhood. It does not, however, calculate welfare effects. It also documents other curious anomalies, such as early Social Security benefit claiming among widows and divorced women, greater disability application rates among divorced women, and persistently higher poverty rates for divorced women. In future work, a structural model of household decision-making that accounts for divorce may be better able to measure the life-cycle effects, and provide stronger evidence of changes in men and women's welfare following divorce.

Chapter 3:
Structured Contributions and Directed Search

14 Introduction

A common thread in the organizational design literature is that a business environment often contains complex, tightly interconnected activities[44] that play important roles in creating and sustaining a firm's competitive advantage (Milgrom and Roberts, 1995a; Rivkin, 2000). Firms facing these environments must consider that the marginal cost and benefit associated with any activity depends on the configuration of others activities (Rivkin and Siggelkow, 2003). If activities are complementary then changing one activity amplifies the return from the other activities. For example, if a firm offers a digital music player, it might increase its sales by offering a music downloading service. Alternatively, two activities are substitutable when changing one activity reduces the return from the other activity. In recent years, the popular press has highlighted that the introduction of tablets may be responsible for declining computer sales (Sherr and Ovide, 2013; Worthen and Sherr, 2012; Wingfield, 2013).

Authors wishing to analyze how organizational design affects a business' response to its environment have oscillated between two modeling extremes. One assumes that a firm can directly maximize its profit because activities interact in a perfectly known framework (Milgrom and Roberts, 1990). The other assumes that the business landscape[45] is complex: a firm can only distinguish its optimal performance through exhaustively searching all possible configurations of activities (Levinthal, 1997). Both literatures require performance-maximizing firms to configure a set of interdependent activities in a consistent manner (Porter, 1996). The outcome of such efforts determines the firm's competitive advantage and its sustainability. The effectiveness of an organization's structure in adapting to a complex business landscape is determined by the interdependencies among activities (Siggelkow and Levinthal, 2003). Can complementarities within these interdependencies alter the effectiveness of an

[44] An activity is a discrete economic process within a firm, such as producing a completed product, training employees, employee benefit programs, or marketing decisions, that can be configured in a variety of ways (Porter, 1985).

[45] The business landscape is the realization of the firm's total performance from every possible combination of activities. Performance could be measured as the firm's profit, or it could any other measure by which the firm determines its success.

organization's design in adapting to a complex business environment?

We find that complementarity or substitutability between activities, more than the quantity of interactions, determines the effectiveness of an organization's design in adapting to a new business landscape. If a business exists within a landscape where all of its interconnected activities are complementary, than a positive improvement in one activity provides positive feedback for other activities. As a result, an organization can speed its adaptation by decentralizing its organizational structure. Additionally, we highlight the key role that the variance of the performance gain from these interactions plays in determining the effectiveness of any organization's structure.

The *NK* model, introduced from theoretical biology, offers a useful platform to model the configuration of firm activities (Levinthal, 1997; Kauffman, 1993). The *NK* model permits the researcher to study whether turbulent or stable environments facilitate certain organization structures (Siggelkow and Rivkin, 2005), whether vacillating between organizational structures can improve firm performance (Siggelkow and Levinthal, 2003), how decision autonomy impacts the relative importance of that activity in optimizing the firm's performance (Ghemawat and Levinthal, 2008), and how the degree of interdependence influences industry structure and dynamics (Lenox, Rockart, and Lewin, 2006, 2007). A critical assumption of the *NK* model, however, is that the contribution of an activity to the firm's total performance is randomly drawn for every possible combination of activities, implying that two activities do not interact in a consistent way. For example, if a firm offers a digital music player, this could interact with the firm's decision to offer a music downloading service or to offer its own brand of headphones. The structure of the *NK* model does not allow for the complementary relationship between the decision to produce digital music players and the decision to offer digital music downloading services to be preserved regardless of whether headphones are offered. Recently, researchers have begun to pay attention to the fact that this assumption may not be realistic. Rivkin and Siggelkow (2007) note that the original *NK* model was developed for "biological systems that evolve by mutations to single, randomly chosen genes, [but] it is dubious for organizational, social, and technological systems in which human agents can employ more sophisticated forms of search."

The economics literature provides an alternative framework to examine the interaction of firm activities, recognizing the consistency of interactions through the complementarity or substitutability of firm activities (Milgrom and Roberts, 1990, 1995b). Under a strict set of assumptions, firms will be able to calculate the optimal activity configuration. In reality, however, the assumptions of Milgrom and Roberts' framework (i.e. the differentiability assumptions, functional forms) do not always hold, because firms operate in a complex and uncertain environment. Moreover, their model focuses only on manufacturing firms, whereas

NK models apply more broadly to any firm that offers a combination of goods and services where the method of aggregation of firm-level activities to determine final profitability may not be fully known.

This paper bridges the gap between the complementary interactions of neatly structured models in economics (Milgrom and Roberts, 1990), and the models used in the organizational strategy field to reflect interactions in complex business landscapes (Siggelkow, 2011). We introduce a method of integrating complementary and substitutable interactions into the determination of the firm's business landscape. With the introduction of consistent complementarities into the business landscape, the natural methods of exploration and adaptation reflect the formation of synergies within the landscape that provide important feedback to the adapting firm.

In examining different methods of search, we allow firms to incorporate information regarding the complementarity of their activities. The purpose of this assumption is to separate when additional information regarding activity interdependencies can be helpful in improving a firm's adaptation. Our results show that a small amount of information regarding a firm's interdependencies can significantly improve the speed and performance gains from local adaptation under certain circumstances. This provides us with a theoretical framework from which to test strategies for organizational adaptation that will permit a firm's management to exploit the "lowest lying fruit" for improving their firm's performance. This approach builds on the prior literature that incorporated managerial understanding of these interactions through cognitive representation of high-dimensional landscapes (Gavetti and Levinthal, 2000), modularization of interactions (Ethiraj and Levinthal, 2004b), and broadened search with patterned interactions (Rivkin and Siggelkow, 2007; Ghemawat and Levinthal, 2008). This literature, however, has yet to incorporate how consistent interactions, whether complementary or substitutable, would impact the business landscape and the consequences of these consistent interactions on the search and adaptation process (Rivkin and Siggelkow, 2007).

A few examples demonstrate the importance of incorporating the complementarity of activities within a business landscape for the effectiveness of an organization's design. Siggelkow (2001) explored the interaction of Liz Claiborne's activities as they moved from a position of market dominance into a more competitive environment. Starting in the early 1990s, Liz Claiborne began to offer retailers the opportunity to reorder specific clothing lines after Claiborne's inventory had run out. The firm, however, changed this part of their production model without accounting for the lead time it required to get the product to the retailer which meant that in many instances there was either an inventory shortage or surplus in Claiborne's warehouses (Siggelkow, 2001). Liz Claiborne's management did not account for "offering reorderings" complementary effect on shipping times. Had the firm accounted for this interaction, it might have avoided costly shifts in its inventory levels.

More recently, the growth of digital music downloads demonstrated the complementarities between the hardware and software divisions of technology companies. Apple's management exploited these aspects and used them to make a digital music library consumer friendly so that a market for legal music downloads could be developed. Initially, Apple did not have the resources to become a large digital music retailer. Other firms, such as Sony, were in a better market position to bridge the computer and music divide. Sony had divisions dedicated to computer hardware and software, in addition to being one of the largest music companies in the world. Sony failed to exploit the missing market for digital music downloads due to the firm's decentralized organizational structure. This is just one example (of many) where knowledge about activity complementarities within a company's structure permitted the management to direct their search in such a way as to improve the firm's performance. In Apple's case, the decision to connect the hardware and software experience led them to recognize the complementary nature of music listening and the retail experience which led to their own direct music downloading service. This service, the iTunes Store, became the most dominant retail source in the music industry by 2008, five years after opening. Examples of entrepreneurs exploiting the complementarity in their businesses abound, including large company examples such as IBM's move from hardware to consulting, Google's monetization of search data for marketing, and small company examples of local restaurants expanding to include catering services, specialized grocery stores, or bakeries.

Our analysis contains the following three components. First, we analyze the effect of structured contributions on the firm's business landscape. Structured contributions mean we specifically model the direction of performance gains from activity interactions, ensuring that they are consistent (e.g. if activities A and B are complementary, they will always be so regardless of other activities). This is in contrast to the existing literature, which assumes the contribution from activity interactions is random. The randomness assumption results in two activities interacting in an inconsistent way. For example, randomness could imply that more marketing improves a firms performance if it sells a product in one or three color options, but not with two color options. We introduce a simple method of endogenizing complementarity and substitutability between activities, and permit for variance in the performance gains from activity interactions. Variation in the performance gain from activity interactions accounts for high-order contextual interactions.[46] We demonstrate that it is the combination of complementarities and performance gain variance that determines the effectiveness of

[46] High-order contextual interactions are interactions between more than two activities (i.e. the performance gain from the interaction of activities A and B is amplified/reduced by activity C).

different organizational designs, and not the quantity of interactions.

Second, we propose different mechanisms for internalizing knowledge of how contributions interact, and will revise the firm's search methodology to endogenize this knowledge. If a firm knows that decisions A and B are complementary, then a positive variation in decision A increases the returns from decision B.[47] Internalizing this knowledge is shown to speed performance conditional on the underlying nature of the interactions, and is suggestive of strategies in turbulent environments. We further study how the nature of activity interactions combine with organizational structure to determine firm performance (Fang, Lee, and Schilling, 2010; Nickerson and Zenger, 2004; Sanchez and Mahoney, 1996).

Finally, we explore the role of performance gain variance, which we call uncertainty, in contextual interactions. Uncertainty reflects what Levinthal (2011) called "mapping of the true underlying decision problem to a simple representation that is, in turn, amendable to deductive reasoning." Managers are not perfectly aware of how their choices are reflected in the performance of the firm, because the firm is believed to exist in a complex and dynamic environment. Expectations over the magnitude and variance of the performance gain from activity interactions influences whether centralization/decentralization of an organization's structure is an optimal strategy in turbulent environments.

This paper makes three contributions to the organizational search literature: (1) it adapts the *NK* model to incorporate consistency between activity interactions so that theoretical models of organizational adaptation have business landscapes capable of reflecting feedback from complementary activities, (2) it provides a framework for understanding which organizational designs benefit from the management learning more about the nature of their activity interactions, and (3) it provides a tractable framework for understanding how complementarities and uncertainty impact the effectiveness of organizational structures in adapting to new or changing environments.

We abstract from interactions across firms, as this would require further characterizing the number of firms within a given market, and the strategic nature of their interactions. Instead, this paper focuses on how the introduction of complementary activity interactions can be used to influence the process of local search and adaptation.

[47]This reflects the notion of strategic complements (substitutes) from game theory, as in Milgrom and Roberts (1990, 1995b). This should not be confused with net complements (substitutes) which do not require positive (negative) cross-second derivatives.

15 Structured Contributions

Firms face a wide array of decisions a few of which include offering multiple product lines, choosing output levels, or outsourcing their human resource departments. The degree of interaction between these activities can increase the difficulty of any manager's task, whether the CEO or a division manager, of optimizing the performance of his or her firm.

The organizational adaptation literature frequently uses the *NK* model to encapsulate the complexity of organizational choices. This literature highlights the belief that economic actors are boundedly rational, and therefore can not carry out a complex optimization problem of the kind frequently found in the literature (Ethiraj and Levinthal, 2004a). Moreover, firms can be large organizations with thousands of activities where many of these activities are discrete or yield nonlinear payoffs and cannot be solved with standard optimization techniques. To this end, the *NK* model provides a tractable, structural framework where simulated managers can explore a business landscape in an effort to optimize the performance of their combined choices. Bounded rationality is reflected in the model by each manager being provided limited observational scope. The manager is only able to observe the total performance of the firm, and cannot identify how the individual decisions contribute to the total performance. For practitioners, this method can provide a way of understanding the likely success of different methods of organizational adaptation, such as the context within which centralization versus decentralization of the organization structure can improve performance (Siggelkow and Levinthal, 2003), or how to balance search versus stability (Rivkin and Siggelkow, 2003; Ethiraj and Levinthal, 2004b).

The *NK* model assumes structured interactions. We introduce an additional assumption: structured contributions. Structured contributions imply that interactions between decisions operate in a consistent manner. For example, if greater marketing expenditures improve a firms performance if it sells a product in one or three color options (i.e. the decisions are complementary), then it should also do so if the firm sells it in two color options. The contribution of each decision is structured in such a way that it provides a directionally consistent performance gain if interactions are complementary, or consistent performance loss if the interactions are substitutable. While the contribution to total performance may be directionally consistent, the way we model it permits the magnitude of these interactions to be variable.

In this section we will provide a brief background on how the traditional *NK* model works and what researchers may achieve using it. We will then introduce a mechanism that will provide for the consistent interaction of complementary and substitutable choices within the *NK* model (Milgrom and Roberts, 1990). Consistent interactions lead to a landscape that has greater logical consistency and will provide feedback to improve the firm's search and adaptation process.

15.1 Modeling Complexity: Landscapes in the NK Model

The wide array of complex decisions that firm managers face are typically modeled using the *NK* model which releases a simulated agent onto a business landscape where the agent searches both incrementally and blindly to improve his or her performance.[48] Unlike an optimization problem, the *NK* model assumes the agent has no knowledge of the underlying nature of his or her firm or market, which is like assuming a manager has no knowledge of how his or her divisions interact to produce the product and how this product yields profit for his or her firm. The manager can observe the value (or performance) of the combined choices' output, but knows nothing about how these choices aggregate. Intuitively, this structure assumes that a manager can only see the firm's balance sheet (i.e. whether a firm earned more or less profit relative to the previous period) and what activities the firm has changed this period. Levinthal (1997) describes this as the process of local adaptation, which typically leads to a few dominant organizational forms within a market. In the next section, we will focus more on this search process, but first we explore how the *NK* model generates the business landscape.

Local adaptation relies on a manager making single changes to his or her activity set. This implies that understanding how the *NK* model generates a business landscape is critical to understanding any caveats to local adaptation results. If a firm's manager only makes incremental changes, then he or she could prematurely end adaptation upon finding a local maxima. As interactions become more complex and local peaks populate the business landscape, the firm is susceptible to premature lock-in. Therefore, the model that determines the business landscape determines the feedback firms receive and respond to while undergoing local adaptation. In modeling complex interactions, the *NK* model as proposed by Kauffman (1993), assumes that each combination of decisions yields its own unique performance level and that this performance can change based upon how other decisions change. Specifically, a firm makes N binary decisions about its organization, such as Apple might choose between offering the iPad in many colors or one color, or it may choose whether or not to have a marketing campaign on college campuses, or it can choose whether or not to have promotional offers (such as offering an iTunes gift card with each sale). For simplicity, we will assume that $N = 10$, despite the fact that it is more likely that a company could have hundreds if not thousands of activities. As a result, the space of possible decisions has a total of 2^N

[48]The reader might find it helpful to consider a hiker trying to find the highest peak in a mountain range; however, blinded by fog, the hiker can only see the path immediately in front of them.

possible organizational combinations (i.e. the specific string of N decisions).

Additionally, each of the N decisions interact with K other decisions such that $K \leq N-1$. In our example, this might be that Apple's choice of offering the iPad in multiple colors may interact with its marketing and promotional choices, but not with other choices such as offering health insurance to its employees. We refer to the structure of these interactions as the *interaction matrix*. Figure 15.1 provides an example where each decision d_i is influenced by itself and K other $d_{j \neq i}$ decisions (for a total of $K+1$ decisions). For example, d_4 in figure 1 is influenced by d_1 and d_6.

The firm will make a set of decisions for every activity (total of N binary decisions) and this organizational combination of decisions, denoted by $\mathbf{d}$, determines the contribution for each decision i, denoted $C_i(\mathbf{d})$, via the interaction matrix. $C_i(\mathbf{d})$ is a random number drawn from a uniform distribution ranging from 0 to 1 for each of the 2^{K+1} decision combinations, or 8 in this specific case.[49] This implies that Apple's color choice (decision i), which interacts with marketing and promotional choices, yields a specific contribution, $C_i(\mathbf{d})$, that in the *NK* model is a random number. The *NK* literature then assumes the firm's normalized performance is the sum of these contributions divided by the maximum possible performance value for any $\mathbf{d}$ (therefore performance is always between 0 and 1):

$$P(\mathbf{d}) = \frac{\Sigma_i C_i(\mathbf{d})}{max_{\mathbf{d}} \left\{ \Sigma_i C_i(\mathbf{d}) \right\}} \tag{15.1}$$

Therefore, Apple's final profitability is determined by the sum of the contributions of all of

[49] Note our delineation in terminology: *organizational combination* is the combination of N decisions that the firm selects, where as *decision combination for i* is the combination of $K+1$ decisions that interact with d_i (including d_i, hence the $+1$)

Figure 15.1: Interaction Matrix, ($K = 2$)

$$\begin{array}{cc} & j \\ & \begin{array}{cccccccccc} d_1 & d_2 & d_3 & d_4 & d_5 & d_6 & d_7 & d_8 & d_9 & d_{10} \end{array} \\ i & \begin{array}{c} d_1 \\ d_2 \\ d_3 \\ d_4 \\ d_5 \\ d_6 \\ d_7 \\ d_8 \\ d_9 \\ d_{10} \end{array} \left[\begin{array}{cccccccccc} 1 & 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 1 \end{array} \right] \end{array}$$

Figure 15.2: Landscape Characteristics - Random Contributions ($K = 2$)

(a) Distribution of Business Landscape

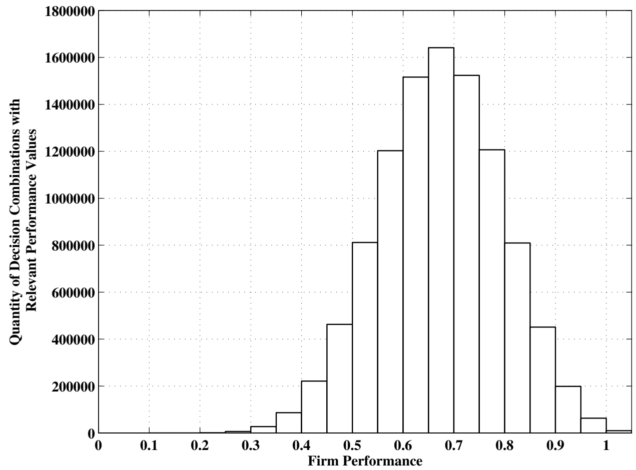

(b) Quantity of Local Peaks by Landscape

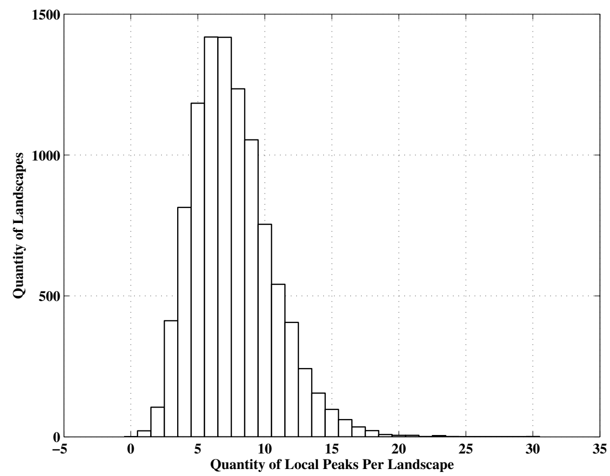

Notes: Simulations done for N=10 over 10,000 landscapes. 100 interaction or influence matrices are drawn and then 100 landscapes are simulated for each of the interaction matrices.

its decisions, $C_i(\mathbf{d})$, such as its iPad color choice, and its marketing, promotional, employee healthcare, and other choices with each given equal weight in the final profit (or performance) of the firm.

An additive contribution structure where all choices are given equal weight creates a distribution of firm-wide performance that is skewed towards a performance of 1 (the maximum possible performance), as seen in figure 15.2a. Due to the random nature of the performance contributions, the landscape is often rugged meaning that it consists of many local peaks - outcomes where no single deviation can improve firm performance, which is illustrated for

Figure 15.3: Increasing Ruggedness of the Landscape

(a) Quantity of Local Peaks by Landscape ($K = 4$)

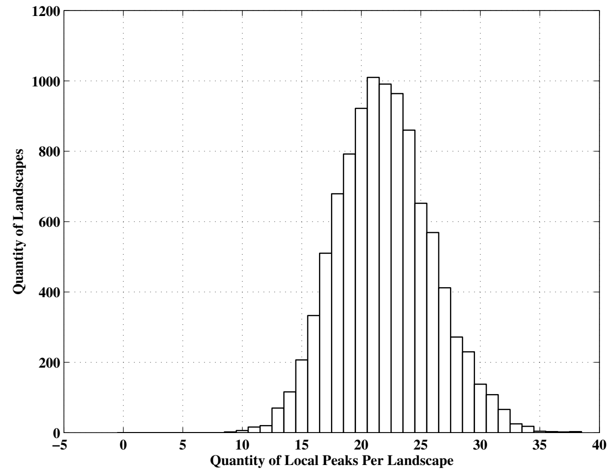

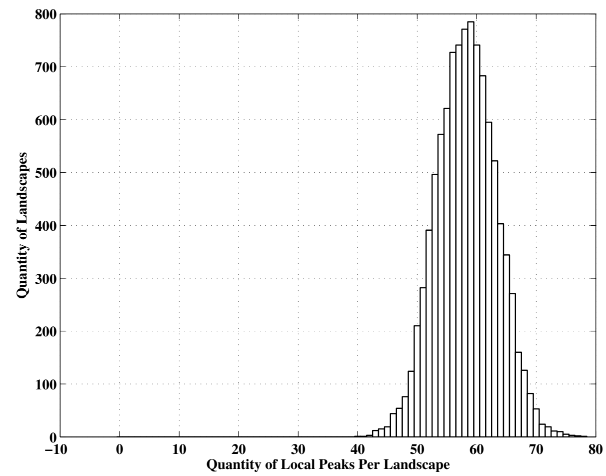

(b) Quantity of Local Peaks by Landscape ($K = 7$)

Notes: Simulations done for N=10 over 10,000 landscapes. 100 interaction or influence matrices are drawn and then 100 landscapes are simulated for each of the interaction matrices.

the case of $K = 2$ in figure 15.2b. Levinthal (1997) showed that as K increased the landscape became more rugged, which is confirmed as we look at the cases of $K = 4$ and $K = 7$ in figure 15.3. This outcome is due to the increased number of decision combinations since the number of organizational combinations is held constant, which is a partial explanation of why researchers maintaining the total number of interactions but altering the pattern of interactions observe increased ruggedness as well (Rivkin and Siggelkow, 2007). Once we introduce consistency across interactions and contextual uncertainty, we will have a more complete picture of the driving factors behind increased landscape ruggedness.

15.2 Consistency across Interactions

In the last subsection, we saw that Apple's color activity (decision i), which interacts with marketing and promotional activities, yields a specific contribution, $C_i(\mathbf{d})$, that in the *NK* model is a random number. This assumption may be appropriate in the context of evolutionary biology, as in Kauffman (1993), where it is unclear how gene types interact with each other, but in a business setting, markets and firms are more likely to have interactions that influence the firm's performance level in a consistent way. Using the Apple example from earlier, if the marketing and color activities are complementary then a performance improvement from marketing should always yield a performance improvement from choice of iPad color. This performance improvement should be robust to changes in other choices as well.[50] Surprisingly, this is not the case with the *NK* model.

In this subsection we introduce a method of endogenizing complementary and substitutable interactions in the *NK* model, so that when contributions are aggregated, the performance of the firm reflects these interactions. We inherit the basic *NK* model setup for §15.1, but permit the interaction matrix to take on a value of 0 for no interaction, -1 for a substitutable interaction, and 1 for a complementary interaction (an example is provided in figure 15.4). If a cell ij in an interaction matrix is drawn randomly as before, we now assign it as a substitutable interaction or a complementary interaction with varying probability (which we will refer to as the percent complementarity in figure 15.2).

Instead of setting $C_i(\mathbf{d})$ by drawing from a uniform distribution, we introduce a new

[50] This line of logic is similar to the contextuality of interactions argument provided by Porter and Siggelkow (2008), but they provide for interactions beyond the direct interactions considered in this paper. Specifically, they permit all choices to interact with one another within a performance function with no uncertainty.

Figure 15.4: Interaction matrix with Structured Contributions ($K = 2$)

$$
\begin{array}{cc}
 & j \\
 & \begin{array}{cccccccccc} d_1 & d_2 & d_3 & d_4 & d_5 & d_6 & d_7 & d_8 & d_9 & d_{10} \end{array} \\
i \;\; \begin{array}{c} d_1 \\ d_2 \\ d_3 \\ d_4 \\ d_5 \\ d_6 \\ d_7 \\ d_8 \\ d_9 \\ d_{10} \end{array} &
\begin{bmatrix}
1 & -1 & 0 & 0 & -1 & 0 & 0 & 0 & 0 & 0 \\
0 & 1 & 1 & -1 & 0 & 0 & 0 & 0 & 0 & 0 \\
1 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\
1 & 0 & 0 & 1 & 0 & -1 & 0 & 0 & 0 & 0 \\
0 & -1 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\
0 & 1 & -1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\
-1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 & 0 \\
0 & 0 & 0 & -1 & 0 & -1 & 0 & 1 & 0 & 0 \\
1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 1 & 0 \\
0 & 1 & 0 & 0 & -1 & 0 & 0 & 0 & 0 & 1
\end{bmatrix}
\end{array}
$$

method of aggregating the contributions of each $d_i \in \mathbf{d}$. First, each decision i starts with its own contribution $C_i(d_i)$ which is drawn from a uniform distribution between 0 and 1. Then each decision $d_{j\neq i} \in \mathbf{d}$ (that interacts with decision i) contributes $\mathbf{I}_{ij} \cdot C_i(d_i) \cdot C_j(d_j)$ to i's overall contribution. If decision j is complementary to i (i.e. $\mathbf{I}_{ij} > 0$), then a positive change in j will lead to a positive change in i, or if j is substitutable for i (i.e. $\mathbf{I}_{ij} < 0$), then a positive change in j will lead to a negative change in i. This setup is consistent with the notion of strategic complements and substitutes in Milgrom and Roberts (1990) seminal paper, but does so in a context that reflects a complex and non-linear relationship between performance level and firm activities.

One of the most important attributes of the *NK* model is that its landscape is so random that it is impossible for any agent to completely learn it, reflecting the concept of bounded rationality. To preserve this feature, we obscure the contribution of each decision by a random variable, γ. If $\gamma = 1$, then with repeated sampling, the agent would be able to credibly distinguish the precise performance contribution of each component, implying that there would be a unique global optima. If γ is a random number, then it introduces variance in the performance gains from activity interactions, which obscured the true contributions of each activity. By obscuring the actual contributions of each component, we introduce uncertainty into the landscape which provides the opportunity for ruggedness.[51] Ruggedness can be increased in the same ways as in the traditional *NK* model, by increasing K and by increasing the number of possible organizational combinations. In equation form:

$$C_i(\mathbf{d}) \quad = \quad C_i(d_i) \cdot \left(1 + \sum_j \left[\mathbf{I}_{ij} \cdot \gamma_{ij} \cdot C_j(d_j)\right]\right) \tag{15.2}$$

Using the structured contributions in equation (15.2), we simulate 10,000 landscapes and influence matrices as we did for the results in §15.1. The choice of what percent of the interaction matrix is complementary has a significant effect of the shape of the landscape. For $K = 2$, the more complementary the interaction matrix, the less right skewed the distribution of performance values across the landscape and fewer peaks. Comparing figure 15.2 to figure 15.5, we choose a 75% complementary matrix to preserve the same approximate landscape distribution, but note that despite this, structured contributions continue to yield many more peaks. This realization helps us understand how dynamic the landscape can be, and it becomes necessary for us to understand how little changes can alter the landscape's

[51] In our model, it is the randomness of γ that prevents the agent from understanding the true underlying model that generates the business landscape. In the standard *NK* model it is the random assignment to every decision combination that instead generates the business landscape. The only difference is that we permit a limited amount of structure that is preserved in the generation of the business landscape.

ruggedness.

Once complements and substitutes are endogenized, we observe that the distribution of the ruggedness of the landscape (i.e. the distribution of peaks) is driven primarily by the variance in the performance gains of the activity interactions in the model, and the mean level of peaks is driven by the type of interaction (complements or substitutes). This is notable, since it has been assumed up until now, that it was the quantity of interactions, not the type of interaction or the assumed uncertainty, that drives landscape ruggedness. This becomes clearer if we hold the level of uncertainty (i.e. variance in the performance gains from activity interactions) and quantity of interactions constant, and alter the complementarity of the interactions. As shown in figure 15.6, as the interaction matrix becomes less complementary (i.e. more of the interactions are strategically substitutable) the ruggedness of the landscape increases in mean quantity of peaks, similar to the increase in K we observed in figure 15.3. Furthermore, if we hold K constant, and decrease the level of uncertainty, we can also decrease the distribution and quantity of peaks. Intuitively, as the landscape becomes more certain, you are in effect making the landscape more consistent and hence smoother. Finally, as K increases in a setting with zero uncertainty (no variance in γ) and a constant level of complementarity, the effect on the distribution is limited. As seen in figure 15.7, as K increases the quantity of peaks remains largely unchanged.

As we will demonstrate in the next two sections, understanding the business landscape is critical because the way activities interact will have a significant effect on how search is conducted. By providing a structured landscape design, we learn that the observation in the literature - that an increase in the quantity of interactions implies an (i.e. K) increase in the ruggedness of the landscape - masks the rich tradeoff between the types of interactions and the uncertainty or randomness that is built into the landscape design. Using the structured contributions framework, as introduced in this section, will allow researchers to tease out the effects of complementary interactions on organizational adaptation in complex and dynamic landscapes.

16 Directed Search

In the previous section, we introduced a way of incorporating complementary and substitutable interactions into the NK model. These interactions allow the business landscape to reflect the formation of synergies within the landscape, which provide important feedback to the adapting firm. In this section, we demonstrate that a firm with a little bit of information, for example, general knowledge regarding the type of interactions between activities, can improve their performance more rapidly than the traditional method of local adapta-

Figure 15.5: Landscape Characteristics - Structured Contributions ($K = 2$)

(a) Distribution of Business Landscape

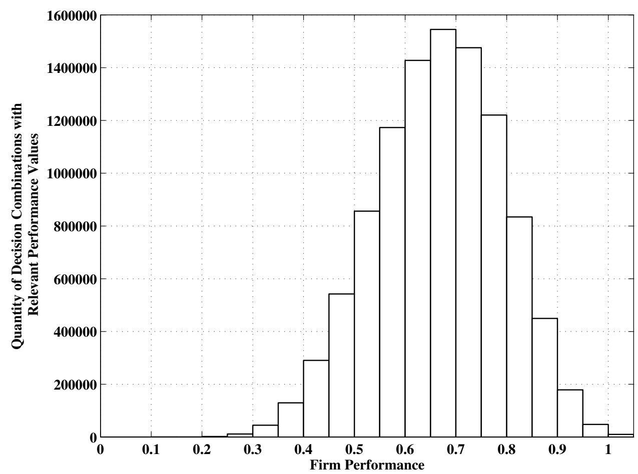

(b) Quantity of Local Peaks by Landscape

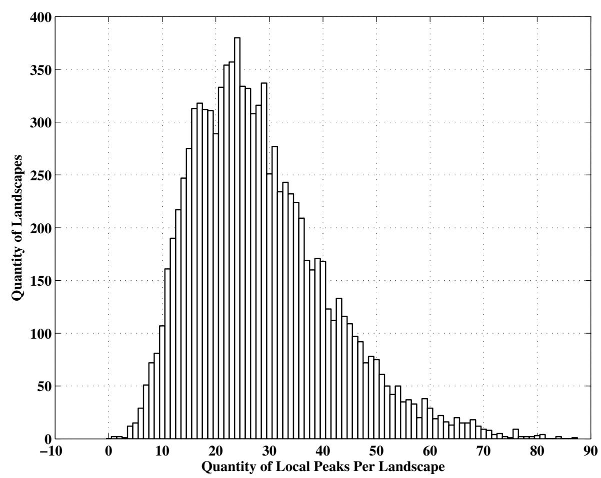

Notes: Results presented assume 75% complementary interaction matrix. The distribution of γ_i is $U[0,1]$. Simulations done for N=10 over 10,000 landscapes. 100 interaction or influence matrices are drawn and then 100 landscapes are simulated for each of the interaction matrices.

Figure 15.6: Increasing Ruggedness of the Landscape with Structured Contributions

(a) Quantity of Local Peaks by Landscape ($K = 2$) with a 25% complementary interaction matrix. (75% substitutable)

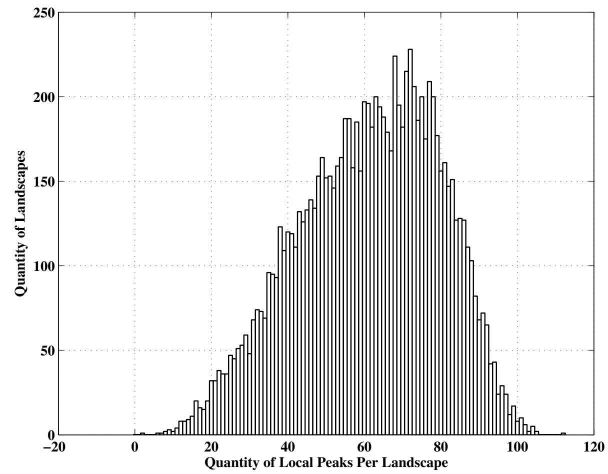

(b) Quantity of Local Peaks by Landscape ($K = 2$) with a 0% complementary interaction matrix. (100% substitutable)

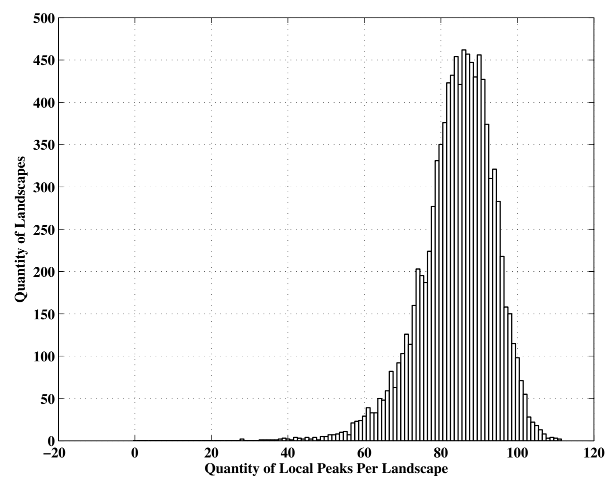

Notes: Simulations done for N=10 over 10,000 landscapes. 100 interaction or influence matrices are drawn and then 100 landscapes are simulated for each of the interaction matrices. The distribution of γ_i is $U[0, 1]$ and is fixed across figures.

Figure 15.7: Increasing Ruggedness of the Landscape with Structured Contributions

(a) Quantity of Local Peaks in a Landscape with no uncertainty ($K = 2$) and a 50% complementary interaction matrix. (50% substitutable)

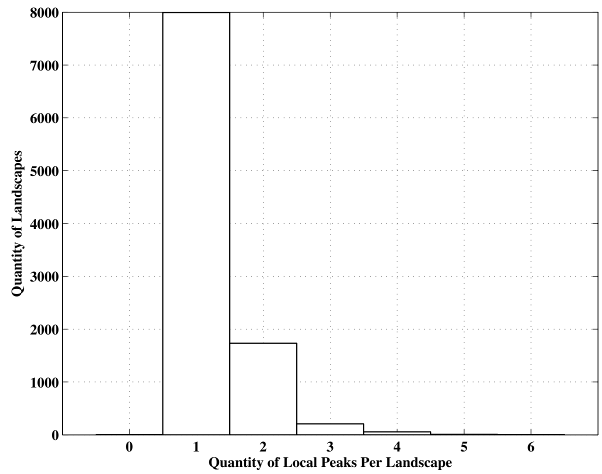

(b) Quantity of Local Peaks in a Landscape with no uncertainty ($K = 7$) and a 50% complementary interaction matrix. (50% substitutable)

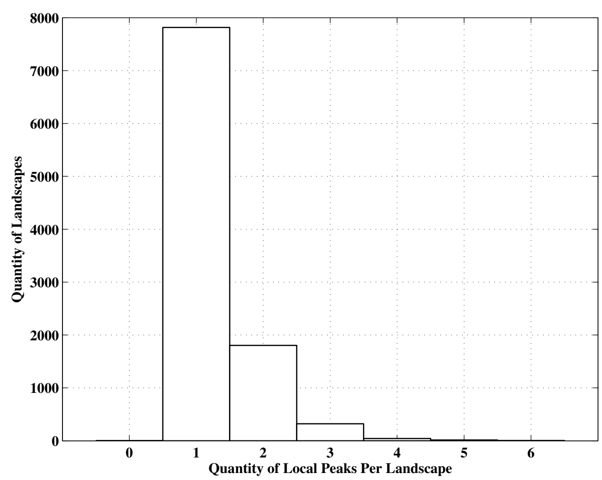

Notes: Simulations done for N=10 over 10,000 landscapes. 100 interaction or influence matrices are drawn and then 100 landscapes are simulated for each of the interaction matrices. The distribution of γ_i is $U[0, 1]$ and is fixed across figures.

tion. Perhaps more surprising, we also show that under certain circumstances, additional information yields no benefit.

In the context of our Apple example, local adaptation suggests that the firm has no knowledge of how the choice of iPad color would interact with a marketing strategy to college students or Apple's employee healthcare. Thus the CEO would randomly select one of these activities, and alter it to observe the impact of this adaptation. If the CEO observes an improvement, he or she is not capable of concluding how that impact reciprocates through their organization. This is like a mountain climber who takes what appears to be a step up the side of a mountain, but is unable to see that the space immediately around him either rises or falls (in essence he is blind). This would be a very frustrating outcome for a CEO (and the mountain climber) as it would likely imply that there movement would be very slow and indirect as each improvement that is taken can result in randomly large or small improvements.

Alternatively, we expect that a CEO has his or her own mental map of how activities interact, or could conduct research to determine how activities interact. With a weak assumption that the CEO has general knowledge of the the interaction matrix - whether the interactions with a specific activity are predominately complementary or substitutable in nature - we can demonstrate that pursuing a directed search strategy may improve the short run performance of the firm. This would be particularly valuable for practitioners who face very volatile landscapes. There are also scenarios, however, where additional information regarding activity interaction yields only limited benefit. For example, we find that when a firm follows a centralized decision-making structure, where the CEO makes all decisions, and the activities within the business landscape are highly substitutable, then there are no gains to directing search. In these circumstances, research into what activities have the most interactions would not be worth learning. Knowing when to invest in learning about activity interactions in a business landscape is an important component of firm management. Our results are suggestive that very simple decision rules for adaptation can yield significant performance gains even with limited information about how activities interact.

16.1 Directed Search Mechanisms

Providing a firm's managment with a little information regarding the activity interactions in its business landscape can greatly improve how it chooses to organize and adapt to changes in the landscape. It is like giving a mountain climber an aerial snapshot of the terrain they are facing - she will not know her position on the landscape, or the exact information about how many peaks exist, or which peak is the tallest, but she will have some idea of what to expect and this will inform how she proceeds to climb the mountain. In corporate strategy, knowledge of what information to collect (i.e. which activities to vary first) would be helpful to practitioners who face the characteristics that the *NK* model is meant to reflect, such as activity complexity, but have more information (or can learn) about how their activities influence their business landscape than the typical *NK* model assumes.

In the organizational design literature, the most frequently chosen search mechanism is local adaptation (Levinthal, 1997; Gavetti and Levinthal, 2000; Rivkin and Siggelkow, 2003). Local adaptation is where only single deviations from the current organizational combination are considered. In most of this literature, the firm accepts the first deviation from its current form that yields improved performance, which is like a CEO considering different alternatives and taking the first available improvement they observe. The process of local adaptation can be varied in different ways, depending on the assumptions of the model. First, we could assume that the CEO considers all N possible single deviations and then chooses the best option, however, this requires the decision-maker to invest in learning about the costs and consequence of N different variations. A second alternative, is to permit the agent to internalize some information that they can use when selecting which activity to invest in learning about. We focus on this second alternative, as search is costly for agents to undertake, so it is infeasible to expect a manager to test all possible single combination deviations if they want to respond quickly to an ever evolving business environment.

In this section we propose a search mechanism that a CEO or division manager could use to speed the process of organizational adaptation and achieve higher average performance. The directed search mechanism described below is chosen to reflect the ability of the firm's leadership to understand the most general interactions between the activities they face. It does not require them to know the magnitude, only the direction, of the interactions. This simple assumption changes the process of how organizational adaptation is conducted by allowing agents to internalize some of their expectations regarding their landscape, and provides a new avenue for refining analyses of organizational design. First, we explore how the current literature treats local search through random search with repeated sampling, and then propose an alternative mechanism based on the agent's limited information regarding the structure of interactions.

16.1.1 Random Search Mechanisms

Random search mechanisms rely on the CEO only observing the aggregate performance of the firm because it is impossible to discern each decision combination's contribution. Moreover, for these mechanisms, we assume that the CEO has no additional information. Two firm-level search mechanisms are:

Random with replacement This is the methodology used in the current literature where a single activity is varied each time with replacement. This assumes that the firm is memoryless, in that it will not be able to recall and learn from its search history.

Random without replacement A single activity is varied each time, and if no improvement is found, then that activity is not repeated again until an improvement is found. We include this mechanism to clearly delineate the gains from non-repetitive search versus the gains from directed search. A "random search without replacement" mechanism makes the assumption that a firm learns from its most recent search history (i.e. it does not retest activities that have previously resulted in poorer performance). This memory is cleared and all activities can be retested following the first successful variation.

16.1.2 Directed Search Mechanisms

Directed search mechanisms assume structured contributions of the type described in §15.2. As a result we can exploit how contributions interact to speed and improve the adaptation process. We introduce only one search mechanism here for brevity, but the *NK* model provides a rich environment in which to test the plausible effectiveness of other proposed search mechanisms. Here, we keep the directed search mechanism simple to illustrate our point that a little information can yield quick and substantial performance improvements relative to a random search mechanisms.

Most Complementary In this method a firm varies the most complementary decisions first (i.e. those with largest sum within a vertical column of the interaction matrix) because these decision are likely to have the most positive impact on the overall performance through the consistent interactions described above. Subsequently the firm varies the next most complementary decisions which are those decisions with the next largest column sum. For example, using figure 15.4, decision d_1 would be the most complementary decision because it has a complementary effect (i.e. a positive change in this decision will increase the marginal returns on the other decisions) on three decisions, and only one substitutable effect (i.e. a positive change in this decision will decrease the marginal returns on the other decisions). Alternatively, decision d_5 is the most substitutable, and hence is the last to be varied in this setup. This process is repeated until no performance improvements are made. After the firm has evaluated each decision once, it would evaluate the most complementary decision again, and then repeat the cycle.

16.2 Directed Search Results

Using the search mechanisms described in §16.1, we simulate the different search mechanisms under increasingly complementary influence matrices for $K = 2$ and the results are shown in figure 16.1. If an interaction matrix is 0% complementary, then this implies that the interaction matrix in figure 15.4 is comprised of only substitutes (i.e. -1) outside of the diagonal. Alternatively, a 25% complementary interaction matrix implies that this is split between 75% substitutable interactions and 25% complementary interactions, and a 100%

Figure 16.1: Averaged Performance Values by Search Mechanism with Structured Contributions ($K = 2$)

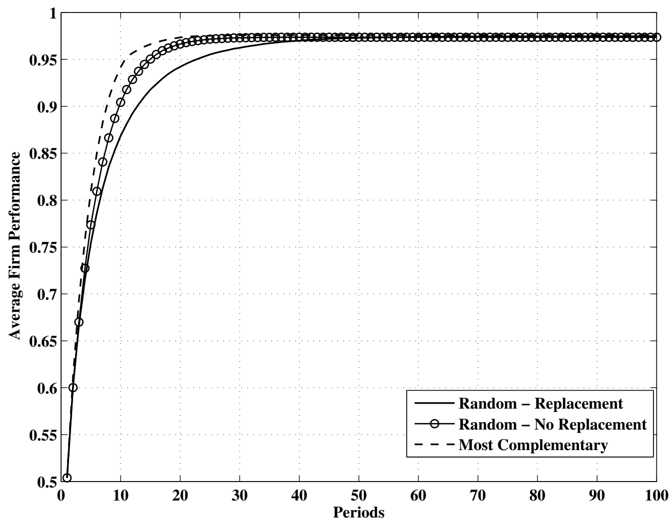

(a) Interaction matrix - 100% Complementary

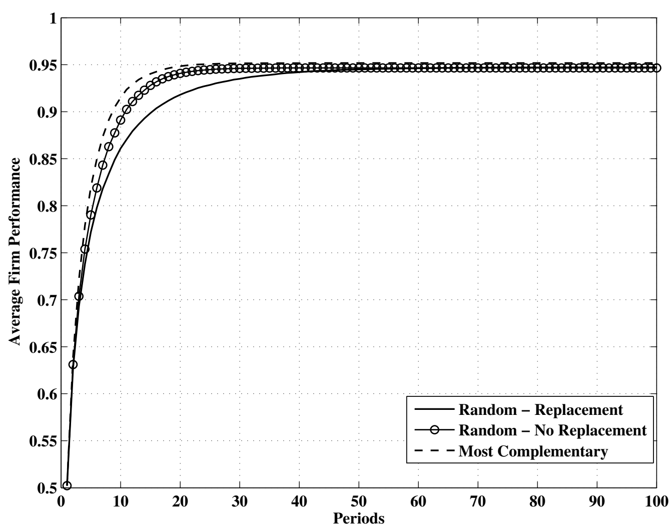

(b) Interaction matrix - 25% Complementary

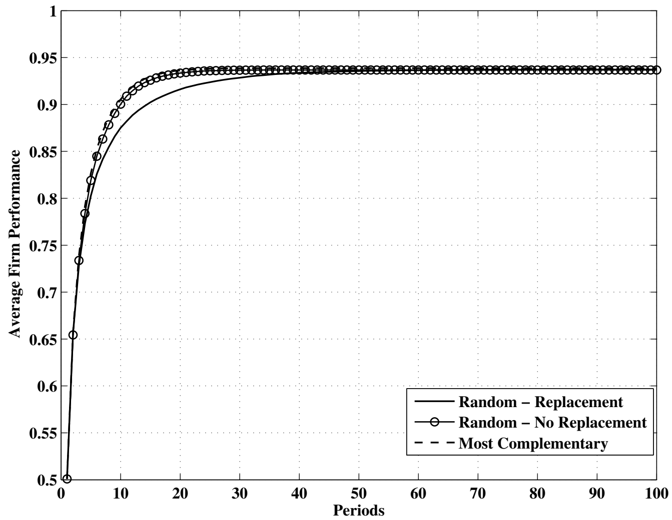

(c) Interaction matrix - 0% Complementary

complementary interaction matrix implies all interactions are complementary (i.e. $+1$).

Non-repeated local search yields substantial performance gains, but these gains are not very informative. These improvements only result from a firm not investing effort in learning about activities that have already been tested. We observe, however, that as the interaction matrix becomes more complementary, there are clear gains from pursuing the directed search strategy. By the 10th period of directed search, a firm pursuing a most complementary directed search mechanism will have, on average, achieved the 95th percentile of performance levels within their business landscape compared to the 91st percentile with non-repeated random search and the 86th percentile with repeated random search.[52] These results suggest that search mechanisms based on a little knowledge of the firm's interaction structure can yield substantial improvements in both speed and level of performance improvement.

Alternatively, consider what happens at the interaction matrix become more substitutable. As the interaction matrix becomes more substitutable, the benefit from learning about the interaction's type diminishes. When the interaction matrix is perfectly substitutable, the firm is equally well-off conducting non-repeated random search. If learning about interactions is costly, or if the firm could possibly be wrong in its belief about how activities interact, then it is optimal for the firm to conduct non-repeated random search.

Heuristics of this type can be very beneficial to practitioners facing complex and dynamic landscapes. They suggest that knowledge about the types of interactions can be beneficial in constructing a strategy for organizational adaptation.

[52] Anytime we are drawing comparisons across different types of influence matrices, we use ranked performance results. The reasoning for this is described in detail in appendix Q.

17 Decentralization with Consistent Interactions

The core idea behind organizational design is that a firm's choice of organizational structure influences its performance. A firm could choose a centralized organizational structure where the CEO makes all of the decisions. Alternatively, the firm could choose a decentralized organizational structure by dividing into divisions and allowing each division manager to simultaneously makes decisions that yield the best performance for his or her division.[53] The choice of an organization's structure will affect how quickly the firm is able to respond to its environment, whether it becomes trapped at local optima during the search process, and the breadth of activities it is able to consider.

A series of papers have highlighted the importance of breadth versus the speed of search (Siggelkow and Levinthal, 2003; Rivkin and Siggelkow, 2003). Breadth versus speed represents a tradeoff between short run gains that are important in very volatile landscapes, and long run performance. A firm that searches broadly is more likely to find a global optimum, but the same method which yields this long-term gain may not be an optimal strategy if the firm faces a very volatile landscape.

Organizational structures provide different advantages. A centralized organizational structure may find higher initial returns from search because the CEO is able to regiment the adaptation process to ensure continuous improvement. It sacrifices the breadth of search. Alternatively, in a very rugged landscape, where a firm may easily find a suboptimal local maxima, a decentralized organizational structure can broadly search, increasing the likelihood of finding the global optima. The broad search of a decentralized organizational

[53]There are many organizational structures (e.g. hierarchial, lateral, liaison; see Rivkin and Siggelkow, 2003), but we present only two here as they represent the most commonly used.

Figure 17.1: Centralized versus Decentralized Organizational Forms ($K = 2$)

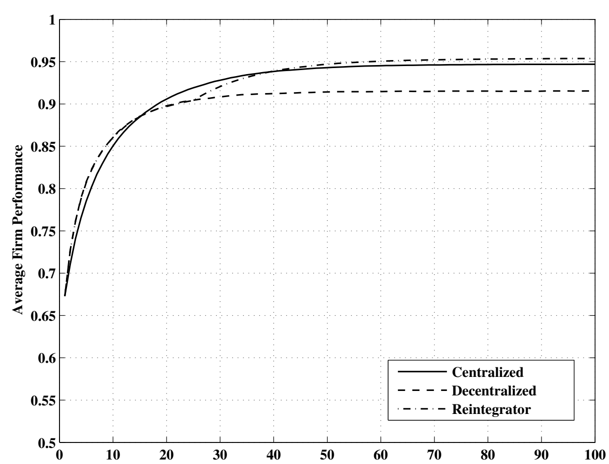

structure has the downside that its divisions may work at odds with each other, stymying the firm's growth. Siggelkow and Levinthal (2003) demonstrated a third organizational type, reintegration. Reintegration is a period of decentralized search followed by centralized search. This organizational structure can improve performance over a centralized organizational strategy, as in figure 17.1. The competitive nature of decentralization permits the divisions to internally optimize before reintegrating, likely placing them in a better location in the business landscape from which to search once the organizational decision-making is centralized.

It is not immediately obvious, but this result relies heavily upon the randomness of the underlying business landscape. If we assume that the underlying interaction matrix is 100% complementary and that the business landscape is generated by structured contributions, then we observe in figure 17.2(a) that decentralization is always the best response. Intuitively, a 100% complementary interaction matrix implies that when one division changes their decision and observes a performance improvement, it will increase the marginal returns on other decisions, including the other divisions' decisions. This means that decentralizing allows a firm to adapt to its landscape faster because each divisions' adaptations reinforce each other.

Additionally, we observe in figure 17.2(b) and (c) that it is not until the interaction matrix becomes very substitutable (i.e. the landscape becomes more rugged) that the centralized organizational structure performs better in the short-run compared to the decentralized organizational structure. In order to generate performance losses from divisions working at odds with each other (as in Siggelkow and Levinthal, 2003) the interactions have to be substitutable and sufficient in number to decrease the marginal returns of the competing divisions.

Only when we include directed search are we able to achieve the short run gains that exceed those from decentralized search. This is because the directed search strategy exploits the nature of the underlying contribution structure and uses this to inform the process of search. It is notable that with a high complementary interaction matrix, like in figure 17.2(a), directed search and decentralization both garner their gains from the complementary nature of activity interactions. These gains, however, are achieve faster with decentralization because two divisions are searching (compared to a single CEO) and hence can explore a larger portion of the landscape.

Figure 17.2: Centralized versus Decentralized Organizational Forms with Structured Contributions *(K = 2)*

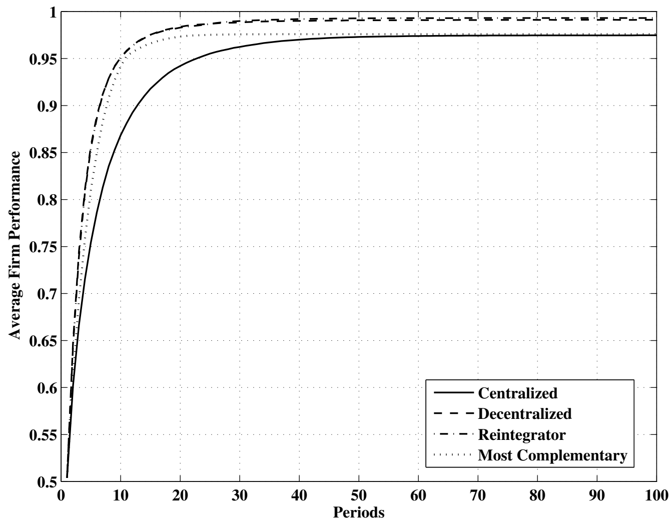

(a) Interaction matrix - 100% Complementary

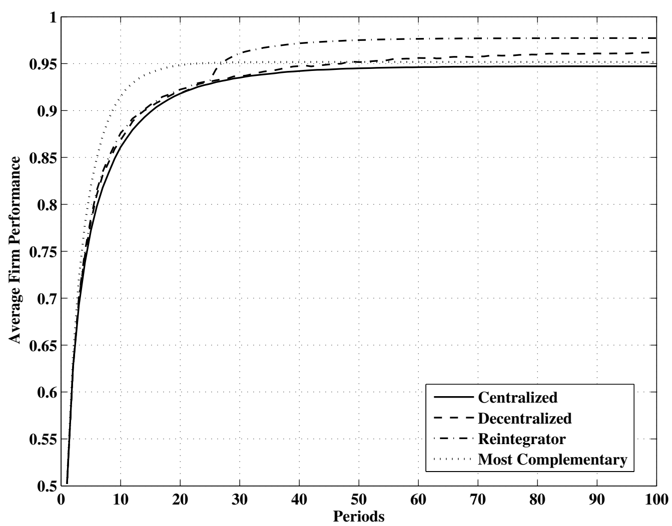

(b) Interaction matrix - 25% Complementary

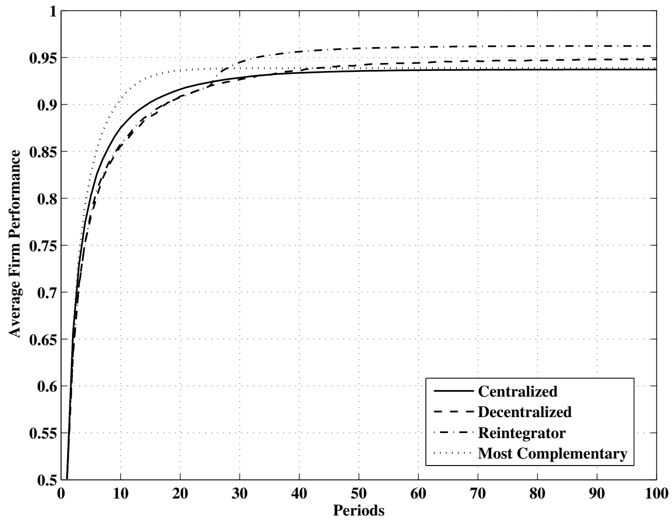

(c) Interaction matrix - 0% Complementary

18 Uncertainty

In §15.2, we introduced the idea of uncertainty along the business landscape. The γ in equation (15.2) represents the uncertainty an agent has regarding the underlying model that generates the business landscape. More accurately, γ can represent the aggregate effect of all the contextual interactions that an agent is unable to foresee because of his or her bounded rationality (Levinthal, 2011; Porter and Siggelkow, 2008). The only assumption we impose is that the effect of the uncertainty surrounding the contextual interactions is unable to reverse the complementary or substitutable nature of the interaction. In fact, γ is capable of amplifying or reducing the interaction effects. It is through this lens that we begin to understand the how complexity and substitutability combine with uncertainty to generate the rugged landscapes previously assumed to be determined based solely on the quantity of interactions, K.

As the mean and variance of γ increases, the underlying parameters of the structural model are further obscured to the agent. While the general interaction structure is preserved, greater uncertainty (high variance of γ) will lead to more local peaks, as in figure 18.2. More local peaks increase the likelihood of premature lock-in before reaching the global optima following centralized search, but agents can avoid this by pursuing decentralized search which permits broad search across the landscape(Levinthal, 1997). The complex interactions of the underlying model that generate the business landscape can lead to high order interactions (interactions between three or more decisions) that can impact the firm's performance. Intuitively, this is like the interaction of Apple's choice of iPad colors and its marketing decision to college students is influenced indirectly by the activity of whether or not to offer Apple computers to high schools. This high level interaction may be part of the true underlying model that generates the business landscape, but is too obscure to be realized by Apple's management. Premature lock-in can be observed in figure 18.1(a)-(c) where increasing uncertainty reduces the average performance of centralized search every period. However, decentralized search escapes the issue of premature lock-in by searching more broadly (see figure 18.1(a) and (b)).

A higher mean will increase the jaggedness of the business landscape. A more jagged landscape implies that deviations will yield more erratic swings in performance. While this effect will not impact the search performance of centralized search (either directed or random), it has significant effects on decentralized search. Specifically, a more jagged landscape will amplify the effects of divisions if they are acting at odds with each other due to the decentralized decision-making process. This can be observed by comparing figures 18.1(b) and (c), where in figure 18.1(b), single deviations have a relatively small effect because of the small mean of γ, but as the mean of γ increases a hundred-fold, then a deviation that is at odds between the divisions will yield a steeper drop in firm-wide performance. Intuitively, if a firm perceives that the net effects of their interactions are large, then decentralized search is inferior to a centralized search methodology because the implicit cost of divisional conflict is greater. It is specifically in these situations where a reintegrator method, which benefits from the initial broad, decentralized search before centralizing (Siggelkow and Levinthal, 2003) yields the best performance improvements.

Uncertainty surrounding the magnitude and variance of contextual interactions contribute both ruggedness and jaggedness to the landscape. It is the tension between these two effects that can ultimately determine whether a landscape is very rugged and whether decentralized or centralized search are more likely to generate short-run performance gains. This theory provides justification for why the literature has observed conflicting results between integration and flexibility in turbulent landscapes. Industries characterized by high complementary activity or non-jagged landscapes are likely to find decentralization optimal in the short-run because divisional decisions tend to reinforce each other (see figure 17.2(a) and figure 18.1(a)). Alternatively, industries characterized by many substitutable interactions or jagged landscapes will find that centralization will be more optimal given environmental turbulence (see figure 17.2(c) and figure 18.1(c)).

19 Conclusion

The wholesale adoption of the *NK* model from the evolutionary biology literature has ignored some key realities of the business landscape, notably, that firms generally have some

Figure 18.1: Averaged Performance Values by Variance of γ and Method of Search with Structured Contributions ($K = 2$)

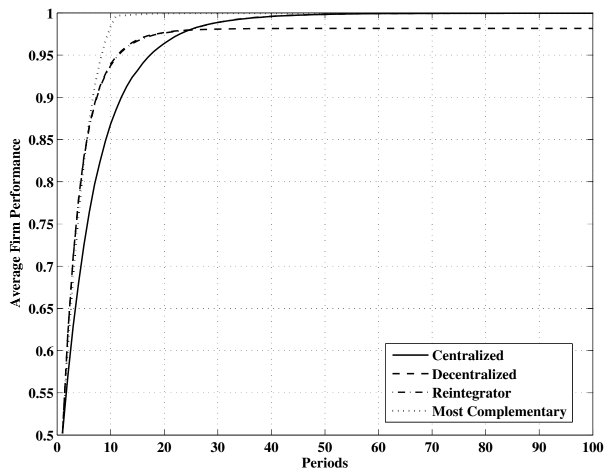

(a) No Variance ($\gamma = 1$)

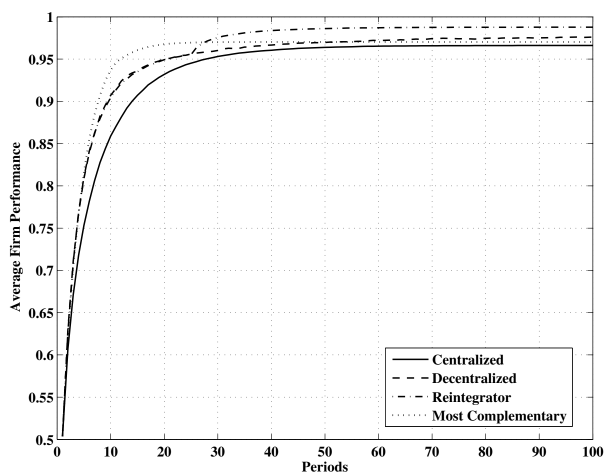

(b) Low Variance ($\gamma \sim U[0, 1]$)

Notes: Results presented assume 50% complementary interaction matrix.

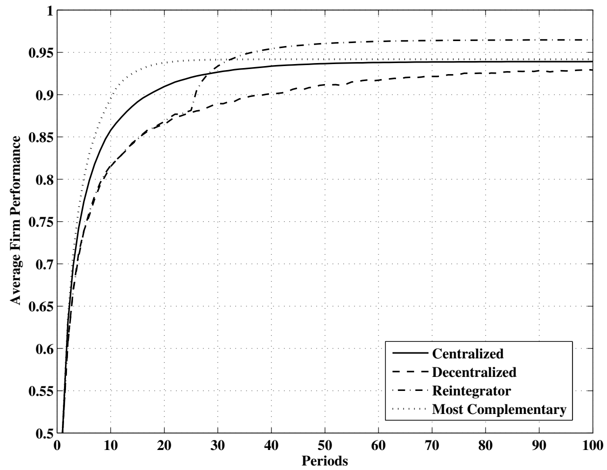

(c) High Variance ($\gamma \sim U[0, 100]$)

Notes: Results presented assume 50% complementary interaction matrix.

Figure 18.2: Quantity of Peaks by Variance of γ with Structured Contributions ($K = 2$)

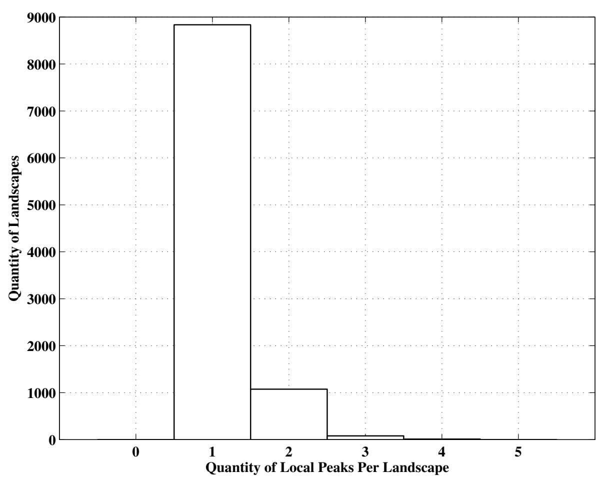

(a) No Variance ($\gamma = 1$)

Notes: Results presented assume 50% complementary interaction matrix.

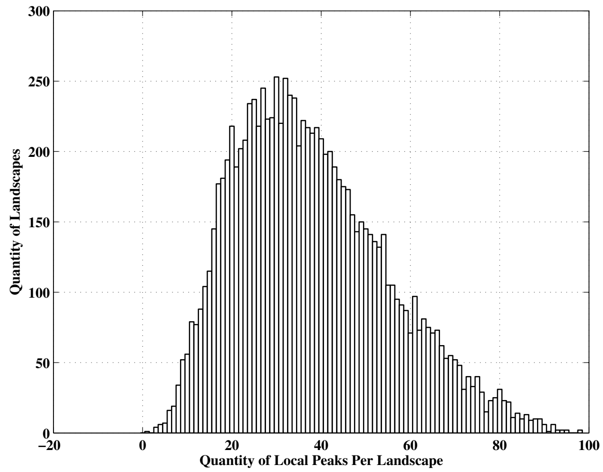

(b) Low Variance ($\gamma \sim U[0, 1]$)

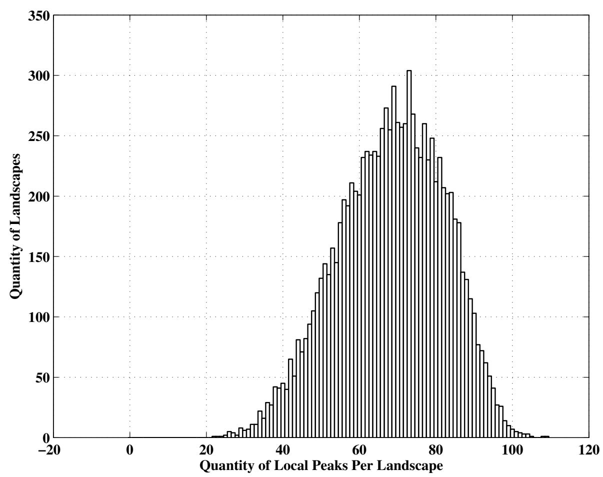

(c) High Variance ($\gamma \sim U[0, 100]$)

Notes: Results presented assume 50% complementary interaction matrix.

knowledge of how decisions are likely to interact with each other, even if they do not know the magnitude of an interaction's influence. We introduce a method for bridging the gap between the simulation models used to model complex interactions with the complementary framework created in the early work of Milgrom and Roberts (1990). Our model introduces the notion of complementarities and uncertain contextual interactions into the *NK* model to help researchers better understand the role of each of these mechanisms in driving both the ruggedness and jaggedness of the business landscape. We have shown that with integrated local search, ruggedness can result in premature lock-in, consistent with past literature, but have discovered that landscape jaggedness, and not the number of peaks (ruggedness), is the mechanism that drives the success or failure of broad, decentralized search.

Building on this new insight, we recognize that competent managers could optimally use even limited information about their firm's activity interactions to improve their search and adaptation process. Thus we provide a simple mechanism that has varying levels of success at improving the speed and performance level relative to the baseline case of non-repeated random search. A natural extension to this current paper would be to expand and test a larger number of simple mechanisms, such as a hybrid random and most complementary mechanism, that could render a more significant improvement in the speed of adaptation.

Future research can build on this basic structure and further explore how centralized and decentralized organizations that misperceive the nature of their interactions are impacted by pursuing strategies based on these misperceptions. This issue is particularly salient given previous work on how the environment often changes due to demand, technological, and institutional factors increasing the likelihood of misinterpreting the interaction of decisions (Henderson and Clark, 1990b; McGahan, 2004; Miller, 1992). Finally, this model can be adapted to incorporate managerial understanding of these consistent interactions through limited cognition (Gavetti and Levinthal, 2000) and modularization of interactions (Ethiraj and Levinthal, 2004b) to see if the tension between complementarities and uncertain contextual interactions will widen or reduce the scope of these adaptation methods.

Our hope is that this paper is an introduction to how complementary organizational structures can be integrated into simulations of complex systems. The development of simple mechanisms will be able to provide real world suggestions to firms' leadership teams as they weigh various methods of adapting their organization's structure to ever changing market landscapes.

Appendices

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.[54] 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.[55] 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 A.1, 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 A.1b, 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

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

[55] 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

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.

Table A.1: 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.

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 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 (2005a), 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:

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

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).[56]

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

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

finer at lower levels of assets since I would expect greater responsiveness at lower levels to changes in 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 υ_j percentage of households (h) with assets below $Q_{\upsilon_j}(A_{ht}, t)$ is defined as

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

   where the quantile index is denoted by j. Put another way, $Q_{\upsilon_j}(A_{ht}, t)$ is the υ_jth age-condition asset quantile. The model analog to $Q_{\upsilon_j}(A_{ht}, t)$ is $\hat{Q}_{\upsilon_j}(t; \theta_0, \chi_0)$ 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

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

   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:

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

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:

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

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

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

for asset quantiles $j \in \{1, 2\}$.[57] 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:

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

   where $L\bar{F}PR_i\left(t, \tau; \theta_0, \chi_0\right)$ 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 $\bar{FT}_i\left(t, \tau; \theta_0, \chi_0\right)$ 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:

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

   I then convert this relationship to an unconditional moment condition:

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

[57] I define $\hat{Q}_{\upsilon_0}\left(age_i; \theta_0, \chi_0\right) = -\infty$ and $\hat{Q}_3\left(age_i; \theta_0, \chi_0\right) = +\infty$

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)

   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:

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

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 $\hat{m}\left(LC_i,u_{is},\theta;\chi_0\right)$ 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:

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

where:

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

such that $D={}^{\partial m(\cdot)}/_{\partial\theta'}\mid_{\theta=\theta_0}$ 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, $D={}^{\partial m(\cdot)}/_{\partial\theta'}\mid_{\theta=\theta_0}$, 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

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

This equation can be rewritten as

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

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

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

Similarly, recall equation (D.2):

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

It can be rewritten as

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

where the first order condition becomes,

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

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* $\in$ {*excellent, very good, good*} or *bad* $\in$ {*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].

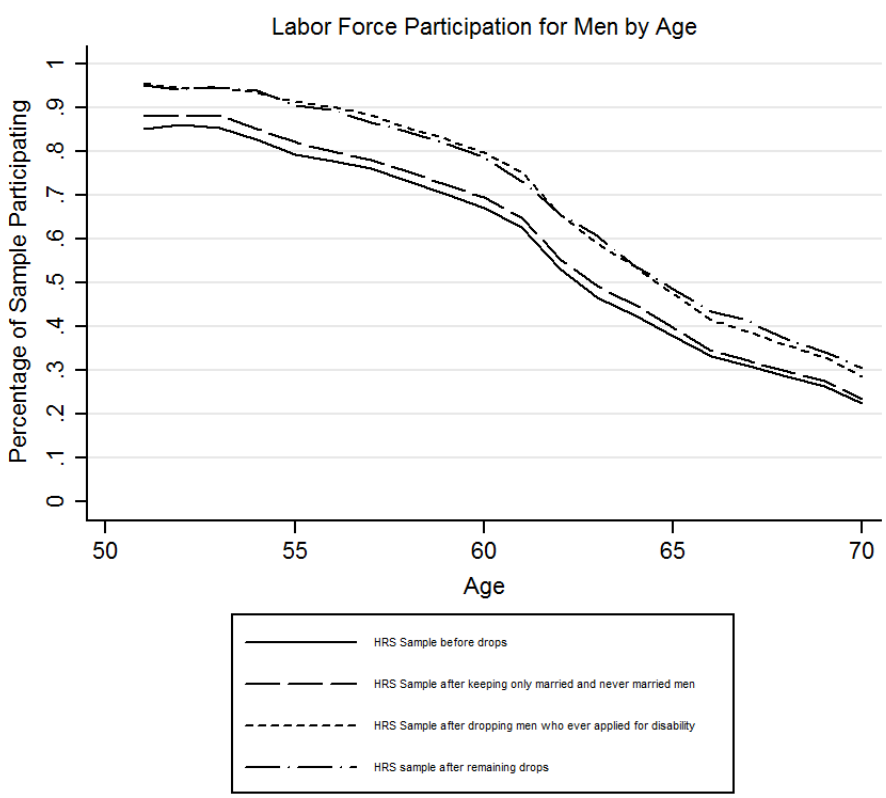

Figure E.1: Sample Selection by Labor Force Participation of Men

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 E.1 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 E.1 provides sample statistics for the entire sample, while table 4.1 in the main text provide the sample statistics for the estimation sample. Finally, table E.2 provides the sample statistics for the model validation sample.

Table E.1: 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 E.2: 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.[58] 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

[58]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.

an 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.[59]

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

[59] 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.

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 investment 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).[60]

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

[60]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.

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.

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.[61] 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.[62]

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

[62] 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 G.1. This is despite median annual hours falling prior to age 57 and then remaining relatively constant for FT-NB jobs (as in figure G.2). 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

Figure G.1: Median Earnings - Non-baseline jobs

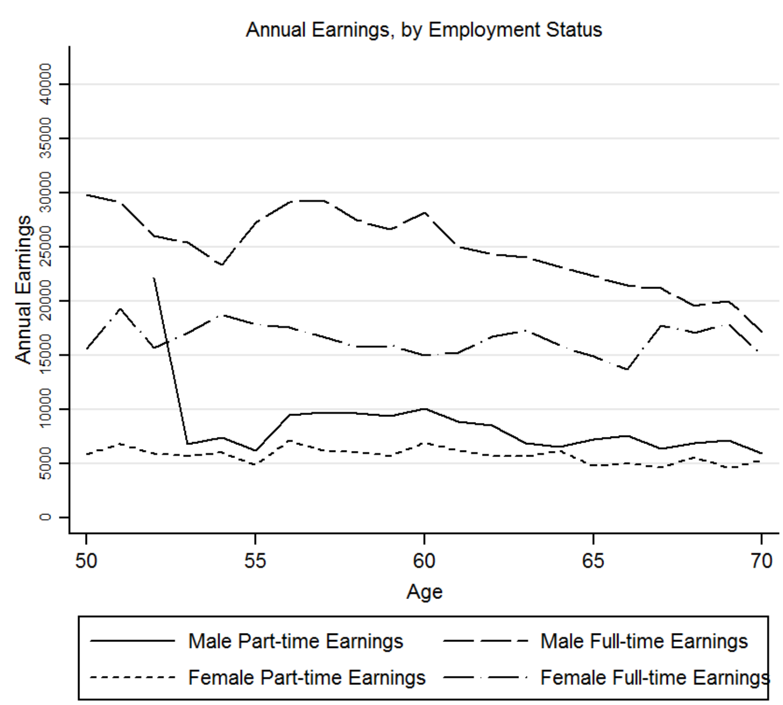

Figure G.2: Median Hours - Non-baseline jobs

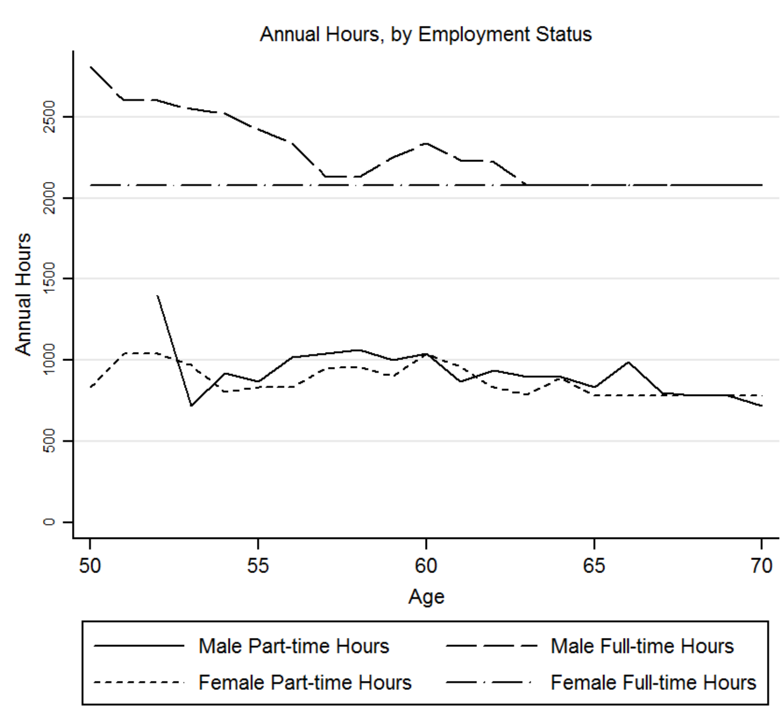

and each subsequent period (up until the maximum working age of 70). Therefore, it is important to assign a feasible wage that a worker might believe is available to her outside of her baseline job (if she is working), or if she was to return to the workforce (if she was not working at baseline).

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 $j \in \{$FT-NB, PT$\}$:

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

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, $\hat{c}_i^j = \bar{\ln w}_i^j - \bar{x}_i\hat{\beta}$ where $\bar{\ln w}_i = \sum_t(\ln w_{i,t}/T_i)$

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 G.3.

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

Figure G.3: Annual Earnings by Employment Status

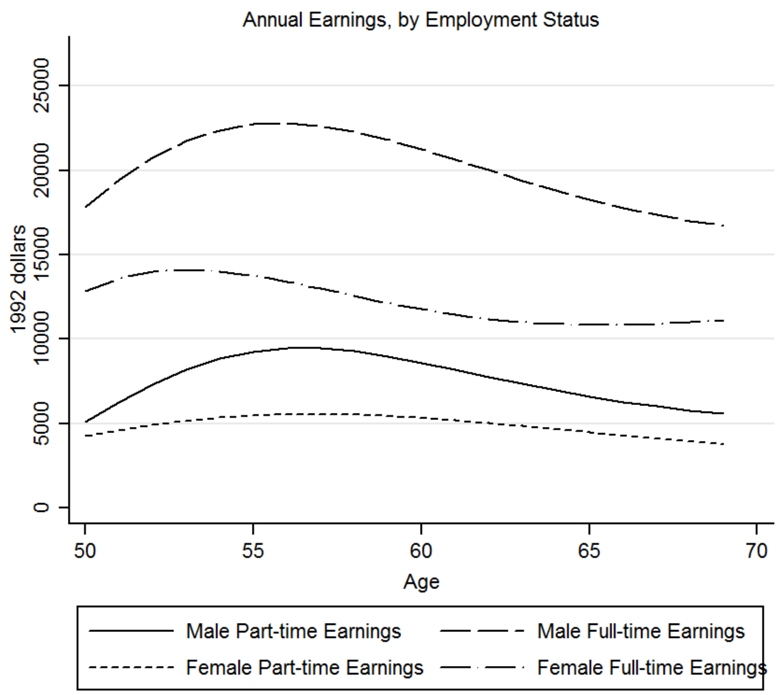

relative to full-time jobs. In fact, since 23% of individuals who have both a PT and FT-NB jobs after baseline have higher part-time wages, it is very likely that changes in observed employment status may be driven by positive shocks during job search.

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 H.1 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 H.2a 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 H.2b shows the comparable result for women.

Figure H.1: Probability of being in bad health
Conditional on previous health status

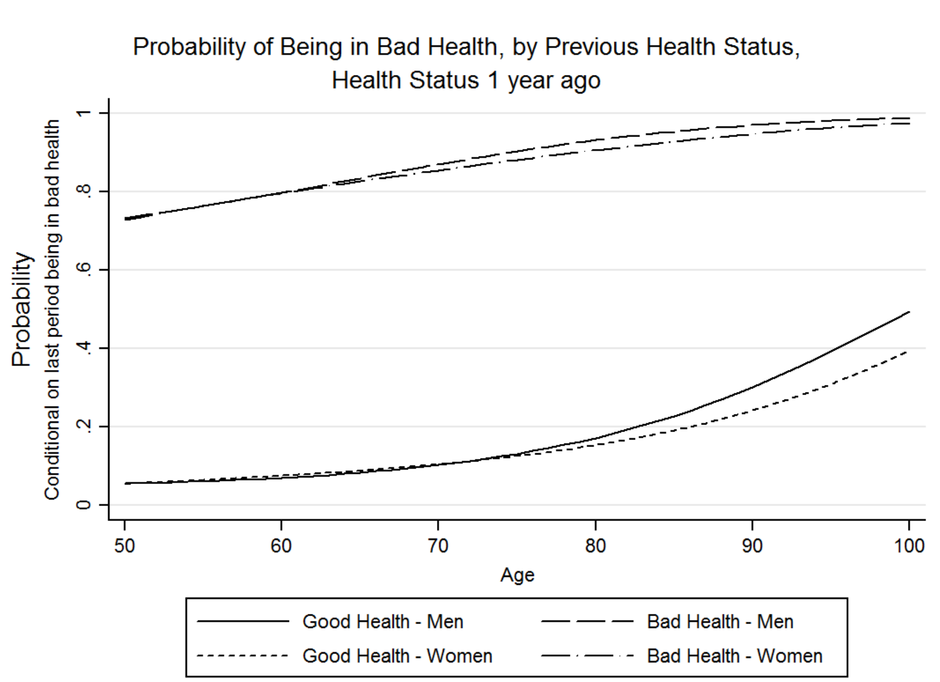

Figure H.2: Probability of Death by Sex
Conditional on previous health status

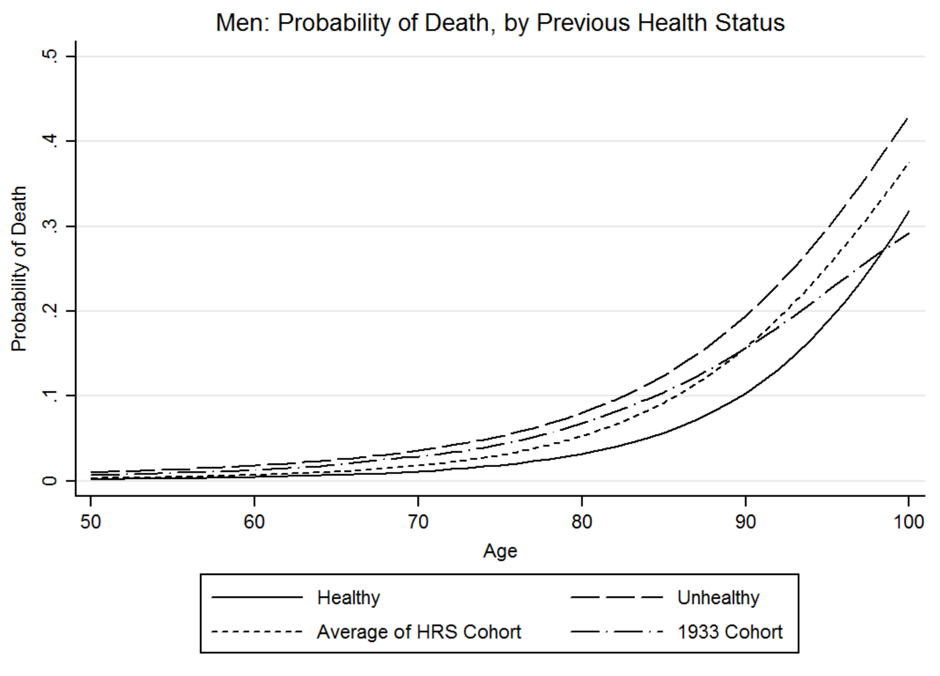

(a) Men

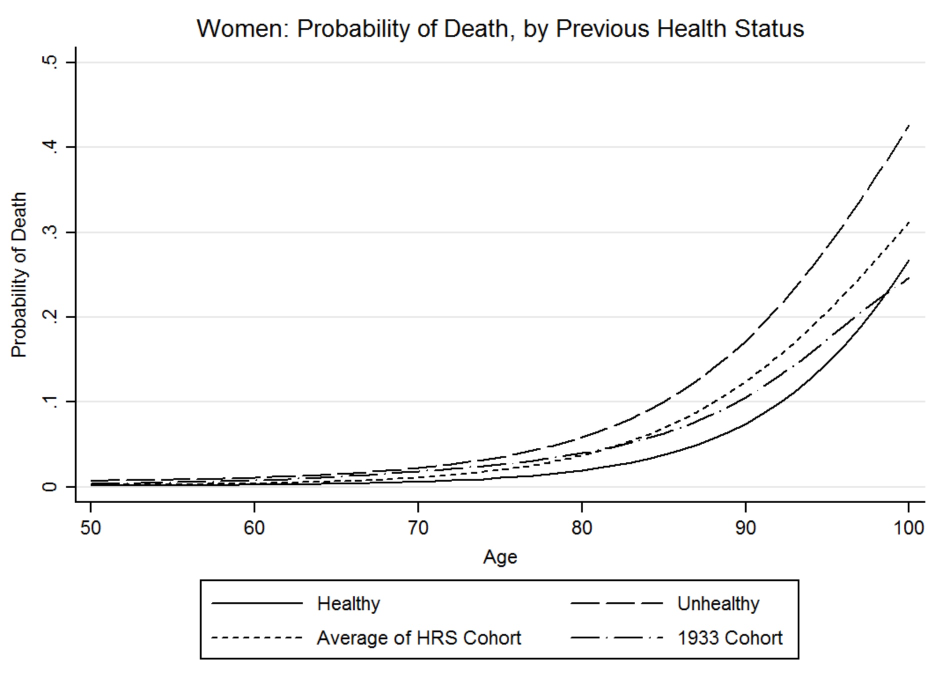

(b) Women

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 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 4.1, 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 K.1: Asset Quantiles (by thirds) by Male Age

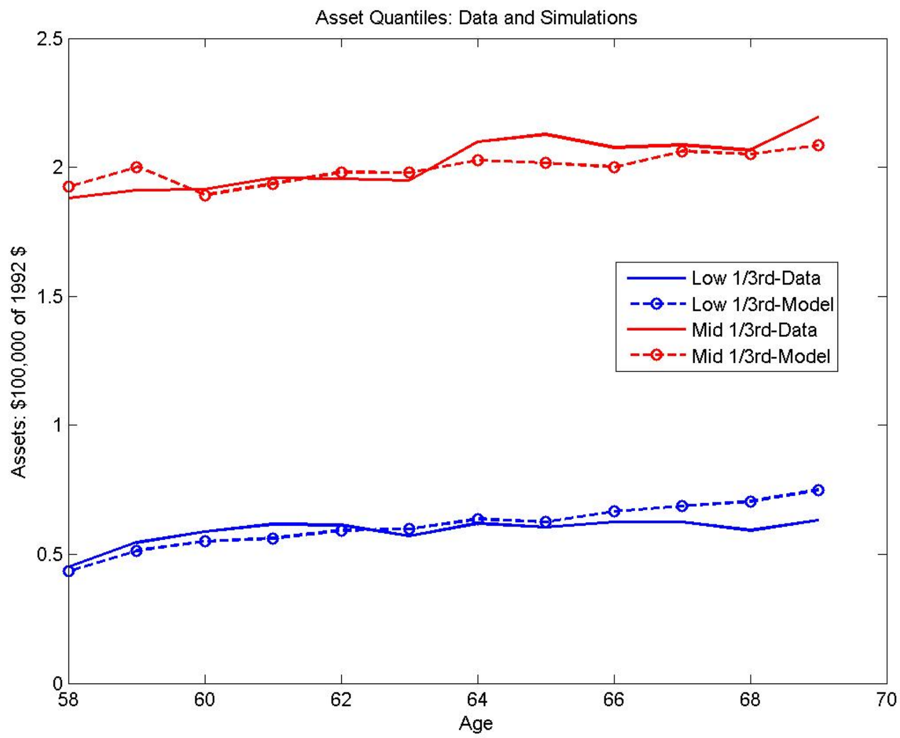

Figure K.2: Asset Quantile Shares by Preference Type

(a) Type 0 - the Out Type

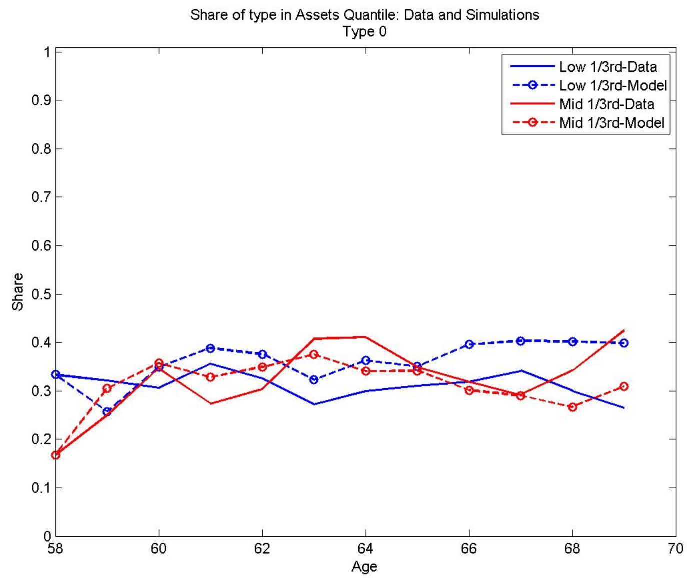

(b) Type 2 - Low Preference for Own Leisure, Low Preference for Spousal Leisure

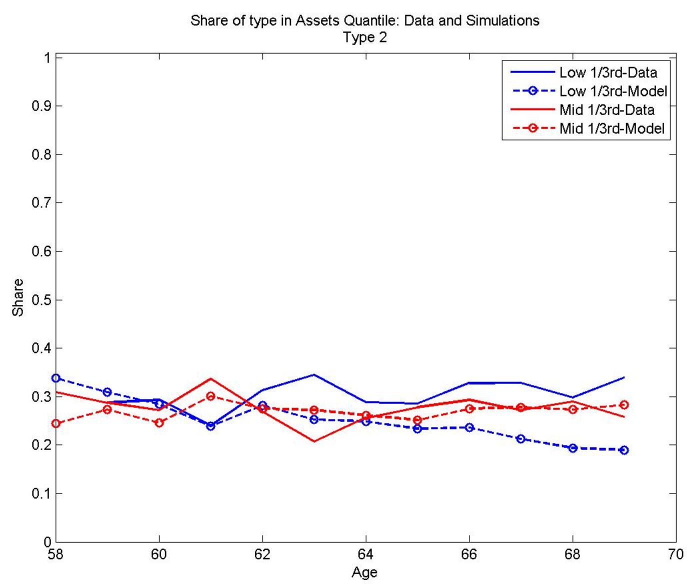

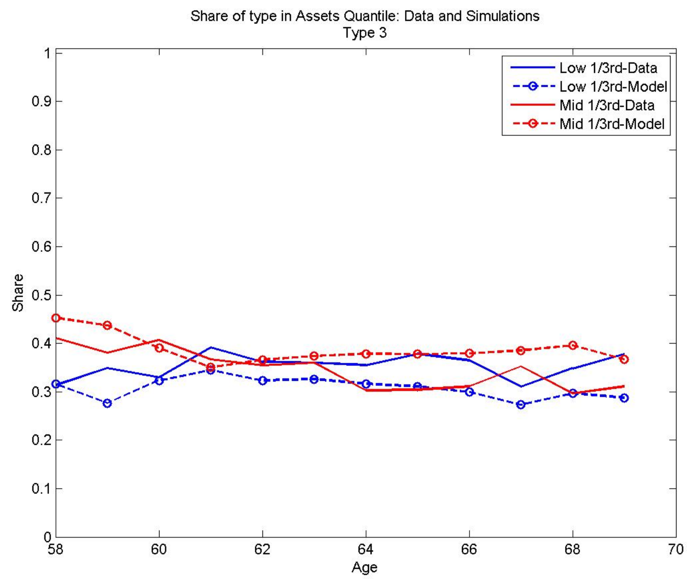

(c) Type 3 - High Preference for Own Leisure, High Preference for Spousal Leisure

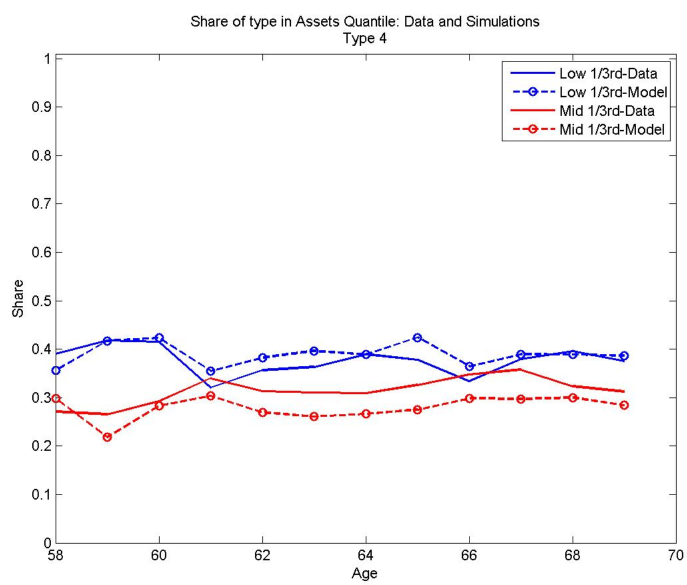

(d) Type 4 - Low Preference for Own Leisure, High Preference for Spousal Leisure

Figure K.3: Men Labor Force Participation by Preference Type

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

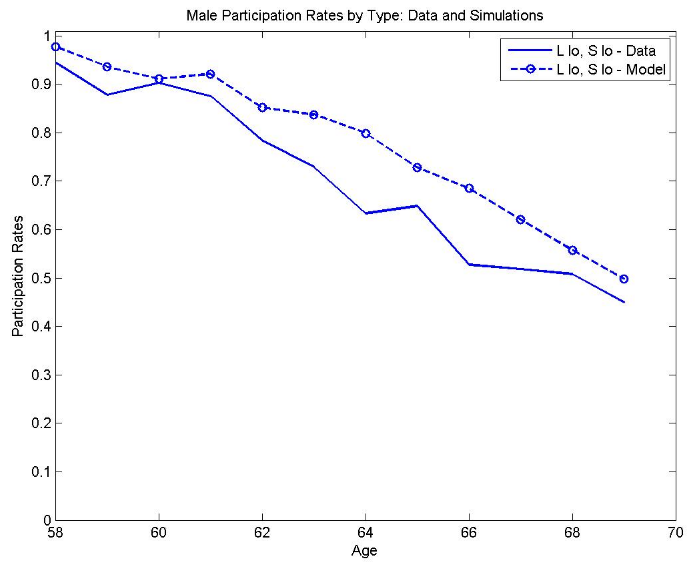

(b) Type 3 - High Preference for Own Leisure, High Preference for Spousal Leisure

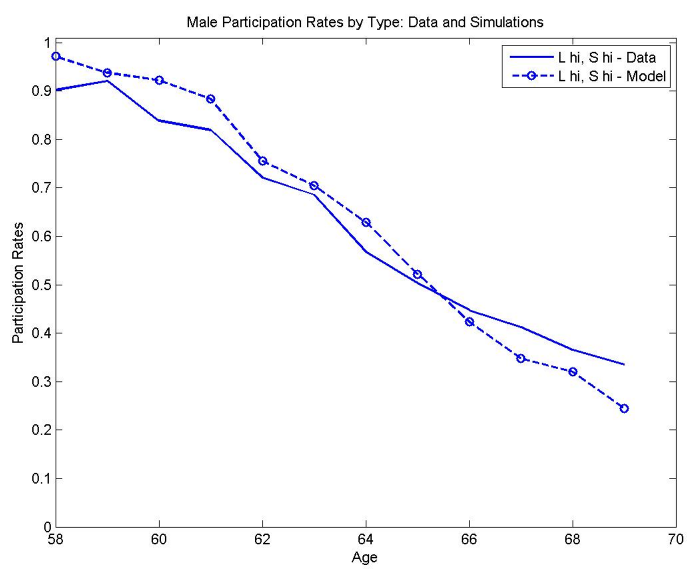

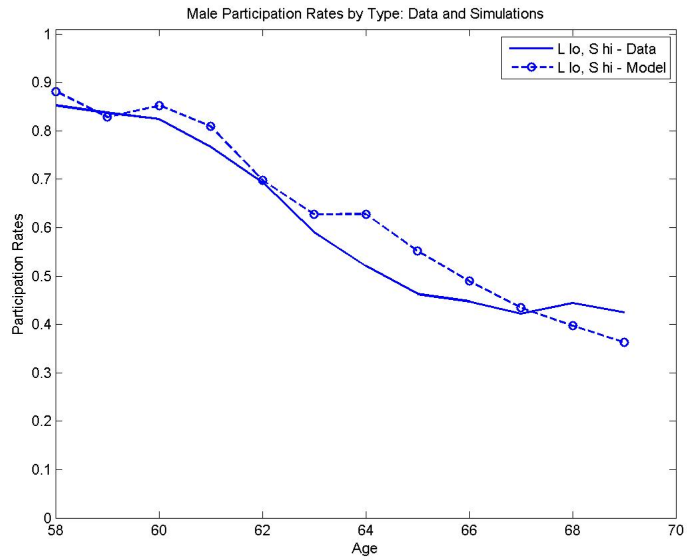

(c) Type 4 - Low Preference for Own Leisure, High Preference for Spousal Leisure

Figure K.4: Women Labor Force Participation by Preference Type

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

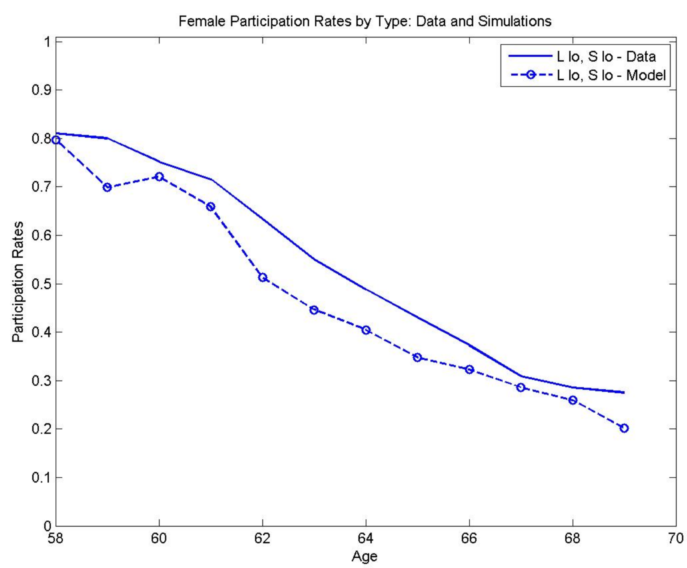

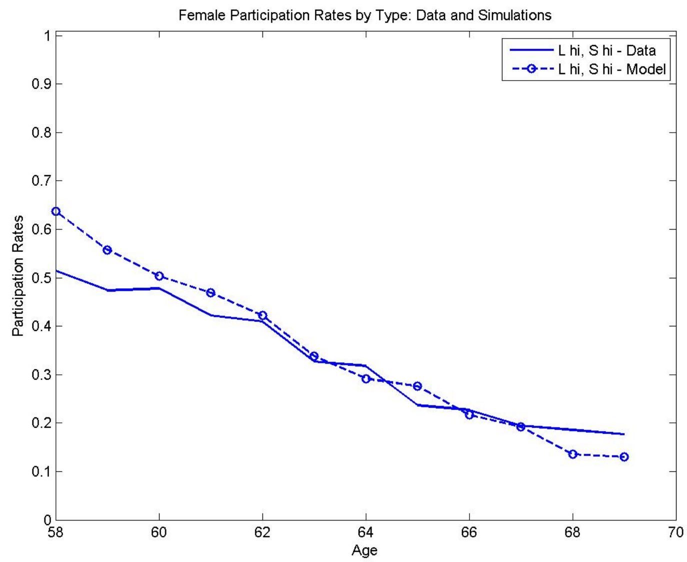

(b) Type 3 - High Preference for Own Leisure, High Preference for Spousal Leisure

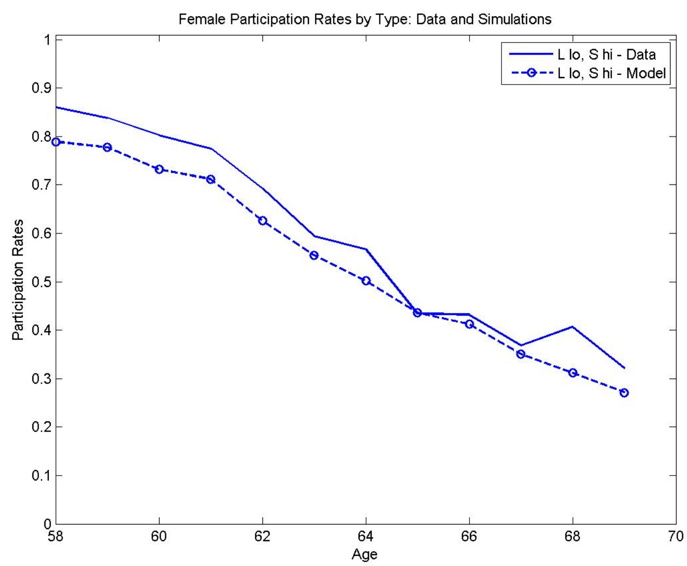

(c) Type 4 - Low Preference for Own Leisure, High Preference for Spousal Leisure

Figure K.5: Men Full-time work by Preference Type

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

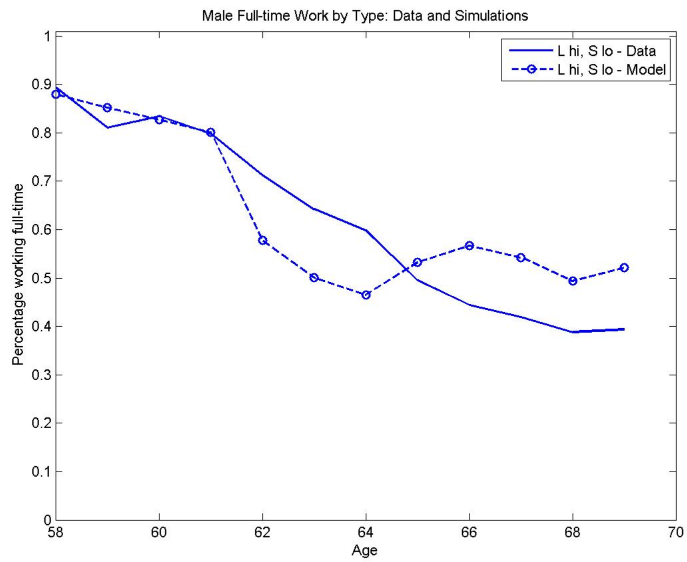

(b) Type 2 - Low Preference for Own Leisure, Low Preference for Spousal Leisure

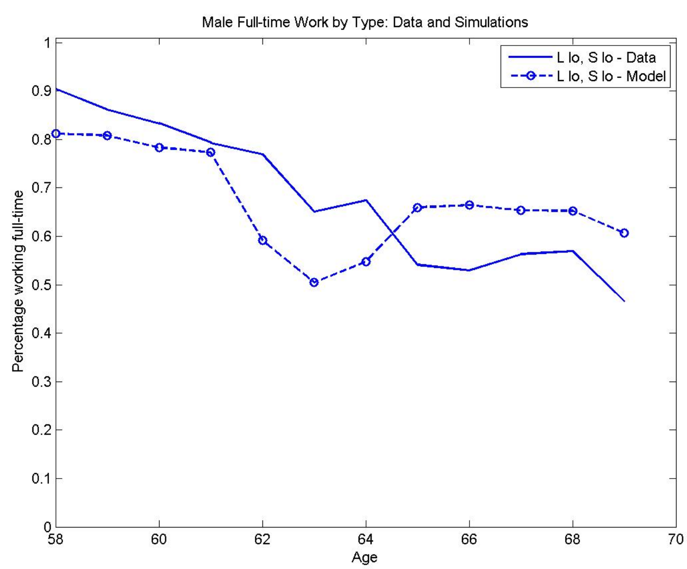

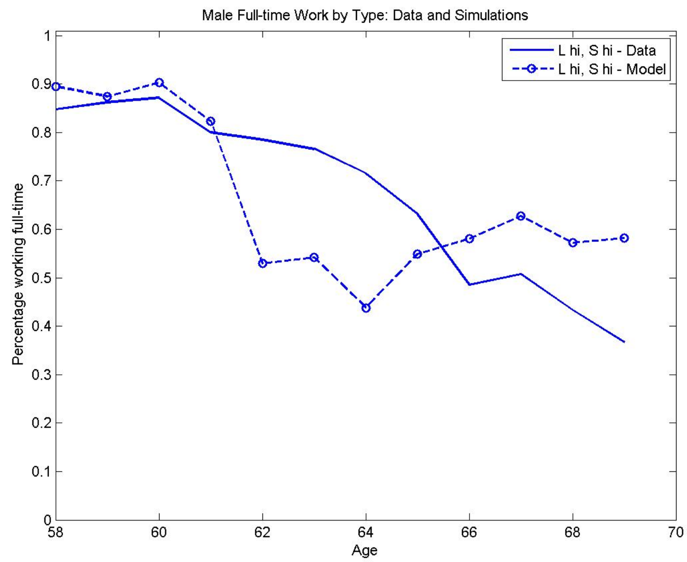

(c) Type 3 - High Preference for Own Leisure, High Preference for Spousal Leisure

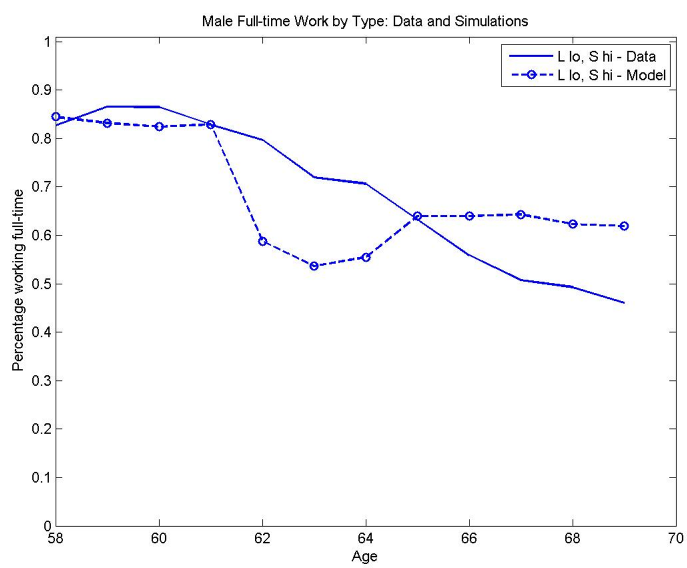

(d) Type 4 - Low Preference for Own Leisure, High Preference for Spousal Leisure

Figure K.6: Women Full-time work by Preference Type

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

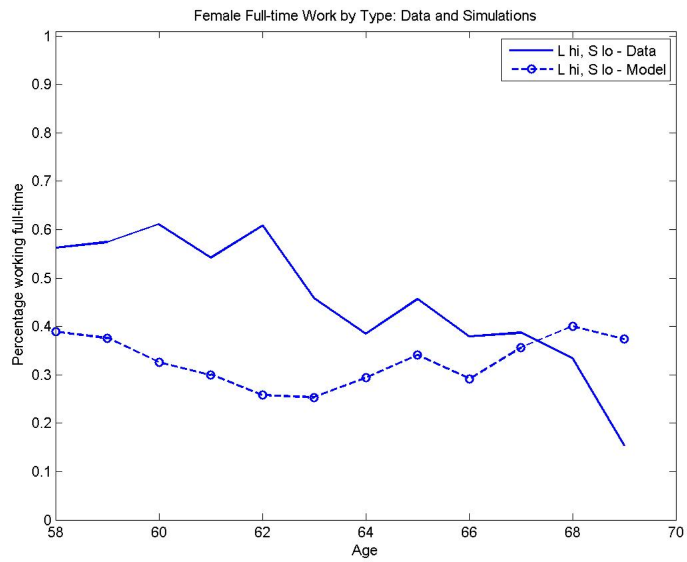

(b) Type 2 - Low Preference for Own Leisure, Low Preference for Spousal Leisure

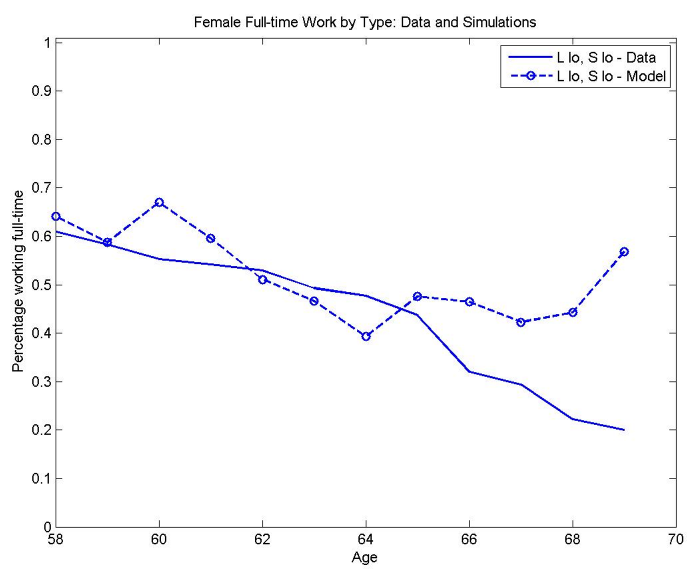

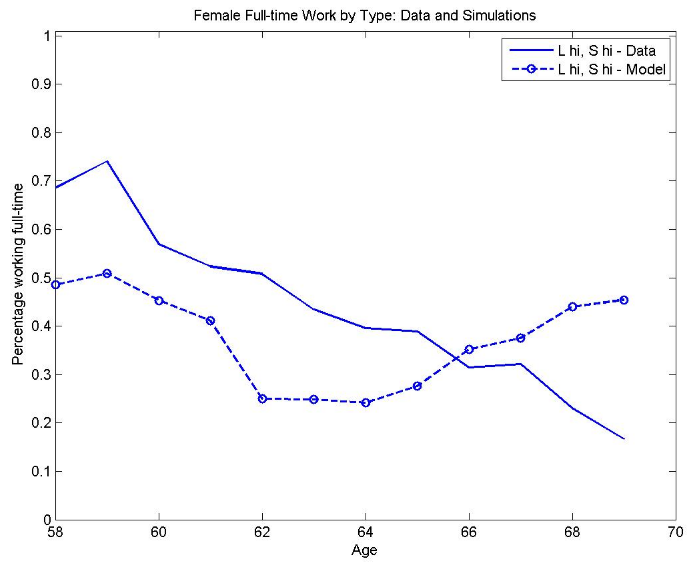

(c) Type 3 - High Preference for Own Leisure, High Preference for Spousal Leisure

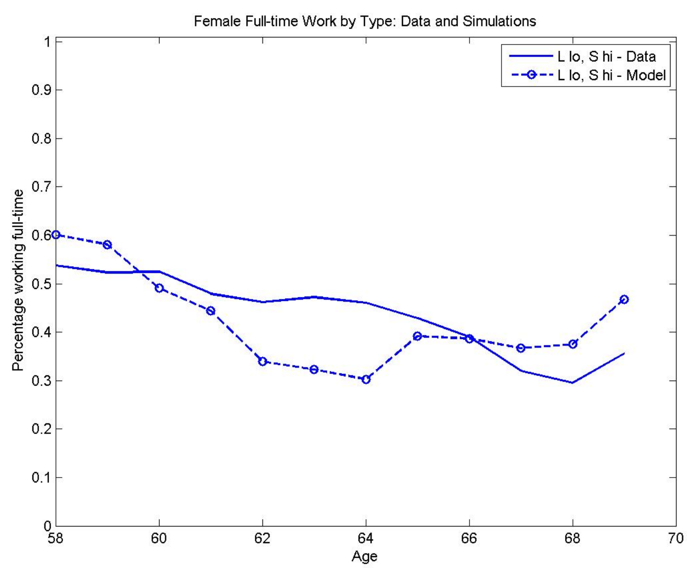

(d) Type 4 - Low Preference for Own Leisure, High Preference for Spousal Leisure

Figure K.7: Participation Rate by Health Status

(a) Men

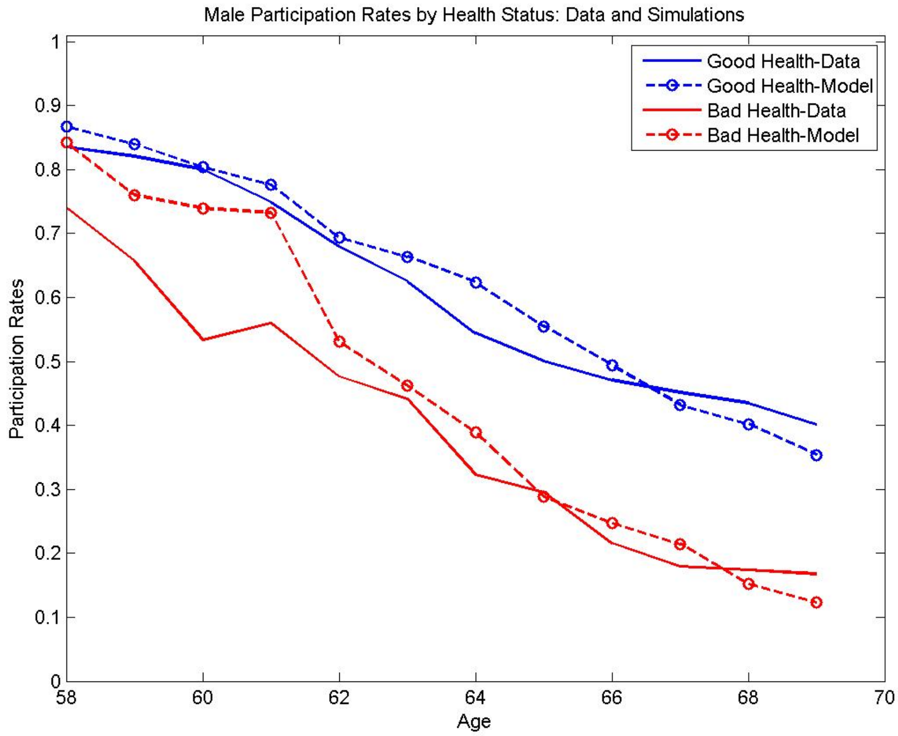

(b) Women

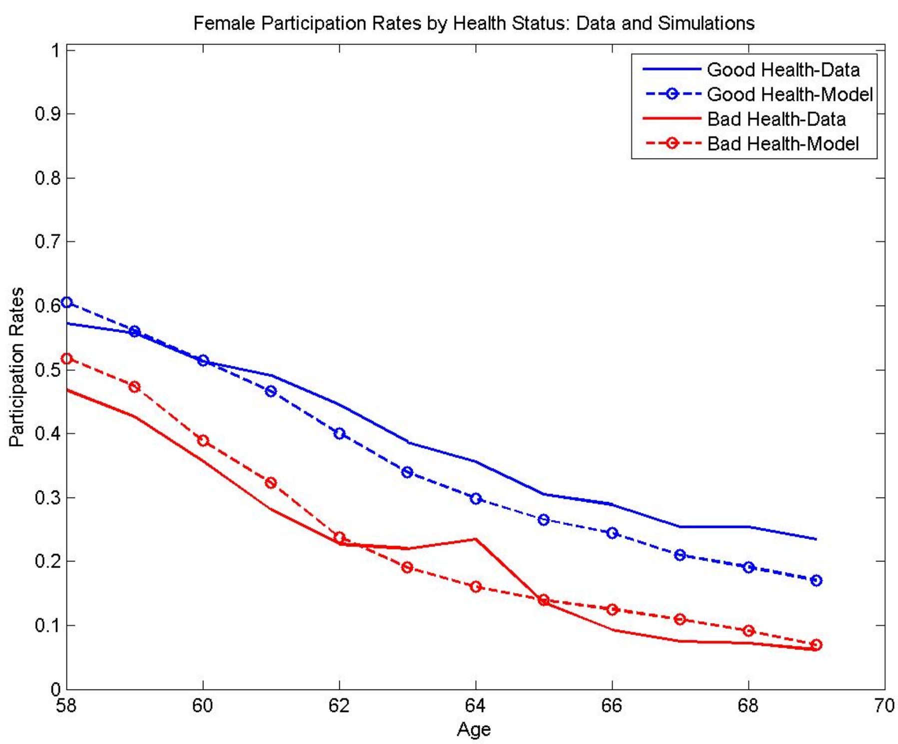

L Calculating Divorce Rates by Cohort

In the Health and Retirement Study (HRS), a respondent is asked for a retrospective marital history, that includes questions on how many times he or she has been married, and when his or her marriages began and ended. Individuals are only asked retrospective marital histories in the first wave that their cohort enters the HRS. Individuals born between 1931-1941 are part of the original HRS cohort (data from 1992-2010), individuals born between 1924-1930 are part of the Children of the Depression cohort (1998-2010), individuals born between 1942-1947 are part of the War Baby cohort (1998-2010) and individuals born between 1948-1953 are part of the Early Baby Boomer cohort (2004-2010). From this information I can approximate the individual's age at divorce for marriages prior to the respondent's participation in the HRS. Furthermore, since the HRS is a longitudinal study, I can track whether or not the individual has become divorced since the HRS began.

Divorces can be observed either through retrospective marital histories or during the survey waves of the HRS. In both cases, so that the proportion of divorces can be viewed as a proportion of those households which are eligible to get divorced, it is important to determine what share of the population is at-risk for the divorce . I restrict my observations to age ranges (i.e. ages 41-50, 51-60, 61-70) and by gender, as individuals' labor supply behavior appears to vary along these dimensions.

Using the retrospective marital history, I create an indicator of whether an individual is at-risk for divorce (i.e. has been married during the restricted age range) which takes a value of 1 if (1) the individual is divorced or widowed within the age range, (2) the individual is married within that age range, (3) the individual was married before and divorced after the age range, or (4) the individual's current marriage pre-existed the age range. Using the retrospective marital histories, I construct a dummy variable for divorce if it falls within the age range of interest. When calculating years married, I use the first divorce in the age range (i.e. if I am examining the age range from 41-50, and an individual is divorced at 41 and again at 49, then I use the length of marriage for the divorce at age 41).

Using the survey waves of the HRS, an individual within the age range of interest is considered at-risk if (1) the individual is currently married, (2) the individual was married since the previous interview wave, or (3) the individual was divorced or widowed since last interview. In the same manner as retrospective histories, I construct a binary variable for divorce within the age range of interest and calculate marital length based on the first divorce within the survey.

M Age Standardization

The statistics reported in tables 9.1 and 9.2 report statistics for the continuously married, divorced, and widowed populations in the HRS who are at-risk for divorce or widowhood between ages 50 and 70. Death, however, is an event that is more likely to happen later in a couple's life, while divorce is an event that is more likely to happen earlier in a couple's life. In tables M.1 and M.2, I standardize the same statistics by age. In order to do this, the data are weighted based on the age distribution of divorcées.

Comparing table 9.1 in the main text, to the weighted and age-standardized table M.1, we observe several differences. First, before the age-standardization, widowed and divorced women appear to have worked the same length on average, but afterward, divorced women appear to have have worked longer by 2.89 years. Age differences with the spouse appear to be much less for continuously married women and divorced women, with divorced women being only 0.18 years younger than their ex-husband. Divorced women also claim their Social Security benefits later than the other groups, once age is controlled for.

Controlling for age results in some substantial changes for men. Continuously married men and divorced men have more education on average, 13.3 years, once we control for age-differences. Other major changes include: widowed men's average years worked is 2.0-2.3 years less than that other groups, remarriage rates increase to nearly 50%, and continuously married and divorced men are much more likely to have worked the previous survey wave. Widowed men's salaries increase substantially once ages are standardized, but this is likely the result of smaller sample sizes amongst widowers.

Comparing table 9.2 in the main text, to the weighted and age-standardized table M.2, we observe a few differences. The difference between the percentage of divorcées and widow(er)s that are working before and after the separation/loss shrink from about 20 to 10 percentage points once ages are standardized. Additionally, the increase in women's expectation of working at 62 and 65 following a separation rises by 2.03 and 2.51 percentage points, respectively.

Table M.1: Summary Statistics from HRS Sample

	Women			Men		
	Population	Divorced	Widowed	Population	Divorced	Widowed
Age (Mean)	53.84	53.81	53.83	54.18	54.19	54.21
Before Separation		56.67	55.57		57.76	58.24
White (%)	0.89	0.92	0.76	0.89	0.90	0.86
Education in Years	12.95	12.83	12.22	13.32	13.31	11.43
Average years worked	23.29	25.67	22.78	33.15	32.86	30.85
Before Separation		27.31	24.53		34.45	33.03
Marriage Length (Mean)	26.91	19.10	28.16	24.78	17.61	24.48
Before Separation		20.25	28.55		18.75	31.06
Remarried after 3 waves (%)		0.31	0.20		0.49	0.52
Age difference with spouse	1.36	0.18	4.59	-3.51	-5.55	-2.00
Number of Children	2.62	2.55	3.00	2.56	2.50	2.34
Worked last survey wave	0.79	0.86	0.70	0.92	0.92	0.80
Before Separation		0.77	0.60		0.80	0.75
Worked in last five years, but not last survey wave	0.09	0.08	0.19	0.05	0.05	0.16
Before Separation		0.11	0.13		0.10	0.19
Not worked in last five years	0.12	0.07	0.12	0.03	0.03	0.04
Before Separation		0.12	0.25		0.10	0.06
Hourly Wage	21.71	18.69	17.04	36.70	29.54	56.56
Before Separation		21.39	16.21		26.90	20.44
Ever Applied Disability (%)	0.13	0.25	0.29	0.15	0.23	0.19
Age begin Social Security	62.46	62.55	61.28	63.06	63.35	62.53
Sample Size	8409	235	239	8475	237	102

Notes: Age's standardized by the age distribution of divorcées. Data weighted using person level weights. Separations occurring between age 50 and 70. Unless otherwise indicated, statistics refer to the first wave.

Table M.2: Before and After Separation - Summary Statistics from HRS Sample

		Women			Men		
		Population	Divorced	Widowed	Population	Divorced	Widowed
	Number in Household	2.80	2.54	2.89	3.02	2.92	3.17
	Before Separation		2.48	2.60		2.73	2.64
	After Separation		1.68	2.05		1.70	2.01
Excluding Housing	Assets ($100,000 - Mean)	2.84	2.81	1.65	2.98	2.87	0.89
	Before Separation		3.21	1.58		3.01	1.28
	After Separation		1.61	1.66		1.67	2.26
Including Housing	Assets ($100,000 - Mean)	4.11	3.78	2.47	4.13	4.12	1.71
	Before Separation		4.42	2.55		4.30	1.81
	After Separation		2.31	2.90		2.56	3.16
	Renting (%)	0.09	0.11	0.14	0.11	0.12	0.15
	Before Separation		0.07	0.14		0.11	0.17
	After Separation		0.40	0.18		0.38	0.16
	1st wv. After Separation		0.36	0.17		0.39	0.26
	Working (%)	0.68	0.73	0.52	0.86	0.88	0.70
	Before Separation		0.62	0.51		0.76	0.63
	After Separation		0.65	0.56		0.73	0.60
	1st wv. After Separation		0.64	0.52		0.65	0.53
	Applied for Disability (%)	0.02	0.04	0.04	0.04	0.06	0.05
	Before Separation		0.06	0.10		0.11	0.07
	Ever Apply	0.13	0.25	0.29	0.15	0.23	0.19
	Prb. Work at 62 (Self-rpt)	36.18	45.23	45.99	51.70	58.13	56.90
	Before Separation		48.60	36.21		58.69	48.89
	After Separation		54.18	32.91		51.19	42.36
	Prb. Work at 65 (Self-rpt)	20.29	26.72	23.55	31.79	30.59	31.21
	Before Separation		26.54	20.38		35.10	25.79
	After Separation		37.29	18.29		36.12	20.67
	Sample Size	8409	235	239	8475	237	102

Notes: Age's standardized by the age distribution of divorcées. Data weighted using person level weights. Separations occurring between age 50 and 70. Unless otherwise indicated, statistics refer to the first wave. Separations occur between waves. The interview wave response before the separation is marked as "Before Separation", and similarly for the response "After Separation". "1st wv. after separation" refers to the interview response in the first full wave since the separation occurred (minimum of 2 years). All assets are reported in 2010 dollars.

N Changes in Household Size around Divorce/Widowhood

When considering old-age divorce/widowhood, we typically think of a household transitioning from a two-person household to two one-person households. For the median household, this is correct (see table N.1). For some households, they may also have a child or parent living at home. This can explain why the average household size is greater than 1 in table N.1.

The Health and Retirement Study collects information regarding the number of other household members at each interview wave. The HRS consistently asks whether children are "at home or temporarily away at college" versus not living at home. It is often unclear, however, what would happen to a divorcing household with shared custody of children, and who comprise the other, non-child, household members. Controlling for the number of children "at home or temporarily away at college" substantially reduces the average number of people in a household. Accounting for the remaining household members is difficult, as the HRS does not consistently track the relationship between the other, non-children members of the household and the respondent.[63]

We might be concerned that household composition is biasing the results we find in this paper. In table N.1, I examine couples that have no other people in the household in the first wave, and before/after the separation or loss. Relative to the full sample in table 9.2, we find that these couples are generally wealthier. Comparing the two samples, we observe that all of the key relationships are preserved. Women receive less than 50% of pre-divorce assets, and widows on average slightly lose assets.

The median household size does not change for widowers following the wife's death. Only three households have remarried or become partnered with another person by the interview wave after the wife's death. Controlling for reported children at home does result in the expected change in median household size. It is possible that widowers are more likely to have their children return home following the wife's death, or to move in with their children.

While assets including housing wealth are stable around the widowing event, the average non-housing assets fall sharply. While non-housing assets decreased for widowers in the entire sample (from \$153,000 to \$145,000), these assets decreased by a much larger amount, from \$256,000 to \$175,000, in the sample of widowers who went from two to one person households. This sharp change is due to large asset changes by two households with substantial assets

[63] It is possible, for example, that divorcing men move in with a girlfriend and the household might now reflect both of their assets, but this scenario would require the man to report being divorced or separatednot, instead of reporting being partnered or married. The HRS interviewer typically asks whether the respondent, *or his/her spouse or partner*, owns any of a particular asset. Therefore it is unlikely that asset measures would be confounded by post-marital relationships.

(greater than $1 million). Dropping these two households from the analysis, we observe a slight increase in assets. There is no evidence to support that asset loss is common following the death of a wife. If anything, there appears to be a median increase in assets, which is suggestive that many of these widowers may receive a payout from a life insurance policy.

Table N.1: Before and After Separation Summary Statistics from HRS Sample (controlling for household size)

		Women			Men		
		Population	Divorced	Widowed	Population	Divorced	Widowed
	Household Size (Median)	2	3	2	2	3	3
	Before Separation		2	2		2	2
	After Separation		1	1		1	2
	Household Size (Mean)	2.89	2.95	2.95	2.89	2.99	3.12
	Before Separation		2.70	2.56		2.76	2.72
	After Separation		1.82	1.95		1.70	1.99
Excluding Housing	Assets ($100,000 - Mean)	2.82	1.97	1.88	2.94	3.32	1.39
	Before Separation		3.26	1.82		3.46	2.56
	After Separation		1.33	1.74		1.74	1.75
	Assets ($100,000 - Median)	1.21	0.97	0.86	1.26	1.51	0.77
	Before Separation		1.35	0.50		1.36	0.51
	After Separation		0.48	0.59		0.71	1.04
Including Housing	Assets ($100,000 - Mean)	3.86	2.95	2.93	3.95	4.44	2.25
	Before Separation		4.27	3.02		4.43	2.77
	After Separation		1.92	2.83		2.60	2.90
	Assets ($100,000 - Median)	2.36	1.71	1.77	2.42	2.66	1.64
	Before Separation		2.44	2.24		2.57	1.48
	After Separation		0.75	1.54		1.30	2.54
	Sample Size	7564	212	233	7609	221	98
	Wave 1 Sample Size	3796	72	88	3813	74	28
	Before/After Sample Size		87	100		82	29

Notes: Author's calculations, Data from Health and Retirement Study. Separations occurring between age 50 and 70. Unless otherwise indicated, statistics refer to the first wave. Separations occur between waves. The interview wave response before the separation is marked as "Before Separation", and similarly for the response "After Separation". All assets are reported in 2010 dollars.

Household size statistics reported using full sample size noted in table 9.2. The full sample size may differ from main text because the household size could not be established for some households. Wave 1 sample size reports households who either never divorce between 50-70, divorce, or widow, and were 2 person households in wave 1. Before/After sample size reports the sample size of households where the household had 2 people before and 1 person after the separation.

Asset statistics referring to the first wave are dependent on the household being comprised of two individuals (Wave 1 Sample Size). Asset statistics referring to before or after separation are dependent on the household being comprised of two individuals before separation and only one individual after separation (Before/After Sample Size).

O Impact of marital separation/loss on retirement benefits

When a household divorces, or one member dies, the change may affect its employer-provided retirement benefits or the household's Social Security benefit. Employee's are generally eligible for two types of retirement benefits: defined benefit (DB), defined contribution (DC) or a combination plan that consists of both. Table O.1 provides a summary of the impact of marital separation/loss on different retirement benefit types.

Table O.1: Retirement benefit options for divorcées and widow(er)s

	Defined Benefit (DB)	Defined Contribution (DC)	U.S. Social Security (SS)
Divorce before retirement	Separation of Accounts Time Rule Formula Flat Dollar/Percentage	Separation of Accounts Rollover to IRA Flat Dollar/Percentage	Better of {50% ex-spouse's benefit, 100% own benefit} if married 10+ yrs
Divorce after retirement	Time Rule Formula Flat Dollar/Percentage	Separation of Accounts Rollover to IRA Flat Dollar/Percentage	
Widowhood before retirement	Lump Sum Payment Survivor Share of Joint Life Annuity	Assumes account ownership	Better of {100% spouse's benefit, 100% own benefit} & early claim age reduced to 60
Widowhood after retirement	Single Life Annuity Joint Life Annuity Certain Period Annuity Cash Refund Annuity	Assumes account ownership	

Notes: Table provides a summary of changes to retirement plans due to a marital separation/loss. A person is retired in this context if they have left an employer and claimed the DB/DC benefit provided by that employer. SS is claimed separately from the employer plans. Terms are defined in the text. This summary is not an exhaustive list, as the options vary by plan, state, and whether it is a public or private employer.

O.1 Laws governing employer-sponsored retirement plans

The Employee Retirement Income Security Act of 1974 (ERISA) establishes rules for employer-sponsored retirement plans.[64] Two important provisions for this paper include the establishment of spousal consent and qualified domestic relation orders.

ERISA requires that the default payout for DB plans includes a surviving spouse provision. This provision required that a beneficiary's surviving spouse must receive at least half of the monthly benefit received while the beneficiary was still alive. The spouse may consent to receiving less than half, but this requires her to provide written confirmation of her intent to forgo survivorship rights.

ERISA also provides for the creation of qualified domestic relation orders (QDRO). QDROs are court orders that recognize the rights of an alternative payee to receive all or part of the benefits from another person's retirement plan. QDROs are restricted to an individual's spouse, former spouse, child, or other dependent. They are separate from divorce orders, and pension plans will only make benefit changes based on QDROs. The retirement plan administrator determines whether an order is a QDRO. By law, a QDRO cannot require a plan to do anything beyond what the plan's administration permits. It is common for plans to provide sample QDRO language that delineate the options the plan makes available in the case of a divorce.

ERISA only applies to private plans. Federal, state and local government entities, such as state public employee retirement systems, establish similar rules for their plan. State pension plans are not required to establish a spousal consent provision.[65] Most state pension plans have established orders similar to QDROs, often referred to as "division of property order" (Ohio), "eligible domestic relations order" (Michigan), or simply a "domestic relations order" (New York).

O.2 Defined Benefit Plans

DB plans guarantee the retiring worker a monthly benefit based on a formula that accounts for years of service, average or final pay, and the age at which the benefit begins. These plans are typically referred to as pension plans, and are common in public sector jobs. The employee who is eligible for the pension is considered a plan member, while his or her

[64] It was subsequently amended in 1984 by the Retirement Equity Act. I refer to the act, including the amendments, as ERISA.

[65] "Nineteen states and the District of Columbia do not require their retirees to notify their spouses or get their spouse's consent when they elect a single-life annuity: Alabama, the District of Columbia, Colorado, Georgia, Indiana, Kentucky, Maryland, Mississippi, Montana, Nebraska, New York, North Carolina, North Dakota, Pennsylvania, Rhode Island, South Carolina, Tennessee, Utah, Vermont, and West Virginia." (Pension Rights Center, 2013)

spouse would be considered a nonmember. Members and nonmembers are only entitled to an annuity, or lifetime, benefit if the member is vested, meaning they have worked for minimal number of years at the employer. Non-vested members are generally able to withdraw their contributions and, depending on the plan, part of the employer's contribution. DB plan administrators establish different rules for before and after an individual begins receiving monthly benefits. Three common methods of dividing marital retirement assets include:

Separation of Accounts. This provision can generally only occur prior to the first receipt of benefits. The California Public Employees' Retirement System defines separation of accounts in the following way.

> The nonmember spouse no longer relies upon the member to determine when benefits become payable. The nonmember spouse may choose to receive a refund of the accumulated contributions while the member continues to work. Should the nonmember elect to refund, they have the right to withdraw, by direct refund or rollover, the contributions and interest credited to their nonmember account, plus interest earned at 6 percent per year through date of payment. The taxable portion of the benefit would be subject to 20 percent federal withholding, unless it is rolled over to an IRA. If the member was not vested on the date of dissolution or dies prior to reaching the minimum retirement age, the nonmember's only right would be to withdraw their contributions by a direct refund or rollover. (California Public Employees' Retirement System, 2013)

Specific provisions for separated accounts vary by plan and state. Employers do not have to provide separate accounts for nonmembers, although this option is common to prevent the pension fund from needing to make lump sum disbursements to nonmembers following divorces or other life events.

Time Rule Formula. This provision occurs before or after the plan's beneficiaries first receipt of benefits. The time rule formula (referred to as the Majauskas Formula in New York) is the most common across states. Benefits to the nonmember spouse are calculated according to the following formula:

$$\text{Nonmember Benefit} = 50\% \times \frac{\begin{array}{c}\text{Total service credits earned}\\ \text{while married and employed}\end{array}}{\begin{array}{c}\text{Total service credits}\\ \text{at retirement}\end{array}} \times \text{Retirement Benefit}$$

The time rule formula need not be the 50% specified above. The actual percentage is agreed to as part of the QDRO. The domestic relation order will also provide for what happens to the benefit when either the member or ex-spouse dies. For example, some plans include a "pop-up" provision that converts the benefit payment to a single life annuity when one of the members dies. In some private plans, the ex-spouse may have the option of receiving their benefits as a lump-sum payment. Plans, however, are not required to offer this provision.[66]

Flat Dollar/Percentage. This provision occurs before or after the plan's beneficiaries first receipt of benefits. As part of the QDRO, the ex-spouse may receive a flat dollar amount or a percentage of the final monthly retirement benefit. For example, an ex-wife might settle for $1,500 of her ex-husband's future retirement benefit. A flat dollar amount is rare, as it would not include the ex-spouse in any COLA adjustments that might be part of the plan.

Should a couple remain married, several common options exist for ways to receive the benefits from a DB plan. These options specify whether the annuity payment includes a survivorship option, and what share of the original benefit the survivor would be entitled to. Benefit options include:

Lump Sum Payment. The nonmember may choose to receive all of the benefits in a single, lump sum payment. This is common in survivorship cases when the dying spouse would not have qualified for a retirement benefit at the age of their death.

Single Life Annuity (SLA). A single life annuity provides a retirement benefit based on the member's years of service, total service credits, their average or final pay, and the age they begin their benefit. The annuity payments last only for the duration of the member's life, and their are no benefits for the survivor. Since the passage of ERISA, choosing a single life annuity requires the written consent of the spouse. In this case, the survivor would receive no benefits.

Joint Life and Survivor Annuity (JLSA). This annuity is the most common form of benefit since the passage of ERISA. A JLSA provides a benefit to the member while they are alive and to a nonmember (e.g. wife) after the member dies. A 100% JLSA would provide a benefit of equal size to a surviving nonmember after the death of the member. A 50% JLSA would provide a benefit of half the size of the original benefit to a surviving nonmember after the death of the member. The nonmember beneficiary generally cannot be changed after the initial assignment since benefit levels are based on the age of the member and nonmember.

Certain Period Annuity. This annuity provides a guaranteed income for a certain

[66] In a legal case involving Continental Airlines, nine pilots were sued by the airline after they divorced and received QDROs assigning their ex-spouses complete pension rights. The ex-spouses proceeded to cash out the retirement plans, costing the airline $10-11 million. The pilots subsequently remarried their ex-spouses. The airline was entering bankruptcy, and it was likely that their pension benefits would have been reduced in the bankruptcy proceedings. In 2011, the 5th U.S. Circuit Court of Appeals upheld a lower-court ruling that employers are prohibited from considering or investigating why their employees get divorced or whether their divorce is fraudulent.

duration (e.g. 5, 10, 20 years), and specifies a beneficiary in case the recipient does not survive the entire certain period.

Cash Refund Annuity. This annuity guarantees that the member will receive total annuity payments equal to his or her contributions (alternatively, the plan's initial value). If the member dies before this occurs, the remaining balance is paid to a specified beneficiary (e.g. wife). The cash refund annuity can be used in combination with a SLA or a JLSA, but doing so reduce the member's monthly benefit level.

The benefits a nonmember beneficiary receives if the plan's member dies before collecting benefits will often depend on the member's years of service. If he has enough years of service, the nonmember beneficiary will often be able to collect a 100% JLSA. If the member did not have enough service credits, the nonmember beneficiary is often constrained to taking a lump sum distribution.

Divorcing couples where the plan's member has not begun collecting their benefit generally can choose from separation of accounts, time rule formula, or flat dollar/percentage payment. The method of division, however, may vary by employer, and so this is not meant to be a comprehensive list. Once the plan's member begins collecting his or her benefits, often the set of options is reduced. Once a payment plan is chosen, changes from a divorce cannot make the retirement plan's finances worse off. This would be grounds for a plan's administrator refusing to certify a court order as a QDRO. For example, if a household chose a SLA and subsequently divorces, the household cannot convert the SLA to a JLSA because that would alter the actuarial plan. Instead the household would be likely to opt for a flat percentage of the retirement benefit.

O.3 Defined Contribution Plans

DC plans, such as 401(k), 403(b), 457, or thrift plans, are voluntary contribution plans where the employer and employee can elect to contribute. Contributions to DC plans can be made pre-tax (i.e. traditional) or post-tax (i.e. Roth). Similar to DB plans, most DC plans have a vesting provision based on years of service if the employer is making contributions. Non-vested members are generally able to withdraw their contributions and, depending on the plan, part of the employer's contribution. Since DC plans have a specific balance, the rules associated with divorce/widowhood are simpler than DB plans.

In the case of divorce, the balance of a DC account can be split between the couple. In some cases the employer will allow the non-employee to remain on their plan, however, the non-employee is often required to rollover the balance to an individual retirement account (IRA). Similar to DB plans, it is also possible for a DC plan's assets to be split in flat dollar amounts or percentages at the end of the employee's service.

In the case of widowhood, DC plans have an assigned beneficiary who receives the balance when the member dies. The default beneficiary is the member's spouse. DC plans, like DB plans, are required under ERISA to provide a survivorship option. A nonmember spouse must provide written permission if the member specifies a person other than his or her spouse as the DC plan's primary beneficiary.

O.4 Social Security

Social Security provides benefits to an earner, and the earner's spouse, ex-spouse, and survivor. A person may collect Social Security benefits as early as age 62, but benefits are reduced when claimed before a person's normal retirement age (i.e. age 65 for individuals born before 1938, and incrementally increasing to 67 for people born after 1959). If an individual claims at 62, his benefit is reduced by 20%, if born before 1937, or up to 30%, if born after 1959.

While alive, a spouse is entitled to the better of her own earner benefit or 50% of her spouse's benefit. If the couple divorces, conditional on having been married for 10 or more years, she may still collect this benefit. In order to claim the spouse's/divorced spouse's benefit, the high income earner must first have claimed their benefit. If the couple has been divorced more than two years and the high earner is at least 62, the divorced spouse does not have to wait for the high earner to claim his benefit in order to begin collecting her divorced spouse's benefit. If the divorced spouse remarries, then she is no longer eligible to receive the divorced spouse's benefit (although she may be eligible to receive a spouse's benefit based on her new spouse's earnings history). If a currently divorced person has been married to multiple individuals for more than 10 years, she will receive the greater of the divorced spouses' benefits. The actuarial reduction for spouse and divorced spouse's benefits is greater than the earner's benefit reduction. If a spouse claims at 62, her spouse's benefit is reduced by 25%, if born before 1937, or up to 35%, if born after 1959. The spouse's benefit is based on the high earner's unreduced benefit (i.e. if he claimed at his normal retirement age).

When the high earner dies, the surviving spouse is entitled to the better of her own benefit or 100% of her spouse's benefit. A surviving spouse is entitled to collect their survivor's benefit as early as age 60. If the couple divorced prior to the high earner's death, the surviving ex-spouse is entitled to a divorced survivor's benefit. The divorced survivor's benefit is equivalent to the survivor's benefit. The rules governing whether an ex-spouse qualifies for a divorced survivor's benefit are similar to the rules governing who qualifies for a divorced spouse's benefit. The actuarial reduction for survivor and divorced survivor's benefits is generally lower than the earner's benefit reduction. If a survivor claims at 60, her survivor's benefit is reduced by 28.5%. The survivor's benefit reduction factor has not changed like the earner and spouse's benefits. The survivor's benefit is based on the dead earner's unreduced benefit if he had not claimed his benefit. If the dead earner had claimed his benefit, then the survivor's benefit would be based on the greater of the dead earner's reduced benefit or 82.5% of his unreduced benefit.

Social Security has a maximum family benefit that can be derived from a single earner's earnings history. This amount is typically around 150% of the earner's benefit level. Divorced spouse and survivor's benefits, however, do not count against there family maximum. For example, an earner could divorce a spouse every 10 years, immediately marry another spouse, and after 40 years have 4 divorced spouses' benefits, 1 spouse's benefit, and his own earner benefit (a possible total benefit of 350% of his normal retirement benefit). When he died, each would convert to a divorced survivor's benefit (a possible total benefit of 500% of his normal retirement benefit).

P Figures and Tables

Figure P.1: Women - Percentage whose most recent marital status is Widowed (by Age)

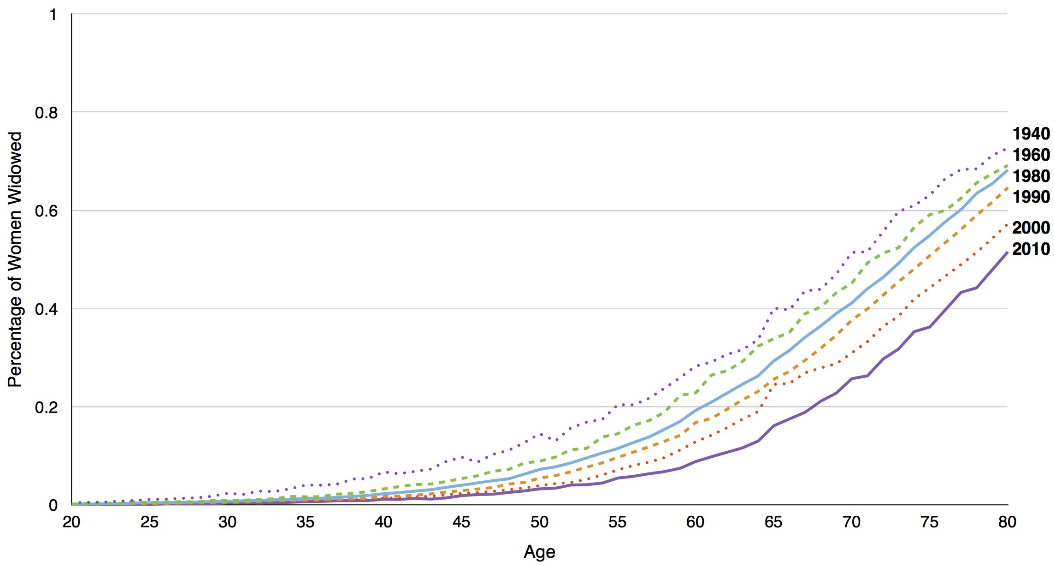

Notes: Author's calculations for 1940-2000 using U.S. Census (1% samples), 2010 using the American Community Survey.

Figure P.2: Men - Percentage whose most recent marital status is married (by age)

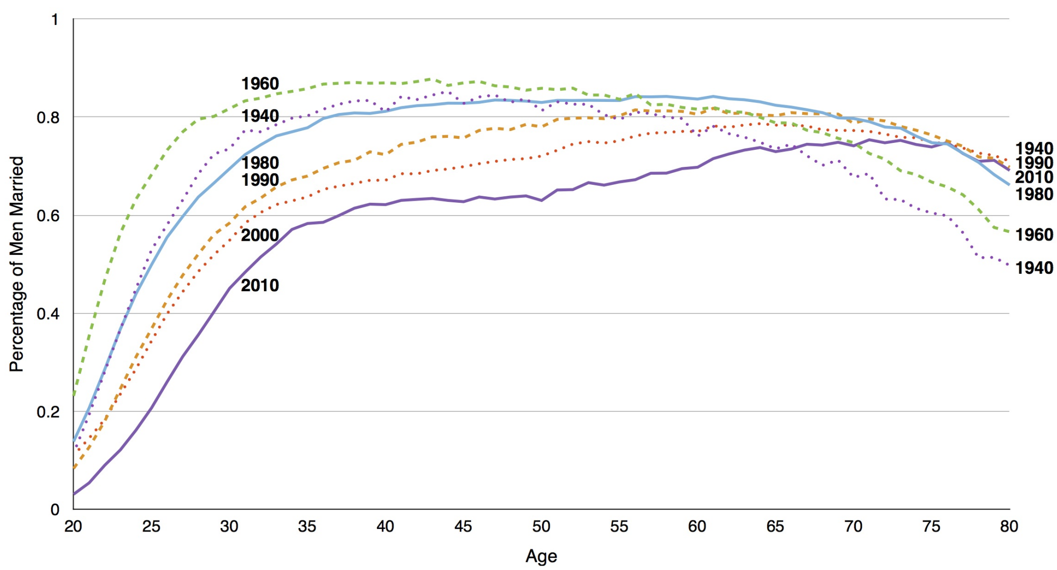

Notes: Author's calculations for 1940-2000 using U.S. Census (1% samples), 2010 using the American Community Survey.

Figure P.3: Men - Percentage whose most recent marital status is divorced (by age)

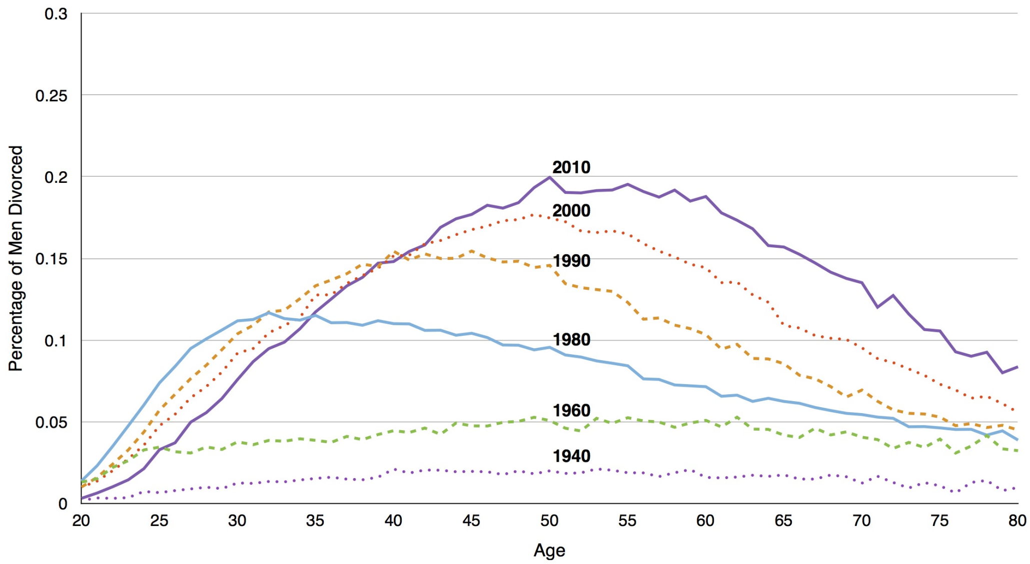

Notes: Author's calculations for 1940-2000 using U.S. Census (1% samples), 2010 using the American Community Survey.

Figure P.4: Men - Percentage whose most recent marital status is Widowed (by Age)

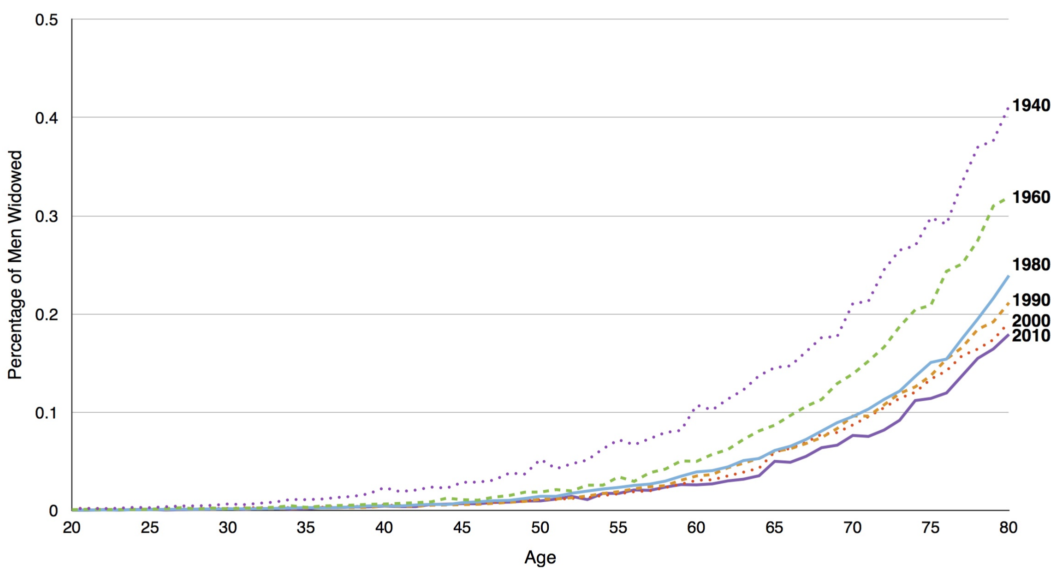

Notes: Author's calculations for 1940-2000 using U.S. Census (1% samples), 2010 using the American Community Survey.

Figure P.5: Effect of Widowhood timing on Women's Labor Force Participation

(a) Labor Force Participation, Ages 50-70

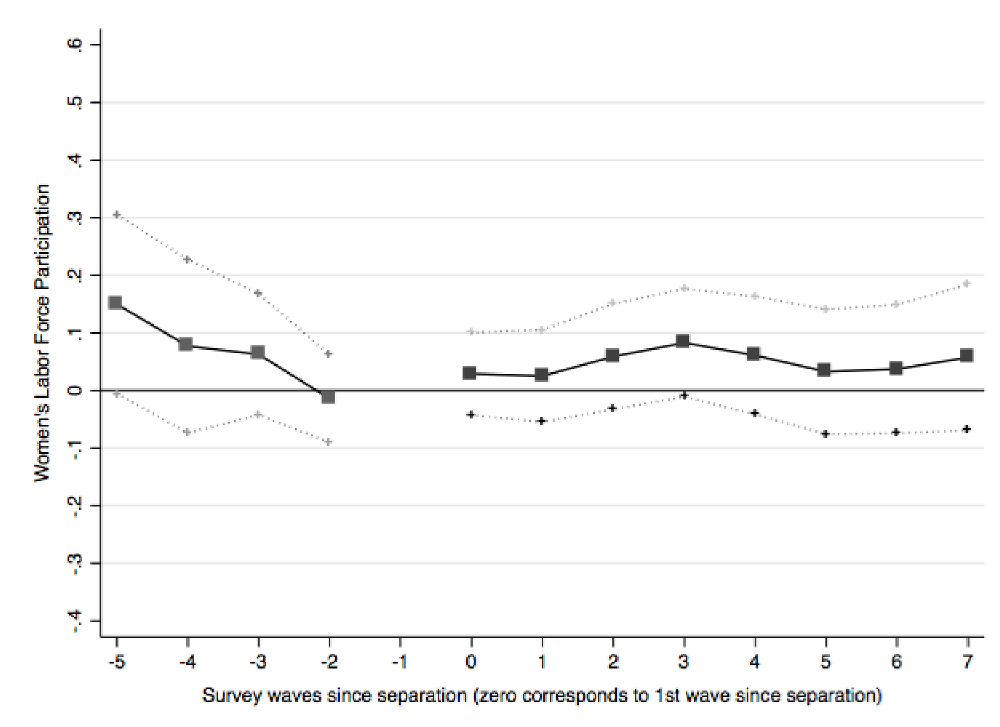

(b) Hours Worked, Ages 50-70

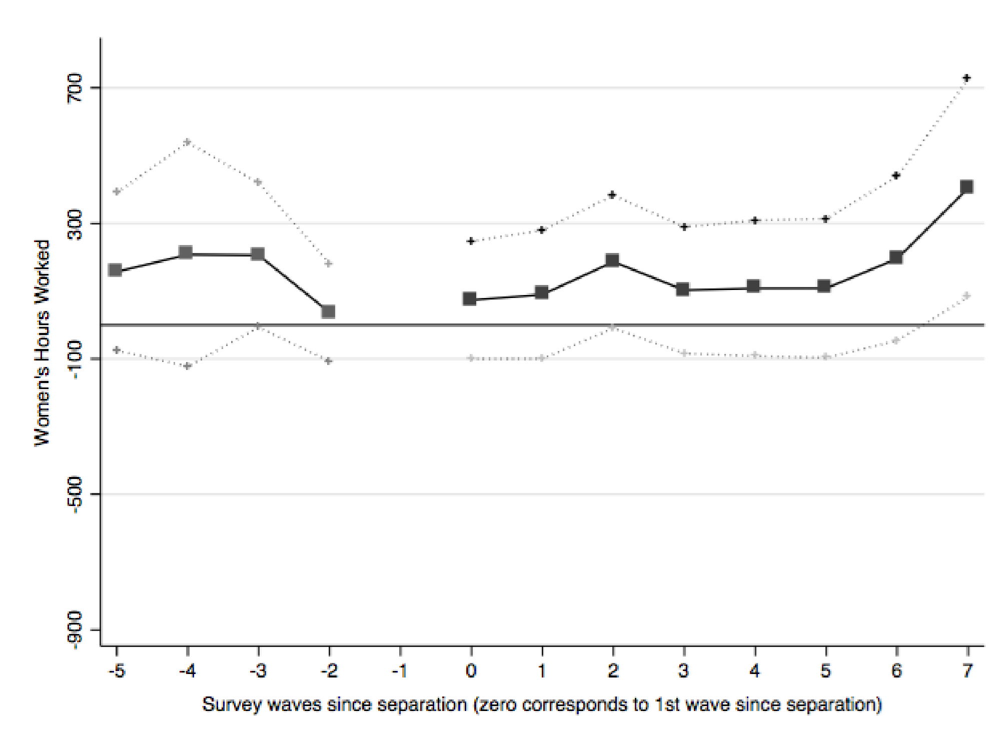

Figure P.6: Effect of Divorce timing on Men's Labor Force Participation

(a) Labor Force Participation, Ages 50-70

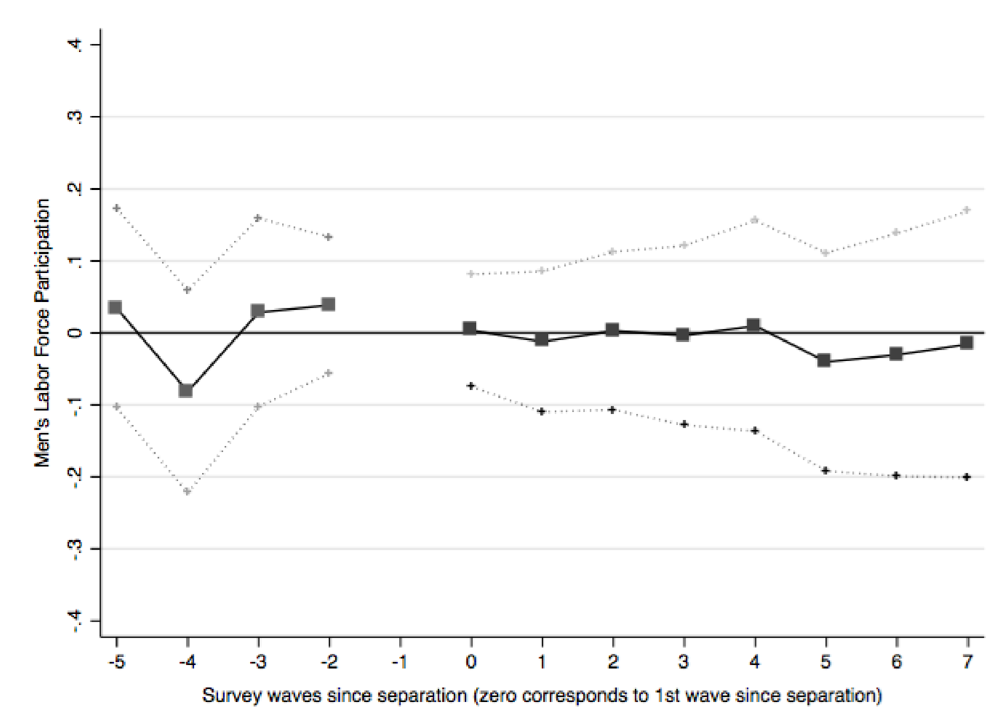

(b) Hours Worked, Ages 50-70

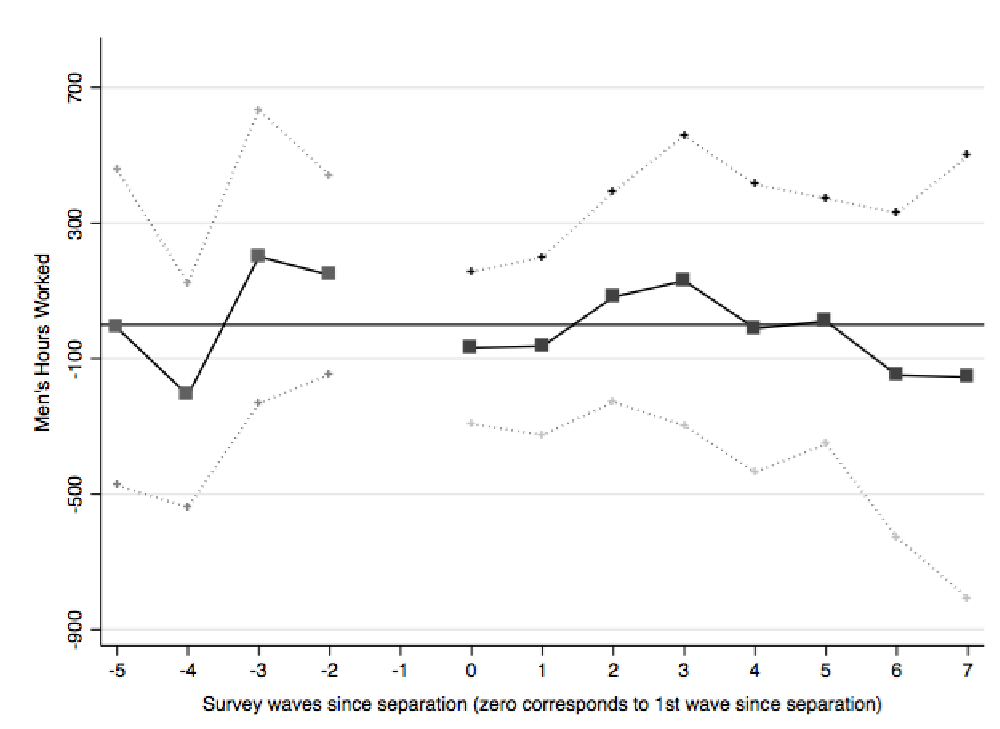

Table P.1: Comparisons of Estimates of Leisure Demand and Labor Force Participation of Women accounting for spousal leisure complementarities.

		Leisure Demand - Divorce						Labor Force Participation - Divorce					
		Original Model			Model with Joint Leisure			Original Model			Model with Joint Leisure		
		Dummy	× Asset Loss	× Sp. Loss	Dummy	× Asset Loss	× Sp. Loss	Dummy	× Asset Loss	× Sp. Loss	Dummy	× Asset Loss	× Sp. Loss
spouse works					-0.0117***						0.0699***		
					(0.00246)						(0.00824)		
spouse works full-time					-0.00734***						0.0241***		
					(0.00260)						(0.00890)		
After Separation (wave)	0th	0.00270	-0.00923***	0.000424	-0.00381	-0.00936***	0.000336	0.0802	0.0227**	-0.00175	0.114**	0.0233**	-0.00132
		(0.0128)	(0.00300)	(0.000414)	(0.0127)	(0.00310)	(0.000418)	(0.0532)	(0.0105)	(0.00126)	(0.0525)	(0.0109)	(0.00126)
	1st	-0.0172	-0.0130***	0.000656	-0.0224	-0.0133***	0.000571	0.0634	0.0272**	-0.000775	0.0908	0.0288**	-0.000350
		(0.0142)	(0.00355)	(0.000486)	(0.0142)	(0.00362)	(0.000489)	(0.0585)	(0.0122)	(0.00143)	(0.0590)	(0.0126)	(0.00144)
	2nd	-0.0338	-0.0288*	0.00127	-0.0378	-0.0292*	0.00118	0.114*	0.0405***	-0.00273	0.136**	0.0423***	-0.00226
		(0.0300)	(0.0159)	(0.000918)	(0.0301)	(0.0160)	(0.000922)	(0.0595)	(0.0126)	(0.00179)	(0.0601)	(0.0129)	(0.00179)
	3rd	-0.0363	-0.00394	0.00131*	-0.0394	-0.00429	0.00123*	0.0879	0.0283*	-0.00114	0.106	0.0300*	-0.000713
		(0.0315)	(0.00460)	(0.000691)	(0.0315)	(0.00462)	(0.000688)	(0.0709)	(0.0164)	(0.00215)	(0.0706)	(0.0166)	(0.00215)
	4th	0.0188	-0.00433*	-0.000963	0.0166	-0.00468*	-0.00105	-0.0517	0.0413**	0.000899	-0.0388	0.0430**	0.00130
		(0.0218)	(0.00247)	(0.000686)	(0.0219)	(0.00248)	(0.000687)	(0.0741)	(0.0171)	(0.00298)	(0.0747)	(0.0174)	(0.00300)
	5th	0.0109	-0.00567*	-0.000808	0.00944	-0.00599**	-0.000901	-0.00796	0.0280	0.00113	0.00130	0.0296	0.00158
		(0.0221)	(0.00301)	(0.000884)	(0.0222)	(0.00301)	(0.000877)	(0.0831)	(0.0223)	(0.00327)	(0.0837)	(0.0224)	(0.00324)
	6th	0.0476**	-0.0196**	0.00531**	0.0473**	-0.0200**	0.00522**	-0.0787	0.0429	-0.0193***	-0.0755	0.0445*	-0.0188***
		(0.0232)	(0.00852)	(0.00243)	(0.0231)	(0.00851)	(0.00243)	(0.0848)	(0.0263)	(0.00555)	(0.0840)	(0.0260)	(0.00549)
	7th	0.0568**	-0.0266***	-0.00705	0.0576**	-0.0269***	-0.00735	-0.0733	0.0789***	0.0511	-0.0755	0.0801***	0.0526
		(0.0255)	(0.00646)	(0.00501)	(0.0255)	(0.00651)	(0.00506)	(0.0898)	(0.0199)	(0.0338)	(0.0889)	(0.0199)	(0.0341)

		Leisure Demand - Divorce						Labor Force Participation - Divorce					
		Original Model			Model with Joint Leisure			Original Model			Model with Joint Leisure		
		Dummy	× Asset Loss	× Sp. Loss	Dummy	× Asset Loss	× Sp. Loss	Dummy	× Asset Loss	× Sp. Loss	Dummy	× Asset Loss	× Sp. Loss
Before Separation (wave)	1st	(control)			(control)			(control)			(control)		
	2nd	9.25e-06			-0.000513			0.0260			0.0288		
		(0.0122)			(0.0123)			(0.0387)			(0.0391)		
	3rd	-0.00791			-0.00897			0.0365			0.0415		
		(0.0190)			(0.0191)			(0.0629)			(0.0631)		
	4th	-0.0250			-0.0268			0.00185			0.0107		
		(0.0230)			(0.0231)			(0.0775)			(0.0779)		
	5th	0.0194			0.0171			-0.00790			0.00365		
		(0.0269)			(0.0273)			(0.0944)			(0.0954)		
Obs.		50,376			50,376			50,806			50,806		
R^2		0.140			0.143			0.152			0.158		

Note: Model 1 (columns 1-3) represents estimates of a fixed effect regresstion of leisure demand on regressors as specified in the text. It is a repeat of an earlier table. Model 1 is then adjusted to account for joint leisure by including dummies for if the ex-spouse was working at the time of the divorce, and if he/she was working full-time. Model 2 follows a similar methodology but is a fixed effect regression with labor force participation as the dependent variable. Standard errors, clustered at the person level, are reported in parentheses. Results for the remaining regressors are available from the author.
*** $p<0.01$, ** $p<0.05$, * $p<0.1$.
"× Asset Loss": interaction of assets lost after separation with dummy for wave before/after separation (Assets in $100,000).
""× Sp. Loss": interaction of spousal income lost after separation with dummy for wave before/ after separation (Spousal income in $10,000 and is multiplied by number of potential years left to work if spouse would have worked until 65).

Table P.2: Comparisons of Estimates of Leisure Demand and Labor Force Participation of Men accounting for spousal leisure complementarities.

		Leisure Demand - Widower						Labor Force Participation - Widower					
		Original Model			Model with Joint Leisure			Original Model			Model with Joint Leisure		
		Dummy	× Asset Loss	× Sp. Loss	Dummy	× Asset Loss	× Sp. Loss	Dummy	× Asset Loss	× Sp. Loss	Dummy	× Asset Loss	× Sp. Loss
spouse works					-0.0223***						0.0947***		
					(0.00275)						(0.00878)		
spouse works full-time					-0.00957***						0.00655		
					(0.00325)						(0.00937)		
After Separation (wave)	0th	0.0251	0.0127***	-0.000170	0.0240	0.0125***	-0.000357	0.00961	-0.0472***	-0.00110	0.0164	-0.0465***	-0.000309
		(0.0274)	(0.00410)	(0.000552)	(0.0277)	(0.00429)	(0.000562)	(0.0648)	(0.0143)	(0.00147)	(0.0677)	(0.0153)	(0.00151)
	1st	0.0296	0.00909**	-0.000203	0.0303	0.00890*	-0.000410	-0.0259	-0.0363**	-0.000120	-0.0242	-0.0356*	0.000793
		(0.0286)	(0.00447)	(0.000632)	(0.0289)	(0.00471)	(0.000648)	(0.0882)	(0.0183)	(0.00155)	(0.0903)	(0.0192)	(0.00158)
	2nd	0.0215	0.00852*	-0.00100	0.0236	0.00834*	-0.00125*	-0.0389	-0.0300	0.00229	-0.0417	-0.0294	0.00336*
		(0.0284)	(0.00464)	(0.000717)	(0.0288)	(0.00487)	(0.000728)	(0.104)	(0.0190)	(0.00194)	(0.107)	(0.0199)	(0.00198)
	3rd	0.0371	0.00549	0.00246*	0.0406	0.00534	0.00214	-0.0552	-0.0210	-0.0125***	-0.0634	-0.0204	-0.0111***
		(0.0375)	(0.00618)	(0.00145)	(0.0377)	(0.00634)	(0.00143)	(0.125)	(0.0213)	(0.00362)	(0.127)	(0.0221)	(0.00354)
	4th	0.0340	0.00562	0.0195***	0.0388	0.00541	0.0184***	-0.0607	-0.0198	-0.0599***	-0.0727	-0.0191	-0.0562***
		(0.0384)	(0.00596)	(0.00405)	(0.0389)	(0.00614)	(0.00409)	(0.132)	(0.0194)	(0.0115)	(0.135)	(0.0202)	(0.0117)
	5th	0.0281	0.00273		0.0339	0.00255		0.0712	-0.0123		0.0557	-0.0117	
		(0.0422)	(0.00732)		(0.0427)	(0.00751)		(0.138)	(0.0229)		(0.140)	(0.0238)	
	6th	0.0437	-0.00996	0.0650**	0.0508	-0.0105	0.0594**	0.0434	-0.0118	-0.263***	0.0234	-0.00990	-0.244***
		(0.0409)	(0.0119)	(0.0253)	(0.0413)	(0.0119)	(0.0254)	(0.145)	(0.0450)	(0.0746)	(0.147)	(0.0464)	(0.0749)
	7th	0.0205	-0.00693	-0.00587***	0.0286	-0.00764	-0.00632***	-0.00948	0.0384	0.0191***	-0.0331	0.0411	0.0210***
		(0.0359)	(0.00753)	(0.00175)	(0.0362)	(0.00758)	(0.00177)	(0.165)	(0.0373)	(0.00579)	(0.166)	(0.0374)	(0.00589)

		Leisure Demand - Widower						Labor Force Participation - Widower					
		Original Model			Model with Joint Leisure			Original Model			Model with Joint Leisure		
		Dummy	× Asset Loss	× Sp. Loss	Dummy	× Asset Loss	× Sp. Loss	Dummy	× Asset Loss	× Sp. Loss	Dummy	× Asset Loss	× Sp. Loss
Before Separation (wave)	1st	(control)			(control)			(control)			(control)		
	2nd	-0.0304			-0.0294			-0.00124			-0.00361		
		(0.0298)			(0.0299)			(0.0392)			(0.0384)		
	3rd	-0.0444			-0.0416			0.106			0.0979		
		(0.0302)			(0.0303)			(0.0998)			(0.101)		
	4th	-0.0499			-0.0453			0.348***			0.330***		
		(0.0372)			(0.0383)			(0.109)			(0.111)		
	5th	-0.0393			-0.0302			0.254**			0.225*		
		(0.0336)			(0.0350)			(0.122)			(0.117)		
Obs.		43,263			43,263			43,693			43,693		
R^2		0.242			0.247			0.229			0.235		

Note: Model 1 (columns 1-3) represents estimates of a fixed effect regresstion of leisure demand on regressors as specified in the text. It is a repeat of an earlier table. Model 1 is then adjusted to account for joint leisure by including dummies for if the ex-spouse was working at the time of the divorce, and if he/she was working full-time. Model 2 follows a similar methodology but is a fixed effect regression with labor force participation as the dependent variable. Standard errors, clustered at the person level, are reported in parentheses. Results for the remaining regressors are available from the author.

*** p<0.01, ** p<0.05, * p<0.1.

"× Asset Loss": interaction of assets lost after separation with dummy for wave before/after separation (Assets in $100,000).

""× Sp. Loss": interaction of spousal income lost after separation with dummy for wave before/ after separation (Spousal income in $10,000 and is multiplied by number of potential years left to work if spouse would have worked until 65).

Table P.3: Probit estimates of retirement using alternative retirement definition (women)

	Women			Women (weighted)		
	(1)	(2)	(3)	(1w)	(2w)	(3w)
SEP	-0.332**	-0.507*	-0.494*	-0.209	-0.843**	-0.773**
	(0.134)	(0.279)	(0.273)	(0.164)	(0.335)	(0.340)
	[-0.069]	[-0.106]	[-0.103]	[-0.045]	[-0.179]	[-0.164]
×AssetLoss		-0.000386	-0.0138		-0.00643	-0.0388
		(0.0318)	(0.0365)		(0.0305)	(0.0391)
		[-0.000]	[-0.003]		[-0.001]	[-0.008]
×SpIncLoss		0.00135	0.00308		0.0144	0.0141
		(0.00888)	(0.0101)		(0.00882)	(0.0128)
×WvsSinceSep		0.0660	0.0630		0.167*	0.131
		(0.0709)	(0.0789)		(0.0856)	(0.0931)
×WvsSinceSep × AssetLoss			0.00735			0.0196
			(0.0202)			(0.0196)
			[0.002]			[0.004]
×WvsSinceSep × SpIncLoss			-0.00225			-5.87e-06
			(0.00591)			(0.00682)
WID	0.219	0.0266	0.0542	0.118	-0.0635	-0.0171
	(0.151)	(0.240)	(0.259)	(0.189)	(0.245)	(0.267)
×AssetLoss		0.0239	-0.141		0.0194	-0.0967
		(0.0421)	(0.162)		(0.0595)	(0.125)
×SpIncLoss		0.0145*	-0.00465		0.0174*	-0.0102
		(0.00863)	(0.0151)		(0.00998)	(0.0151)
×WvsSinceWid		0.0641	-0.0165		0.0407	-0.0105
		(0.0861)	(0.124)		(0.0960)	(0.119)
×WvsSinceWid × AssetLoss			0.100			0.0403
			(0.103)			(0.0778)
×WvsSinceWid × SpIncLoss			0.0810**			0.101***
			(0.0331)			(0.0308)
EARN	1.14e-07	1.07e-07	1.22e-07	6.64e-07	6.34e-07	6.47e-07
	(4.55e-07)	(4.55e-07)	(4.54e-07)	(6.25e-07)	(6.22e-07)	(6.22e-07)
ASSETS	0.00124	0.00122	0.00119	-0.000351	-0.000397	-0.000433
	(0.00152)	(0.00152)	(0.00153)	(0.00217)	(0.00217)	(0.00218)
N	8,225	8,225	8,225	6,209	6,209	6,209
Pseudo R^2	0.0842	0.0847	0.0861	0.0737	0.0752	0.0766

Notes: Dependent variable is retirement in the next period measure by annual hours of work. If an individual transitions from working more than 300 hours per year to less than 300 hours per year, the individual is treated as retired (the individual is still required to be working full-time in his or her first interview wave). Marginal effects included in square brackets for select covariates. Retirement is treated as an absorbing state, so observations after retirement are dropped from the estimation. "weighted" indicates that data is weighted using person-level weights corresponding to the specific HRS interview wave. Regressions are conducted only on individuals between ages 50 and 70 who worked full-time in their first HRS interview wave, were at-risk for a divorce/widowhood between ages 50 and 70, did not remarry, and did not ever apply for disability. Descriptions of the included covariates are in the text.

Table P.4: Probit estimates of retirement using alternative retirement definition (men)

	Men			Men (weighted)		
	(1)	(2)	(3)	(1w)	(2w)	(3w)
SEP	-0.132	0.127	0.115	-0.178	0.107	0.0802
	(0.128)	(0.207)	(0.220)	(0.160)	(0.287)	(0.294)
	[-0.028]	[0.026]	[0.024]	[-0.036]	[0.021]	[0.016]
×AssetLoss		0.00111	-0.0332		-0.0280	-0.0736
		(0.0274)	(0.0491)		(0.0410)	(0.0660)
×SpIncLoss		-0.00747*	-0.00439		-0.00685	-0.00367
		(0.00424)	(0.00529)		(0.00491)	(0.00571)
×WvsSinceSep		-0.0345	-0.0362		-0.0427	-0.0416
		(0.0711)	(0.0804)		(0.0967)	(0.115)
×WvsSinceSep × AssetLoss			0.0263			0.0371
			(0.0264)			(0.0320)
×WvsSinceSep × SpIncLoss			-0.00280			-0.00361
			(0.00267)			(0.00307)
WID	0.347	0.426	0.549*	0.447*	0.393	0.660
	(0.245)	(0.295)	(0.299)	(0.265)	(0.436)	(0.436)
	[0.074]	[0.091]	[0.117]	[0.091]	[0.080]	[0.134]
×AssetLoss		-0.0917	0.00847		-0.0752	0.122
		(0.0673)	(0.133)		(0.0601)	(0.118)
		[-0.020]	[0.002]		[-0.015]	[0.024]
×SpIncLoss		0.00361	-0.000579		0.00466	-0.000472
		(0.00549)	(0.0137)		(0.00539)	(0.0138)
×WvsSinceWid		-0.238	-0.413		-0.146	-0.517
		(0.199)	(0.254)		(0.244)	(0.395)
		[-0.000]	[-0.000]		[-0.000]	[-0.000]
×WvsSinceWid × AssetLoss			-0.206			-0.370*
			(0.216)			(0.202)
			[0.000]			[0.000]
×WvsSinceWid × SpIncLoss			0.00413			0.00345
			(0.0114)			(0.0117)
EARN	1.81e-07	1.79e-07	1.81e-07	1.14e-07	1.13e-07	1.15e-07
	(1.33e-07)	(1.33e-07)	(1.33e-07)	(1.21e-07)	(1.20e-07)	(1.21e-07)
ASSETS	0.000738	0.000724	0.000735	0.00181	0.00180	0.00181
	(0.000959)	(0.000959)	(0.000960)	(0.00180)	(0.00180)	(0.00179)
N	11,424	11,424	11,424	10,693	10,693	10,693
Pseudo R^2	0.0942	0.0949	0.0952	0.0886	0.0892	0.0899

Notes: Dependent variable is retirement in the next period measure by annual hours of work. If an individual transitions from working more than 300 hours per year to less than 300 hours per year, the individual is treated as retired (the individual is still required to be working full-time in his or her first interview wave). Marginal effects included in square brackets for select covariates. Retirement is treated as an absorbing state, so observations after retirement are dropped from the estimation. "weighted" indicates that data is weighted using person-level weights corresponding to the specific HRS interview wave. Regressions are conducted only on individuals between ages 50 and 70 who worked full-time in their first HRS interview wave, were at-risk for a divorce/widowhood between ages 50 and 70, did not remarry, and did not ever apply for disability. Descriptions of the included covariates are in the text.

Table P.5: Probit estimates of retirement - robustness checks (women)

	Women				Women (weighted)			
	(2lc)	(3lc)	(2ps)	(3ps)	(2lc)	(3lc)	(2ps)	(3ps)
SEP	-0.747***	-0.679**	-0.829**	-0.895**	-0.836***	-0.804**	-0.908*	-1.115**
	(0.277)	(0.281)	(0.417)	(0.434)	(0.304)	(0.326)	(0.496)	(0.495)
×AssetLoss	0.0574*	-0.0333	0.0582	-0.0699	0.0276	-0.0849	0.0290	-0.138
	(0.0314)	(0.0471)	(0.0361)	(0.0559)	(0.0304)	(0.0765)	(0.0353)	(0.103)
×SpIncLoss	0.00881	0.0119	0.0111	0.0162*	0.0196***	0.0245***	0.0214*	0.0308***
	(0.00661)	(0.00741)	(0.00885)	(0.00860)	(0.00706)	(0.00925)	(0.0112)	(0.0106)
×WvsSinceSep	0.138**	0.0448	0.156**	0.0541	0.132	0.0245	0.188**	0.105
	(0.0694)	(0.0906)	(0.0696)	(0.0912)	(0.0812)	(0.115)	(0.0923)	(0.103)
×WvsSinceSep × AssetLoss		0.119***		0.148***		0.186**		0.228***
		(0.0408)		(0.0480)		(0.0728)		(0.0839)
×WvsSinceSep × SpIncLoss		-0.00470		-0.00473		-0.0114**		-0.0136**
		(0.00328)		(0.00354)		(0.00515)		(0.00540)
WID	0.0183	0.0363	0.0660	0.0832	-0.136	-0.120	-0.0635	-0.0491
	(0.240)	(0.247)	(0.239)	(0.246)	(0.265)	(0.270)	(0.263)	(0.269)
×AssetLoss	0.0456	0.00310	0.0469	0.00376	0.0189	0.0204	0.0211	0.0223
	(0.0514)	(0.101)	(0.0515)	(0.101)	(0.0656)	(0.108)	(0.0652)	(0.109)
×SpIncLoss	0.00454	-0.00355	0.00465	-0.00335	0.00755	-0.00980	0.00745	-0.00969
	(0.0112)	(0.0130)	(0.0112)	(0.0129)	(0.0114)	(0.0145)	(0.0113)	(0.0144)
×WvsSinceWid	-0.0476	-0.0646	-0.0497	-0.0663	-0.0671	-0.0700	-0.0680	-0.0703
	(0.0905)	(0.0951)	(0.0908)	(0.0954)	(0.104)	(0.101)	(0.104)	(0.102)
×WvsSinceWid × AssetLoss		0.0150		0.0158		-0.0284		-0.0278
		(0.0590)		(0.0590)		(0.0594)		(0.0594)
×WvsSinceWid × SpIncLoss		0.0317		0.0315		0.0426*		0.0423*
		(0.0243)		(0.0243)		(0.0248)		(0.0247)
EARN	-3.03e-07	-3.22e-07	-2.92e-07	-3.08e-07	-2.73e-07	-3.06e-07	-2.58e-07	-2.83e-07
	(5.23e-07)	(5.25e-07)	(5.25e-07)	(5.26e-07)	(7.70e-07)	(7.75e-07)	(7.69e-07)	(7.74e-07)
ASSETS	0.00575***	0.00574***	0.00556***	0.00554***	0.00746***	0.00745***	0.00703**	0.00701**
	(0.00188)	(0.00188)	(0.00185)	(0.00185)	(0.00281)	(0.00281)	(0.00279)	(0.00278)
N	8,022	8,022	8,022	8,022	5,944	5,944	5,944	5,944
Pseudo R^2	0.136	0.138	0.135	0.137	0.128	0.130	0.127	0.130

Notes: #lc corresponds to the same model # run for table P.5, except it accounts for the labor force participation of a married spouse, and whether that spouse is working full-time. #ps corresponds to the same model # run for table P.5, except that it accounts for the effect of pensions and Social Security by including dummies for whether the ex-spouse is retired, whether the ex-spouse is at least 62, and continuous variables for the time rule formula for divorcing households, and the relative value of survivor's benefit to the household benefit if each member claimed at 62 (or immediately if older). Dependent variable is retirement in the next period measured as a transition from employment in the current period (full-time, part-time, unemployed) to retirement in the next period (defined by the next period being retired, partly retired, or not in labor force). Marginal effects included in square brackets for select covariates. Retirement is treated as an absorbing state, so observations after retirement are dropped from the estimation. "weighted" indicates that data is weighted using person-level weights corresponding to the specific HRS interview wave. Regressions are conducted only on individuals between ages 50 and 70 who worked full-time in their first HRS interview wave, were at-risk for a divorce/widowhood between ages 50 and 70, did not remarry, and did not ever apply for disability. Descriptions of the included covariates are in the text.

Table P.6: Probit estimates of retirement - robustness checks (men)

	Men				Men (weighted)			
	(2lc)	(3lc)	(2ps)	(3ps)	(2lc)	(3lc)	(2ps)	(3ps)
SEP	-0.284	0.00734	-0.267	0.0716	-0.274	-0.0529	-0.210	0.0635
	(0.227)	(0.238)	(0.244)	(0.312)	(0.295)	(0.303)	(0.309)	(0.365)
×AssetLoss	0.0265	0.0327	0.0256	0.0397	0.0332	0.0387	0.0303	0.0437
	(0.0195)	(0.0461)	(0.0236)	(0.0441)	(0.0240)	(0.0445)	(0.0271)	(0.0439)
×SpIncLoss	-0.00614*	-0.0188*	-0.00686*	-0.0180*	-0.00647	-0.0154	-0.00655	-0.0153*
	(0.00346)	(0.00979)	(0.00369)	(0.0104)	(0.00485)	(0.00939)	(0.00428)	(0.00916)
×WvsSinceSep	0.0668	-0.0544	0.138	0.00883	0.0200	-0.0827	0.0775	-0.0404
	(0.0955)	(0.103)	(0.111)	(0.138)	(0.0989)	(0.113)	(0.120)	(0.151)
×WvsSinceSep × AssetLoss		-0.00507		-0.00558		-0.00773		-0.00948
		(0.0240)		(0.0216)		(0.0255)		(0.0237)
×WvsSinceSep × SpIncLoss		0.00535		0.00456		0.00484		0.00463
		(0.00337)		(0.00351)		(0.00356)		(0.00347)
WID	0.467*	0.649**	0.513*	0.695**	0.439	0.843**	0.475	0.880**
	(0.281)	(0.270)	(0.281)	(0.270)	(0.376)	(0.351)	(0.375)	(0.350)
×AssetLoss	-0.102*	0.0197	-0.102*	0.0190	-0.0692	0.168	-0.0707	0.166
	(0.0548)	(0.145)	(0.0548)	(0.145)	(0.0474)	(0.115)	(0.0475)	(0.115)
×SpIncLoss	0.000778	-0.00869	0.000819	-0.00867	0.00253	-0.00860	0.00261	-0.00861
	(0.00466)	(0.0131)	(0.00466)	(0.0132)	(0.00461)	(0.0136)	(0.00460)	(0.0136)
×WvsSinceWid	-0.348	-0.649**	-0.349	-0.653**	-0.218	-0.886**	-0.218	-0.890**
	(0.241)	(0.323)	(0.241)	(0.324)	(0.274)	(0.407)	(0.275)	(0.409)
×WvsSinceWid × AssetLoss		-0.321		-0.321		-0.532***		-0.532***
		(0.249)		(0.249)		(0.195)		(0.195)
×WvsSinceWid × SpIncLoss		0.00989		0.00994		0.0106		0.0107
		(0.00963)		(0.00963)		(0.00945)		(0.00946)
EARN	1.42e-08	1.37e-08	2.94e-08	2.89e-08	2.76e-09	2.24e-09	2.81e-09	2.48e-09
	(1.62e-07)	(1.62e-07)	(1.55e-07)	(1.55e-07)	(1.24e-07)	(1.24e-07)	(1.24e-07)	(1.24e-07)
ASSETS	0.00352**	0.00353**	0.00362**	0.00363**	0.00436**	0.00437**	0.00441**	0.00442**
	(0.00156)	(0.00156)	(0.00157)	(0.00157)	(0.00177)	(0.00177)	(0.00180)	(0.00180)
N	10,461	10,461	10,461	10,461	9,854	9,854	9,854	9,854
Pseudo R^2	0.141	0.142	0.141	0.142	0.127	0.128	0.127	0.128

Notes: #lc corresponds to the same model # run for table P.5, except it accounts for the labor force participation of a married spouse, and whether that spouse is working full-time. #ps corresponds to the same model # run for table P.5, except that it accounts for the effect of pensions and Social Security by including dummies for whether the ex-spouse is retired, whether the ex-spouse is at least 62, and continuous variables for the time rule formula for divorcing households, and the relative value of survivor's benefit to the household benefit if each member claimed at 62 (or immediately if older). Dependent variable is retirement in the next period measured as a transition from employment in the current period (full-time, part-time, unemployed) to retirement in the next period (defined by the next period being retired, partly retired, or not in labor force). Marginal effects included in square brackets for select covariates. Retirement is treated as an absorbing state, so observations after retirement are dropped from the estimation. "weighted" indicates that data is weighted using person-level weights corresponding to the specific HRS interview wave. Regressions are conducted only on individuals between ages 50 and 70 who worked full-time in their first HRS interview wave, were at-risk for a divorce/widowhood between ages 50 and 70, did not remarry, and did not ever apply for disability. Descriptions of the included covariates are in the text.

Q Performance Comparisons across Landscapes

In this paper, anytime we are drawing comparisons across different types of influence matrices, we use ranked performance results. Performance measures can not be compared directly across separate aggregation regimes (i.e. landscapes generated from different γ or $\mathbf{I}_{ij}$ in equation (15.2)), because the distribution of performance values differs significantly based on the firm level aggregation measure that is chosen. Figure (15.5)(A) shows the distribution of the performance landscape with a 75% complementary influence matrix (i.e. $\mathbf{I_{ij}}$), whereas figures Q.1(A) and (B) show the distribution of the performance landscape at both extremes in complementarity of the influence matrix. Due to the differences in distributions, comparing these two landscapes directly would likely result in a misinterpretation of the results because performance is not normalized to some consistent measure.[67]

When Kauffman introduced the *NK* model in his Kauffman (1993) book, he originally ranked the performance levels and treated them only ordinally; however, he later begins to treat these values as cardinal without a discussion why. It is possible that this might make sense in the case of evolutionary biology, but there is no clear reason why business strategists should expect a performance of 0.8 versus a performance of 0.88 to be 10% better. In fact, given figure Q.1(B), we might expect this change in performance to be better than over 10% of the possible performance levels or, alternatively, much smaller for figure Q.1(A). As a result, we propose that treating performance measures as ordinal measures and rank ordering the results is the only correct way of comparing one performance measure to another. In this way, a rank ordering permits the author to argue that a performance level of x could be better than a certain percentage of other possible organizational combinations.

[67]For example, in figure (Q.1)(B), 0.8 would correspond to the average *possible* performance level, whereas in figure (Q.1)(A), the same 0.8 would correspond to a possible performance level in the top 10% of all possible performance levels.

Figure Q.1: Landscape Characteristics - Structured Contributions ($K = 2$)

(a) Distribution of Business Landscape with a 100% complementary interaction matrix. (0% substitutable)

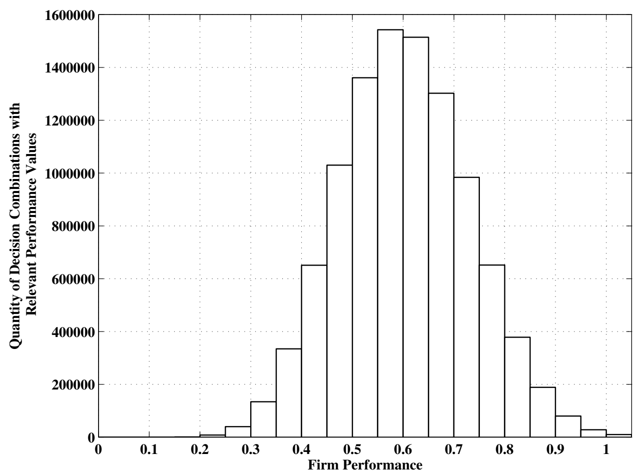

(b) Distribution of Business Landscape with a 0% complementary interaction matrix. (100% substitutable)

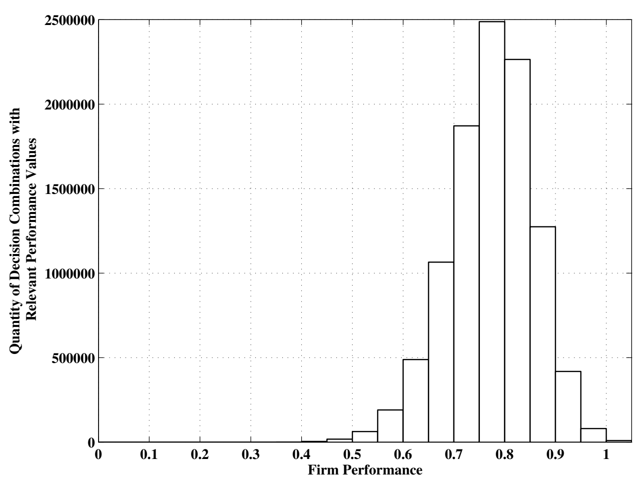

References

AARP (2004): "The Divorce Experience: A Study of Divorce at Midlife and Beyond," report by Xenia P. Montenegro,.

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

Altmeyer, A. (1966): *Formative Years of Social Security*. University of Wisconsin Press.

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

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

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

Blundell, R., and T. MaCurdy (1999): "Labor Supply: A Review of Alternative Approaches," in *Handbook of Labor Economics*, ed. by O. Ashenfelter, and D. Card, vol. 3. Elsevier Science.

Browning, M., P.-A. Chiappori, and A. Lewbel (2013): "Estimating Consumption Economies of Scale, Adult Equivalence Scales, and Household Bargaining Power," *Review of Economic Studies*, Forthcoming.

Butrica, B. A., D. Murphy, and S. R. Zedlewski (2007): "How Many Struggle to get by in Retirement?," Center for Retirement Research (CRR) Working Paper, WP 2007-27. http://crr.bc.edu/wp-content/uploads/2008/02/wp _2007-27-508.pdf.

California Public Employees' Retirement System (2013): "A guide to CalPERS Community Property," Discussion paper.

Citro, C. F., and R. T. Michael (1995): *Measuring Poverty: A New Approach*. National Academy of Press, Washington, DC.

Clarke, S. C. (1995): "Advance report of final divorce statistics, 1989 and 1990," *Monthly Vital Statistics Report*, 43(9), 1–31.

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

Coile, C., P. Diamond, J. Gruber, and A. Jousten (2002): "Delays in Claiming Social Security Benefits," *Journal of Public Economics*, 84(3), 357–385.

——— (2007): "Future Social Security Entitlements and the Retirement Decision," *The Review of Economics and Statistics*, 89(2), 234–246.

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

De Nardi, M., E. French, and J. B. Jones (2010): "Why Do the Elderly Save? The Role of Medical Expenses," *Journal of Political Economy*, 118(1), 39–75.

Ethiraj, S. K., and D. A. Levinthal (2004a): "Bounded Rationality and the Search for Organizational Architecture: An Evolutionary Perspective on the Design of Organizations and Their Evolvability," *Administrative Science Quarterly*, 49(3), 404–437.

——— (2004b): "Modularity and Innovation in Complex Systems," *Management Science*, 50(2), 159–173.

Fang, C., J. Lee, and M. Schilling (2010): "Balancing exploration and exploitation through structural design: The isolation of subgroups and organizational learning.," *Organization Science*, 21(3), 625–642.

Fernandez, R., and J. C. Wong (2013): "Divorce Risk, Wages, and Working Wives: A Quantitative Life-Cycle Analysis of Female Labor Force Participation," .

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

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

Friedberg, L., and A. Webb (2005): "Retirement and the Evolution of Pension Structure," *Journal of Human Resources*, 40(2), 281–308.

Gavetti, G., and D. A. Levinthal (2000): "Looking Forward and Looking Backward: Cognitive and Experiential Search," *Administrative Science Quarterly*, 45(1), 113–137.

Ghemawat, P., and D. A. Levinthal (2008): "Choice Interactions and Business Strategy," *Management Science*, 54(9), 1638–1651.

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

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

Gourinchas, P.-O., and J. A. Parker (2002): "Consumption over the Life Cycle," *Econometrica*, 70(1), 47–89.

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

Gustman, A. L., and T. L. Steinmeier (1986a): "A Structural Retirement Model," *Econometrica*, 54(3), 555–584.

——— (1999): "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.

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

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

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

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

Haider, S., and G. Solon (2000): "Nonrandom Selection in the HRS Social Security Earnings Sample," in *Labor and Population Program, Working Paper Series 00-01*. RAND.

Hausman, J. (1981): "Labor supply, how taxes affect economic behavior," in *Tax and the economy*, ed. by H. Aaron, and J. Pechman. Brookings Institution.

Henderson, R. M., and K. B. Clark (1990b): "Architectural innovation: The reconfiguration of existing product technologies and the failure of established firms," *Administrative Science Quarterly*, 35(1), 9–30.

Hubbard, R. G., J. Skinner, and S. P. Zeldes (1995): "Precautionary Saving and Social Insurance," *Journal of Political Economy*, 103(2), 360–399.

Hurd, M. D., J. P. Smith, and J. M. Zissimopoulos (2002): "The Effects of Subjective Survival on Retirement and Social Security Claiming," in *NBER Working Paper 9140*. National Bureau of Economic Research.

Johnson, W. R., and J. Skinner (1986): "Labor Supply and Marital Separation," *American Economic Review*, 76(3), 455–469.

Kapteyn, A., P.-C. Michaud, J. P. Smith, and A. V. Soest (2006): "Effects of Attrition and Non-Response in the Health and Retirement Study," in *Labor and Population Program, Working Paper Series WR-407*. RAND.

Kauffman, S. (1993): *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press, New York.

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

Kennedy, S., and S. Ruggles (2013): "Breaking up is Hard to Count: The Rise of Divorce and Cohabitation Instability in the United States, 1980-2010," Minnesota Population Center Working Paper Series, WP 2013-01.

Knapp, D. M. K. (2013): "The Effect of Social Security Auxiliary Spouse and Survivor Benefits on the Household Retirement Decision," mimeo.

Lenox, M., S. Rockart, and A. Lewin (2006): "Interdependency, Competition, and the Distribution of Firm and Industry Profits.," *Management Science*, 52(5), 757–772.

——— (2007): "Interdependency, Competition, and Industry Dynamics," *Management Science*, 53(4), 599–615.

Levinthal, D. A. (1997): "Adaptation on Rugged Landscapes," *Management Science*, 43(7), 934–950.

——— (2011): "A behavioral approach to strategy-what's the alternative?," *Strategic Management Journal*, 32(13), 1517–1523.

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

McGahan, A. (2004): *How Industries Evolve: Principles for Achieving and Sustaining Superior Performance*, Harvard Business School Press Books. Harvard Business School Press, Boston, MA.

Milgrom, P., and J. Roberts (1990): "The Economics of Modern Manufacturing: Technology, Strategy, and Organization," *American Economic Review*, 80(3), 511–528.

——— (1995a): "Complementarities and fit Strategy, structure, and organizational change in manufacturing," *Journal of Accounting and Economics*, 19, 179–208.

——— (1995b): "The Economics of Modern Manufacturing: Reply," *American Economic Review*, 85(4), 997–999.

Miller, D. (1992): "Environmental fit versus internal fit," *Organization Science*, 3(2), 159–178.

Nickerson, J., and T. R. Zenger (2004): "A knowledge-based Theory of the Firm - the Problem-Solving Perspective," *Organization Science*, 15(6), 617–632.

Peterman, W. B. (2012): "Reconciling Micro and Macro Estimates of the Frisch Labor Supply Elasticity," Federal Reserve Board of Governors.

Porter, M., and N. Siggelkow (2008): "Contextuality Within Activity Systems and Sustainability of Competitive Advantage," *Academy of Management Perspectives*, 22(2), 34–56.

Porter, M. E. (1985): *Competitive Advantage*. Free Press, New York.

Porter, M. E. (1996): "What is strategy?," *Harvard Business Review*, 74(6), 61–78.

Rivkin, J. W. (2000): "Imitation of Complex Strategies," *Management Science*, 46(6), 824–844.

Rivkin, J. W., and N. Siggelkow (2003): "Balancing Search and Stability: Interdependencies Among Elements of Organizational Design," *Management Science*, 49(3), 290–311.

——— (2007): "Patterned Interactions in Complex Systems: Implications for Exploration," *Management Science*, 53(7), 1068–1085.

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

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

Sanchez, R., and J. Mahoney (1996): "Modularity, Flexibility, and Knowledge Management in Product and Organization Design," *Strategic Management Journal*, 17(Winter), 63–76.

Sass, S. A., W. Sun, and A. Webb (2013): "Social Security Claiming Decision of Married Men and Widow Poverty," *Economics Letters*, 119, 20–23.

Sherr, I., and S. Ovide (2013): "Computer Sales in Free Fall," Discussion paper, Wall Street Journal Online, http://online.wsj.com/news/articles/SB10001424127887324695104578414973888155516.

Shoven, J. B., and S. N. Slavov (2013): "Recent Changes in the Gains from Delaying Social Security," in *Working Paper 19370*. National Bureau of Economic Research.

Siggelkow, N. (2001): "Change in the Presence of Fit: The Rise, the Fall, and the Renaissance of Liz Claiborne," *Academy of Management Journal*, 44(4), 838–857.

——— (2011): "Firms as Systems of Interdependent Choices," *Journal of Management Studies*, 48(5), 1126–1140.

Siggelkow, N., and D. A. Levinthal (2003): "Temporarily Divide to Conquer: Centralized, Decentralized, and Reintegrated Organizational Approaches to Exploration and Adaptation," *Organization Science*, 14(6), 650–669.

Siggelkow, N., and J. W. Rivkin (2005): "Speed and Search: Designing Organizations for Turbulence and Complexity," *Organization Science*, 16(2), 101–122.

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

Stevenson, B., and J. Wolfers (2007): "Marriage and Divorce: Changes and their Driving Forces," *Journal of Economic Perspectives*, 21(2), 27–52.

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

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

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

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

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

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

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

Weir, D. R., and R. J. Willis (2000): "Prospects for Widow Poverty," in *Forecasting Retirement Needs and Retirement Wealth*, ed. by O. S. Mitchell, P. B. Hammond, and A. M. Rappaport, chap. 7. The Pension Research Council of the Wharton School of the University of Pennsylvania.

Wingfield, N. (2013): "PC Sales Still in a Slump, Despite New Offerings," Discussion paper, New York Times, http://www.nytimes.com/2013/04/11/technology/data-show-steep-decline-in-pc-shipments.html.

Worthen, B., and I. Sherr (2012): "PC Sales Go Into a Tailspin," Discussion paper, Wall Street Journal Online, http://online.wsj.com/news/articles/SB10000872396390444657804578048981461803220.