Social Security, Endogenous Retirement and Intrahousehold Cooperation

Giovanni Gallipoli
University of British Columbia

Laura Turner
University of Toronto

March 20, 2012

Abstract

This paper studies the retirement incentives induced by the U.S. Social Security system in a framework which allows for different degrees of cooperation and strategic interaction between spouses. We develop a model in which spouses maximize a joint household utility function, subject to the additional constraint that—conditional on the partner's optimal response— neither one finds it optimal to deviate from the best constrained household allocation by either changing labor force status or by claiming early benefits. We show that accounting for "non-cooperative" behavior through this additional constraint can rationalize various choices of older US couples observed in the 1932-42 cohort of the Health and Retirement Study. In particular non-cooperative behavior helps in the explanation of two recurring puzzles in the retirement literature: (i) the clustering of benefit claiming at the early age of 62; (ii) the joint benefit claiming of couples. We contrast our findings to those from a standard model, extended to include a process for declining health and complementarities in leisure between spouses, and show that the latter can rationalize neither early nor joint claiming behavior whenever individuals can make benefit claiming decisions and labor force participation decisions separately. Finally, we assess the importance of variable health costs and of bequest motives in determining old age decisions. Our non-cooperative model offers interesting insights on welfare and policy.

VERY PRELIMINARY AND INCOMPLETE

1 Introduction

In recent years, economists have made a great deal of progress in resolving the traditional "puzzles" associated with U.S. retirement behavior among married workers, of which the primary are (1) the very large peak in regular retirement at age 62, the first year at which Social Security retirement benefits are available independent of health status, and (2) the tendency of husbands and wives to retire and claim benefits concurrently, with fully 59% of married couples born between 1932 and 1942 in the Health and Retirement Study (HRS) claiming benefits within two years from each other. The profile or benefit claims, capturing the "spike" in benefit claiming at 62, and the tendency of virtually *all* workers, male and female, to claim benefits before their 66th birthday, are plotted in figure 1. The upper panel of figure 1 shows the benefit claiming hazard of male (left) and female (middle) workers, and the cumulative hazards for both sexes over the age range 60-69 (right). The lower panel shows the same figures for single males and females. Table 1 shows the distribution of the difference between spouses within a household in the benefit claiming age and in age of final labour market exit for the same HRS sample.[1]

We argue that the tendency toward both early and joint benefit claiming among couples is, at the aggregate level, inconsistent with standard forward-looking optimizing behavior among a majority of households, as is the fact that the retirement hazards for marrieds are—comparing the top and bottom rows of table 1—effectively similar to the patterns for singles. The major reason for this inconsistency is the U.S. Social Security Administration

[1]The sample includes households in which the spouses are less than five years apart in age, neither one holding defined benefit pension wealth. From the perspective of simple eyeballing, however, the graphs look very similar when wider age ranges are considered (as pointed out by Rivas (2010) among others) and/or when db pension holders are included.

(SSA)'s rules governing survivor or widow benefits. For the majority of married households in which the husband's Primary Earning Amount (PIA), the benefit he is entitled to claim at the "regular retirement age" of 65 (in 2000), is substantially larger than the wife's at retirement,[2] delaying the main earner's benefit claim beyond age 62, or even 65, should be very attractive. Claimable benefits appreciate at roughly roughly 7% a year for each year of delayed claiming (a rate that is nearly actuarially fair averaged across the entire population)[3]. A main earner bequeaths his entire Social Security *benefit* to a surviving spouse at his death, with no adjustment made for the surviving spouse's own benefit claiming age. This has the effect of making the (expectation of) the surviving spouse's mortality profile the relevant schedule for maximization of Social Security wealth, at least if the household is otherwise unconstrained and can use assets to smooth consumption during the "retirement window", the couples' 60s.[4]

Table 1: HRS: Husband's minus wife's date of SS benefit claiming and final labour market exit in years

	-6	-4	-2	0	2	4	6
benefit claim (share in %)	2.9	14.9	14.7	58.9	4.9	4.8	1.3
labor force exit (share in %)	14.7	14.2	8.0	37.5	6.8	7.2	11.6

Two major elements have emerged in the literature as explanations of the early and joint retirement phenomena. These are are (1) the role of declining health for workers in late middle age and (2) the role of spousal complementarities in leisure. Recent dynamic structural

[2]In roughly 59% of couples in our HRS sample who both retire at 62, the husband's retirement benefit is more than 50% larger than the wife's.

[3]Specifically, using the Social Security rules that prevailed in the year 2000, benefits appreciate at a rate of roughly 6% per year between the ages of 62 and 65, and by 8% per year between ages 66 and 69.

[4]This argument is complicated when either the husband or his spouse holds a defined benefit (db) pension because of the Social Security offset rules at age 65 and above, that are features of many pensions, provide an incentive to claim benefits at age 62 (Stock and Wise (1990)) that work against the incentives discussed herein. As db pension incentives are in general quite difficult to model, we follow Rivas (2010) by excluding bd pension holders from our analysis.

Figure 1: Benefit claiming hazards by age in the 1932-42 HRS cohort

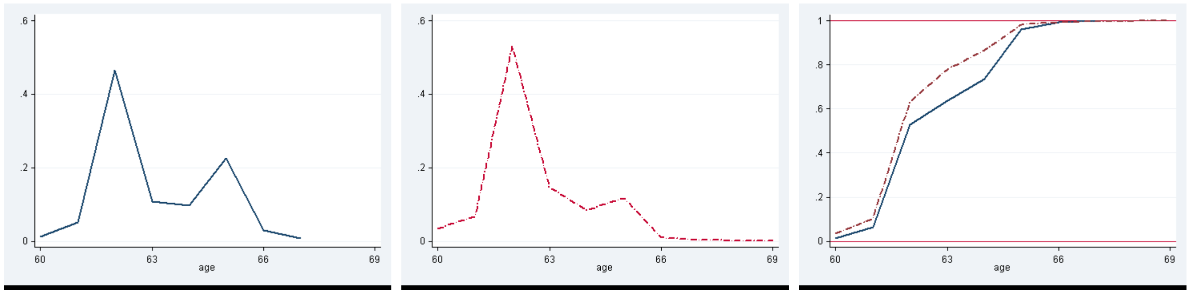

(a) married men (b) married women (c) married men and women

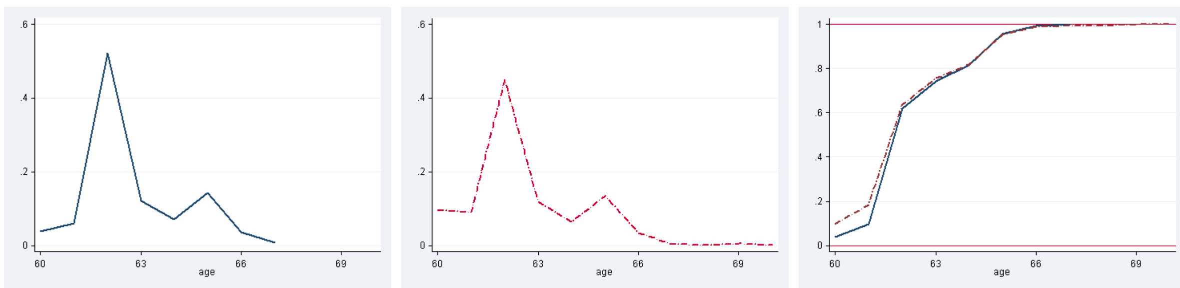

(d) single men (e) single women (f) single men and women

models of Social Security developed by Imrohoroglu and Kitao (2010), for the individual agent problem, and Van der Klaauw and Wolpin (2008) for the couple's problem, have focussed largely on the former, while a series of authors including Maestas (2001), Coile (2004), Gustman and Steinmeier (2000), Schirle (2008) and Rivas (2010) (among others), have emphasized the latter, using both structural and reduced form analysis. Both the enjoyment of shared leisure and poor health can clearly induce labor market exit by raising the relative marginal disutility of work. As well, if being in poor health decreases life expectancy, it will induce single agents to move their benefit claim dates forward in time. However, neither explanation is very satisfactory with regard to explaining the *benefit claiming*, as opposed to retirement, decisions of *married* households. A household with a sick husband who expects to die soon after retirement will optimally want to leave his healthy wife as large a benefit as possible in order to maximize the household's discounted payout from Social Security.[5] Complementarities in leisure, meanwhile, are may be a major factor in explaining the pattern of joint retirements among spouses, but—as shown in the second row of Table 1—the profile of intracouple differences in the age of final labour market exit is effectively more dispersed than the profile of differences in the date of benefit claims.

Several structural papers of retirement (including Van der Klaauw and Wolpin (2008), Rivas (2010), Gustman and Steinmeier (2004) and an earlier draft of this one) circumvent this problem by effectively "bundling" the retirement (labour force exit) and benefit claiming decisions so that individuals automatically claim benefits at the time of labor market exit, or vice versa. Figure 2 plots the pdf and cdf of the difference in age of benefit claiming and final labor market exit within our HRS sample over the range -10 difference in years up to +10 difference in years. Most people do indeed time their final labor market exit and benefit claiming age to roughly coincide, with a median difference of -.6 years for men and -.9 for women. At the 10th percentile, however, the difference between the dates is -3.8 years and at the 90th percentile it is 4.0 years for men, with very similar numbers for women.[6] The evidence suggests that, in fact, the Social Security benefit claiming decision and the final labour market exit decision are not taken as a single bundled decision, but in fact should be treated as separate and essentially independent choice variables.

[5]Back of the envelope calculations based on a 4% interest rate and health-specific mortality profiles calculated using US Life Tables and health and mortality data from the PSID suggest that the optimal benefit claiming date of a husband in poor health married to a wife in good health is 69, one year later than the optimal age if both husband or wife is healthy. The total Social Security wealth cost associated with retiring at 62 instead of 69 is 6.4%.

[6]The sample is very similar to the sample used to create figure 1, but restricted to include only married men and women whose final exit was at 60 or older, so as not to confuse early retirement with more general non-participation in late middle age.

Figure 2: Age of benefit claiming minus age of labour market exit: married men and women

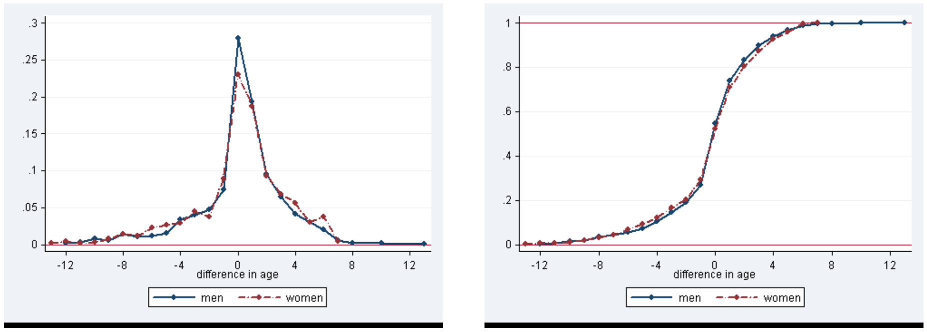

(a) PDF: men and women

(b) CDF: men and women

We show that the observed patterns of labor market exit and benefit claiming among couples without defined benefit (db) pension wealth are much harder to reconcile with life cycle theory in a model in which the labour force participation (what we call "career") decisions and Social Security benefit claiming decisions are made independently. To rationalize existing patterns of retirement and benefit claiming, we develop a "non-cooperative" model of retirement decision-making, i.e. one in violation of the dictates of a standard "sharing rule" governing household decisions in Rivas (2010) and Van der Klaauw and Wolpin (2008). Our model is similar in spirit, and is highly indebted, to the framework developed by Gustman and Steinmeier (2004) and Gustman and Steinmeier (2009). We expand on their model by assuming that, in each period, the household jointly optimizes its forward-looking utility (choosing how much to jointly consume and save, how much each spouse works, and whether or not each spouse should apply for Social Security benefits) subject to a set of incentive compatibility constraints ensuring that neither spouse prefers to unilaterally change his or her labour supply or benefit claim status given the household allocation. Complicating the problem is the fact that any unilateral deviation in own career or benefit claiming status is likely to provoke a corresponding response from the spouse. Under a few restrictions imposed on the nature of deviations, an incentive compatible allocation is one in which there is no alternative allocation which is a pure Nash equilibrium in benefit/career decisions, unless the allocation itself constitutes a pure strategy Nash equilibrium, and where, as in Gustman and Steinmeier (2009), the payoffs from the Nash gains are the continuation value functions for each spouse. In the next section, we attempt to show that the problem is more intuitive than it might sound as a model of household behavior.

Two further features of the model are very important for explaining our result, both of which have received a good deal of recent attention in the literature on retirement and wealth determination. First, as argued recently by DiNardi et al (2010) and Kopecky and Koreshkova (2009), individuals face an increasing, and increasingly variable, process for health expenditures as they age, which is in addition to evolving physical health status. These health expenditures, which are not funded under US Medicare, may be paid for either out of pocket or, subject to a wealth means-test, through public insurance, specifically Medicaid. Among the very elderly, the share receiving Medicaid support is in fact quite large. One of the main Medicaid- (rather than Medicare-) eligible old age health expenditures is nursing home care. Kopecky and Koreshkova (2009) show that two thirds of nursing home residents in the US are supported by Medicaid, while 40% of Americans who reach their 65th birthday eventually spend time in a nursing home.

While much recent research has been devoted to the implications of late-life health expenditure risk on wealth accumulation and inequality, in our model it also has implications for the timing of benefit claiming. The effect is two-fold. First, if a female spouse faces a high probability of finishing her life living on fixed income, as required to meet the Medicaid means test each period, the value of the survivor benefit under Social Security becomes heavily discounted. Second, the expectation that accumulated assets will eventually be "wasted" on health insurance, that would otherwise be publicly funded, reduces the ability to substitute liquid assets for Social Security wealth in the early retirement years. Thus, early labor market exit and benefit claiming by the main earning spouse can effectively increase the household's marginal propensity to consume out of wealth, shifting effective resources into the present, even when, as is the case for the majority of modest-earning households, the couple holds non-zero assets at retirement. Since husbands, who have lower life expectancies, are typically interested in shifting consumption into the present, this imperfect substitutability increases the incentive of lower-income husbands to retire early.[7]

The second feature of the model that is important to our results is the assumption a bequest motive that is binding upon wealthier couples, in particular those who do not expect a surviving spouse to deplete the household's assets on medical expenditures. In the model, even a modest bequest motive, as appears consistent with previous evidence by, for example Hendricks [2007], can induce relatively young couples (i.e. around age 62) to convert social security wealth into liquid assets that can be bequeathed be children or other heirs at the death of the surviving spouse. Using a standard functional form for the bequest motive, richer couples have larger marginal bequest motives, relative to the marginal utility of consumption, than poorer households. Since Social Security wealth is a relatively small component of total wealth for these couples, the conflict over the timing of optimal benefit

[7]The intertemporal allocation conflict was discussed by Lundberg et al. (2003).

claiming is smaller relative to poorer couples, and the shared bequest incentive induces wealthier couples to claim early.

We show that our non-cooperative or strategic retirement model does a good job of predicting the cross-sectional patterns presented in this section: the benefit claiming hazards among married and single households, including the spike in retirement at 62, which is broadly constant across wealth status, marital status and gender; the distribution of joint benefit claiming and joint retirement within couples; and the individual timing of benefit claiming and final labor market exit. In our benchmark model we assume that couples have purely egoistic preferences. In a forthcoming extension, we show that the results are robust, but only modestly improved, to allowing couples to benefit from complementarities in leisure when both are retired. Specifically, strategic behavior within couples is not incompatible with the idea that couples benefit from shared leisure time (as they do, for example, Gustman and Steinmeier (2004)). However, given the wealth distribution among retirement-age couples in the US, complementarities alone are not sufficient to explain the observed patterns of benefit claiming in a standard cooperative unitary model of retirement when participation and benefit claiming decisions are decoupled.

Finally, we want to use our model to engage in policy experiments. The experiments we consider are (1) eliminating survivor/widow benefits under Social Security, (2) eliminating the possibility of early retirement (before age 65 in 2000), (3) replacing half of the Social Security budget with a universal demogrant for the over-65 population, similar to the Canadian Old Age Security benefit, and (4) expanding Medicare to cover all "catastrophic" health costs (effectively including most nursing home care) for the over-65 population, regardless of wealth. The results of these experiments are forthcoming. Preliminary evidence suggests that the efficacy of these interventions varies substantially depending on whether or not households are cooperative joint utility maximizers or not. In fact, the benefits of all four interventions are substantially greater when members of couples make strategic retirement decisions, forcing the household to choose life cycle allocations that differ from first best.

The rest of the paper proceeds as follows. We present the household problem and the nature of the strategic intracouple retirement program in section 2. In section 3, we introduce the model economy and describe the sources of risk and the policy environment faced by individuals and households. In section 4, we present the main (very preliminary) results. We compare the predictions of our "strategic" model against those of a similarly estimated standard unitary cooperative model and show that the strategic model is a better predictor of the cross-sectional individual and joint retirement patterns of married households, particularly for husbands, between ages 60 and 69.[8]. In a forthcoming extension, we compare the welfare results from the two models under the existing Social Security program (the 'benchmark') and four policy experiments described above.

2 The household optimization problem

Households enter our model when the head is age 49 (if female) or 50 (otherwise). Females in the model are one year younger than males, the mean difference in age between husbands and wives in our HRS sample.[9] Men (women) survive to a maximum age of 101 (100) years.

[8] Disability benefits are available to individuals prior to age 62.

[9] The sample is restricted to households in which the spouses are less than five years apart in age. In the full same, the mean age difference between heads and spouses is roughly three years.

90% of households are married at age 50, and there is no divorce thereafter.[10] A model period is one year. Lifespan of individuals and households is uncertain ex-ante. $\varphi^j_{h,i}$ (hereafter φ_i for brevity) is the time invariant probability of living from age j to $j+1$ for an individual of gender i and health status h. When one member of a couple dies, the spouse lives in singlehood until death. Households differ ex-ante by their levels of assets, and the social security entitlements and wages of the members. 25% of households have negative assets of -\$10,000, 25% of households have zero assets, 25% hold assets of \$50,000, and the remaining 25% hold assets of \$250,000. The distribution of wages are chosen to match a correlation of .28 between the ln wages of wives and husbands, and a correlation of .17 (.39) between asset holdings and wages of men (women) at age 50. Individuals value consumption c and leisure l. They have gender-specific time discounting factors β_i and face a time-invariant interest rate $r = .04$.

At every age j, males and females find themselves in one of a theoretically infinite number of seven-dimensional variable states. The dimension are (1) accumulated non-human wealth holding a; (2) wage/productivity level w; (3) health and health expenditure state h, (4) accumulated Social Security credit from previous work e; (5) marital status $ms \in \{0, 1\} \equiv \{single, married\}$, (6) labour force status ls, and (7) benefit claim status ss. The last two states are described in more detail below.

There are two major sources of life cycle risk: productivity risk, which affects w, and risk of poor health, which imposes variable out-of-pocket medical costs, increases the marginal

[10] This is a possibly problematic assumption; however, divorce rates among couples over 50 are much lower than among younger couples in the US, and we assume that marital "sorting" has taken place prior to age 50 and that, even when spouses do not cooperate, there are always sufficient gains to marriage that those who have sorted into marriage which to remain.

utility of leisure, and increases mortality risk. In addition to a potential additive non-pecuniary utility gain, which we do not model, being married has two potential direct effects on welfare. The first arises due to the economies of scale in shared household-level consumption and is estimated in the model along with to match joint spousal labor supplies. Second, married individuals may receive utility from leisure enjoyed jointly.

Up to age 75, ls can take values of C for "career-job", NC for "non-career job" or R, not working ("retired"). For workers, wages are exogenously lower in non-career jobs, and individuals are more likely to be hit with a medical expenditure shock, conditional on physical health status, while in these jobs. Individuals can choose to move from a career to a non-career job by quitting the career job; while in a career job individuals work 45 hours per week, but may choose drop into non-career work which pays a lower hourly wage on average, and also requires only 35 hours per week. Non-career jobs transition into career jobs at rate p_C, which differs by gender. Social Security claim status, ss, also takes three values, NA for not receiving or claiming benefits, B for receiving benefits (either disability or retirement benefits), and A applying for disability benefits, prior to age 62.[11] While applying for (disability) benefits before age 62, individuals must choose $ls = R$. Their applications are successful at rate $p(j, h)$, i.e. the probability of receiving disability benefits depends on health status. It also increases with age in keeping with the Social Security Adminstration's rules for evaluating claims (see Low and Pistaferri (2009)). Applicants who apply from the career job state forfeit their job and fall off the career track into non-career work. Individuals

[11] We assume that once individuals reach 62 they automatically receive benefits even if they would prefer to apply for disability benefits. Specifically, we assume that individuals who retire between 62 and 65 receive the regular retirement age benefit with probability $p(h)$ and the age-adjusted benefit with probability $1 - p(h)$.

applying for regular benefits after age 62 can apply from ls states R or NC and receive benefits automatically (moving to ss state B with probability 1). We rule out applications for Social Security from career jobs because fewer than 4% of career workers simultaneously claim Social Security benefits, which is seldom optimal due to the Social Security earnings test, and to cut down on computational burden.

2.1 Life cycle optimization

Below, we sketch out the dynamic optimization problem for single and married households, beginning with single households whose problem is does not depend on assumptions about intrahousehold cooperation.

2.1.1 Singles

A single individual with gender i and age j has a generalized value function V_j^i, across ls and ss states. The state vector is $x = \{a, e, w, h, 0\}$. The dynamic optimization problem of the individual in a given ls and ss state (with corresponding value function $V_j^{i,ls,ss}$) and state vector x is:

$$V_j^{i,ls,ss}(x) = \max_{\{c,a'\}} u\big(c, l(ls,h)\big) + \beta_i \varphi_i \mathtt{E}_j[V_{j+1}^i(x')] \tag{1}$$

where the expectation $\mathtt{E}$ at age j is taken with respect to h and w. V is maximized subject to budget set B_S:

$$(\tilde{T} - l(h, ls))w(ls) + (1 + r)a + b(a, e, ss, ls) - \chi(h, ls) = I$$

$$I = c + a^{'} \quad a' \geq \underline{a} \qquad if \quad I - \underline{a} > \underline{c}$$

$$c = \underline{c} \quad a = \underline{a} \qquad if \quad I - \underline{a} \leq \underline{c}$$

where I is household period income net of health costs $\chi(h)$ and $\underline{c}$ is the minimum level of consumption guaranteed by the government. $\tilde{T}$ is a weekly time endowment. b is the benefit entitlement in the period, a function of assets, Social Security wealth, Social Security claim status and career status (hours worked). $\underline{a}$ is the minimum level of assets a person can hold. Since some people enter the model with debt, this value can be less than zero. However, we assume that households over 50 (whether single or married) cannot take on new debt, and so $\underline{a}_{kj} = min(0, \underline{a}_{kj-1})$ for household k at age j. Finally,

$$V^i_{j+1} = \max_{\Omega_j(ls,ss)} \left\{ V^{ls,ss}_{j+1} \right\} \tag{2}$$

and $\Omega(j, ls, ss)$ is the set of feasible $\{ls_{j+1}, ss_{j+1}\}$ combinations available given $\{ls_j, ss_j\}$ which are described above.

2.1.2 Marrieds

In a coopertive model of the household the dynamic problem for married couples is essentially the same as that for singles. We define the individual value function for husbands ($i = m$) and wives ($i = f$) Υ^i_j. The married household has a global value function U_j, a weighted sum of the individual members' Υs. U can be understood as the value function maximized by a household social planner whose task is to implement a predetermined marriage contract.

$$U_j(x_M) = \lambda \Upsilon_j^f(x_M) + (1-\lambda)\Upsilon_j^m(x_M) \tag{3}$$

The household's state vector is $x_M = \{a, e_m, e_f, w_m, w_f, h_m, h_f, \lambda\}$. Assets are accumulated at the household level, but social security benefit accumulation, wages and health are individual-level, $\imath-$ indexed states. The utility weight λ is chosen in the initial period of the couple household's existence. A typical assumption, initially proposed by Manser and Brown (1980) is that λ is chosen through cooperative (Nash) bargaining at the time of marriage, in which both household members' threat points are $V_1^{i,ls,ss}$ from the single person's optimization problem defined above. The exact value of λ is then partly a function of the partners' unobserved relative psychic gains from marriage: how much the husband and wife respectively value being married for non-economic reasons. In practice, we choose a constant λ for all couples that generates realistic mean labour supplies for husband and wives, and the model collapses to the simple unitary model.

The household budget set for marrieds, B_M is:

$$\begin{aligned}(\tilde{T} - l(h_f, ls_f))w_f(ls_f) + \tilde{T} - l(h_f, ls_f))w_f(ls_f) + (1+r)a \\ + b(a, e_f, e_m, ss_f, ls_f, ss_m, ls_m) - \chi^M(h_m, ls_m, h_f, ls_f) = I\end{aligned}$$

$$\begin{aligned} I = c + a^{'} \quad a' \geq \underline{a} \qquad if \quad I - \underline{a} > \underline{c}^M \\ c = \underline{c}^M \quad a = \underline{a} \qquad if \quad I - \underline{a} \leq \underline{c}^M \end{aligned}$$

B_M is analogous to B_S for a two-member household, and consumption is assumed to be a public, household good. The minimum consumption level for married households differs from that for singles.

The household planner's dynamic program and choice over ls and ss states of the household members is similar to the problem solved by singles, except that the planner chooses the optimal ls and ss for *both* members concurrently. There are therefore 81 potential $\{ls_m, ss_m, ls_f, ss_f\}$ combinations, though some (such as $\{C_m, A_m, C_f, NA_f\}$, in which the husband is applying for disability benefits from a career job) are ruled out by assumption. Each member of the household faces the same transition possibilities between ls and ss states for given age and individual state vector x as his or her single counterpart defined above, though the *payoffs* within each ls state depend not on x but on x_m.

2.1.3 Non-cooperation and strategic retirement

In the household problem described above, each partner receives a payoff of Υ^i for $i \in \{f, m\}$. These are not, in general, the payoffs that each spouse would receive if he could make decisions about his retirement and benefit claiming unilaterally, or if he could "dictate" an allocation to the household (e.g. for the husband, the scenario where $\lambda = 0$). The standard assumption of the efficient unitary model of the household is that members of the couple agree to sacrifice potential personal gains in order to maximize a joint outcome, and are fully able to commit to maintaining the outcome. We now consider the case where this no longer holds, and in which members of the couple may deviate from the joint outcome by choosing ls and/or ss unilaterally. Since each ls or ss choice provokes a potential response from the spouse, a deviator will take his partner's optimal $\{ls, ss\}$ response into account when choosing what to do. Defining individual payoffs under non-commitment to the planner's solution as $\hat{\Upsilon}$, the sposues, at (head's) age j and planner's allocation $Y = \{c, a', ls_m, ss_m, ls_f, ss_f\}$, jointly solve:

$$\hat{\Upsilon} j^{f} = \max_{\hat{\Omega}^{f}} \Upsilon_{j}^{f,\hat{ls_f},s\hat{s}_f;\hat{ls}_m,\hat{s}s_m|\hat{ls}_f,\hat{s}s_f}$$

$$\hat{\Upsilon} j^{m} = \max_{\hat{\Omega}^{m}} \Upsilon_{j}^{m,\hat{ls_m},s\hat{s}_m;\hat{ls}_f,\hat{s}s_f|\hat{ls}_m,\hat{s}s_m} \tag{4}$$

which says that each spouse chooses his retirement and benefit claim status conditional on his expectations about what his partner will do given his choice. In general, it will not be the case that $\hat{\Upsilon}^i$, and corresponding choices $\{ls_i, ss_i\}$ will correspond to the planner's choices, which immediately suggests that allocations are different when spouses behave strategically. Note also that, in the non-cooperative world, $\hat{\Upsilon}$ is also the value function that enters the non-naive household planner's problem U, which means that allocations are likely to be different even when individuals end up following retirement paths that do not deviate from the planner's. In fact, though not implied by the previous expressions, if the planner is fully non-naive about spousal "outside options" in a given year, the household problem boils down to one in which, in addition to satisfying B_M, the household planner also must satisfy the conditions that $\Upsilon_f = \hat{\Upsilon}_f$ and $\Upsilon_m = \hat{\Upsilon_m}$.[12]

To make the problem tractable, we exploit the fact that both the decision to claim benefits and to quit a career job are binding (irreversible) decisions. We further assume that the decision to apply for SSDI benefits is binding once taken in a period. For younger workers, this is a not a realistic assumption, but most SSDI applications are made by men and women in their late 50s and early 60s who may not easily be able to transition back into non-career work as job search durations

[12] The problem presented is somewhat similar to a married individual's optimization problem in an economy with divorce as an outside option (see, for instance, Mazzocco and Yamaguchi (2007) or Chiappori et al. (2002)), with one crucial difference: once it becomes the optimal choice for i to divorce his spouse, he no longer cares about how $-i$ responds. However, strategic retirement decisions effectively suggest a form *internal* threat point as proposed in Lundberg and Pollak (1991). In this case, spouse i's optimal choice depends on $-i$'s optimal response to his own retirement decision. This is true whether or not i's non-cooperative decision coincides with the planner's allocation.

tend to increase with age should application no longer appear optimal. This assumption allows us to assume that ss and ls decisions are made sequentially, with a unique $\{ss^m, ss^f\}$ given for each ls decision: a spouse claims benefits if her payoff to doing so under the current allocation is higher than her payoff from non-claiming, with no Nash indeterminacies (see, for instance, the problem in Gustman and Steinmeier (2004)). We therefore focus on the per-period one (for NC and R spouses) or two (for at least one C spouse)) Nash game that spouses play over $\hat{ls}$ (the solution to 4). We first impose three further restrictions on the outcome of the game:

1. Assumption #1: Spouses do not play mixed strategies.

2. Assumption #2: If no (pure) Nash equilibrium exists in a subgame, the spouses revert to the planner's (jointly optimal) solution for that subgame.

3. Assumption #3: If more than one (pure) Nash equilibrium exists in a subgame, the players revert to the planner's solution for the subgame.

In a couple consisting of non-career workers, or retirees under 75, the partners play a simple 2x2 static Nash game, and assumptions #1-#3 together imply that the couple deviates from the λ-weighted jointly optimal outcome only in the event that at least one spouse has a dominant strategy over the feasible choices of ls (with its implied ss), and that the Nash equilibrium outcome that results differs from the planner's. Figure 3 shows four types of outcome that can arise: the first two give rise to a potential deviation and the other three do not. In each case, the payoffs are given by $(wife, husband)$, with A>B>C>D for the wife and a>b>c>d for the husband. Figure 1(a) shows the case in which the NC husband retires regardless of the wife's action. The wife retires only if the husband also does so: the Nash equilibrium is $\{R, R\}$. Subfigure 1(b) shows a case in which both partners have dominant strategies to retire, leading to outcome $\{R, R\}$. In both 1(a) and

1(b), the Nash solution to the 2x2 game is likely to differ from the planner's ls allocation. In 1(b), the outcome is actually (ex-post) Pareto inefficient, in that both spouses would be better off if they played the cooperative solution $\{NC, NC\}$.[13] In figure 1(c), both $\{R, R\}$ and $\{NC, NC\}$ are Nash equilibria. In figure 1(d), there is no pure strategy Nash equilibrium. Both of these scenarios lead to implementation of the planner's solution.

Assumptions #1-#3, that prevent deviant outcomes in figures 1(c) and 1(d),are motivated by three ideas: (1) the single-shot Nash game played once per period is actually an approximation to an uncoordinated repeated static game that is played within a period until a stable outcome is reached; (2) retirement is not a binding state for individuals under 75; at any point, an individual can simply opt back into non-career work so long as he is not in state ss state A; (3) spouses find uncertainty over their ls status within a model period costly and therefore avoid it, converging instead on what they know to be the jointly optimal outcome.

The 2x2 case holds for couples whose members are either both currently retired or in non-career jobs. The problem is slightly more complicated if at least one spouse is in a career job, $ls^i = C, i \in \{f, m\}$. This is because, unlike applying for benefits, retiring from a career job (action Q) *is* a binding decision that cannot be undone. The career worker therefore controls the nature of the play, creating a sequential game in which she moves first. Figure 4 shows potential solutions to the game that arise when the wife is a career worker and the husband a bridge worker. The wife solves this game by backward induction, first resolving the 2x2 game that results if she retires, and then comparing this outcome to that which occurs if she does not retire. If the wife is better off playing the 2x2 game with her husband, she quits her career job. Otherwise she remains in her career job and her husband chooses his best ls choice conditional on her choice. In the first panel

[13] The possibility of ex-post inefficient outcomes, in addition to the ex-ante inefficiency introduced by the possibility of non-cooperation, is one way in which the non-cooperative model differs from a collective model with divorce hazards since a divorce is never ex-post Pareto inefficient.

Figure 3: Non-cooperative outcomes for non-career workers

Wife NC \ Husband NC	NC	A
NC	(A,c)	(D,a)
A	(B,d)	(C,b)

(a) Husband dominant retirement strategy

Wife NC \ Husband NC	NC	A
NC	(B,b)	(D,a)
A	(A,d)	(C,c)

(b) NE: Husband and wife dominant retirement strategies

Wife NC \ Husband NC	NC	A
NC	(A,a)	(D,b)
A	(B,d)	(C,c)

(c) Multiple NE: Revert to planner's solution

Wife NC \ Husband NC	NC	A
NC	(B,a)	(C,b)
A	(A,d)	(D,c)

(d) No pure NE: Revert to planner's solution

of figure (4), the wife's payoff if she does not quit is D, since in this case her husband plays A. However, if she quits, the couple will converge on the Nash equilibrium $\{NC, NC\}$, which raises her payoff to B. She will therefore quit. Note that the outcome of the supergame is ex-post inefficient, even though the solution to the subgame is not: regardless of λ, the planner would always prefer $\{C, NC\}$ to $\{NC, NC\}$.

By contrast, in the situation represented in figure 2(b), the wife will not quit, even though she would prefer her husband make a different ls choice while she remains in her career job (and her husband's choice may violate the marriage contract). If she quits, she is guaranteed a payoff $C < B$. In figure 2(c), the wife's quit choice is ambiguous. In the 2x2 game, the couple converges on the outcome dictated by the marriage contract, which could be either $\{NC, R\}$ or $\{A, NC\}$, depending on the value of λ. The wife's choice will be Q if the 2x2 outcome is $\{A, NC\}$, and NQ otherwise.

A symmetric game, and similar logic, governs the optimal retirement of a career husband with a bridge-working or currently retired wife.

Figure 4: Non-cooperative outcomes for a career and a non-career worker

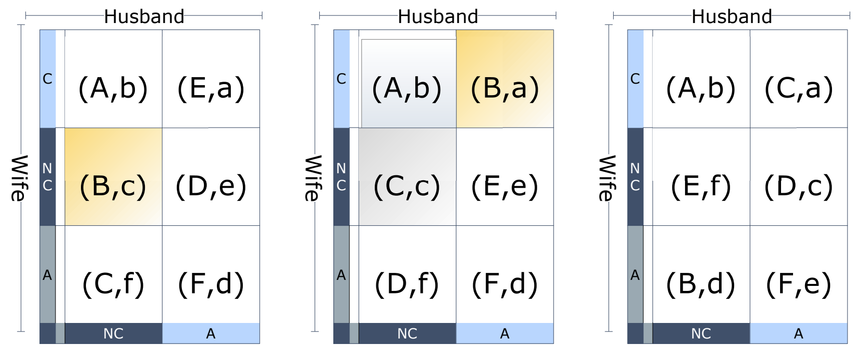

(a) Wife quits (b) Wife does not quit (c) Wife's quit depends on λ

The final game that arises in the model is played between a career husband and a career wife. The resulting 3x3 Nash game can be conceived as a two-stage sequential game. At the first stage,

Figure 5: Non-cooperative outcomes for career workers

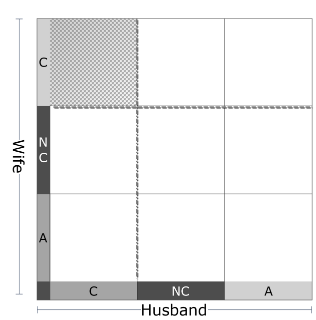

both partners chose whether or not to quit (Q) or not quit (NQ) conditional on their spouse's quitting decision. If the husband, for instance, plays NQ at the first stage, the wife faces a choice over the 2x3 game induced by quitting (the area under the thick dotted line passing to the south of the hatched $\{C, C\}$ square in figure 5) or remaining in $\{C, C\}$. The husband has a symmetric choice set when the wife plays NQ at the first stage. If the husband quits, the wife has a choice between the 3x2 game described in the previous paragraph (playing NQ) and the 2x2 game described at the beginning of this section (playing Q). If she is indifferent between these choices, that must mean that quitting is the outcome of the 3x2 game between a career wife and non-career/applicant husband given x_M. In this case, $\{Q, Q\}$ is the solution to the first stage of the game and the couple locates and plays the 2x2 game defined in the box at the bottom right of figure 5. Thus, indifference over outcomes in the first stage game is resolved in favour of quitting. If neither spouse finds it optimal to quit conditional on the other not quitting, then $\{NQ, NQ\}$ is the solution to the game.

Finally, with the given assumptions, it is easy to calculate the $\{\hat{ls}, \hat{ss}\}$ outcomes for different points in the couple's state space x_m and household intertemporal allocation $\{c, a'\}$, given x_m. The planner's problem problem is therefore complicated by the existence of incentive comparability constraints for each spouse, requiring that $\Upsilon_f = \hat{\Upsilon}_f$ and $\Upsilon_m = \hat{\Upsilon}_m$. Under our assumptions, this implies that the household allocation is itself a Nash equilibrium given current states and age of the household. Obviously, this restricts the feasible choices set for the household not just at retirement but throughout the pre-retirement (and earlier) years. In particular, it will highlight conflicts between the spouses over rates of effective time preference (effectively, the rate at which household resources are consumed), as well as conflicts over which spouse sacrifices leisure to increase household income.

3 Numerical implementation

In this section, we outline the estimation and calibration of the numerical model economy: the processes for risk faced by agents and the policy environment they face. Individuals are subject to three main types of risk: (1) productivity (wage) risk; (2) health risk, which has both a physical and financial risk component; and (3) mortality risk that depends on health and age.

3.1 Life cycle effects and sources of risk

3.1.1 Career and non-career jobs and productivity shocks

Workers who are not currently retired may be in either a career- or a non-career job. In the PSID and HRS, we identify as "career workers" those who satisfy at least one of the following conditions:

1. are employed at the time of the interview, and have at least three years of full-time tenure in their current job
2. have worked full time, have not experienced a spell of unemployment, and have not changed occupations over a three year period.

The second condition allows us to identify as career workers those who may have engaged and acted on the results of on-the-job-search. The restriction that workers move into new jobs within their own industries is consistent with the evidence in Kambourov and Manovskii (2009), who show that human capital tends to be occupation- rather than individual- or firm-specific. Roughly 50% of workers in their 40s and 50s are identified as career workers; this share falls to around 2% by age 70.

Because of the difficulty involved in obtaining structural estimates of the evolution of wages conditional on selection both into the labor market and between career and non-career jobs, we estimate the following reduced-form equation separately for men and women:

$$
\begin{aligned}
\ln w_{it}^{g} = \beta_0^g + \beta_1^g age_{it} + \beta_2^g age_{it}^2 + \beta_3^g age_{it}^3 \\
+ \beta_4^g career_{it} + \beta_5^g career \times age_{it} + \epsilon_{it}
\end{aligned} \tag{5}
$$

$$
\epsilon_{it} = \mu_{it} + \nu_{it}; \qquad \mu_{it} = \rho^g \mu_{it-1} + \varphi_{it}; \qquad \nu_{it} \sim N(0, \sigma_{\nu}^{g^2}); \qquad \varphi_{it} \sim N(0, \sigma_{\varphi_{it}}^{g^2})
$$

for $g \in \{f, m\}$. "career" is a binary variable taking the value 1 if the worker is a career worker and 0 otherwise. The resulting 12 reduced form parameters β, autocorrelation coefficients ρ and variances σ_{ν}^2 for $g \in \{m, f\}$ are then used as targets in the estimation of the model.

To run this regression, we first pool individuals from the 1976-1993 waves of the PSID with the first eight waves of the HRS (1992-2006), and keep all individuals born between 1932 and 1942. This allows us to follow this cohort back in time to a minimum age of 35. Since the PSID and HRS represent different populations (the PSID represents the whole US population circa 1969), we use data from the overlapping year, 1992, to estimate the likelihood of any individual in the PSID sample also appearing in the HRS, and assign the resulting propensity scores as weights on the PSID observations.

3.1.2 Health and medical costs

Health shocks increase the disutility of labor by increasing the amount of effective time an individual loses, by a factor $\delta(h)$ while completing a unit of career or non-career work. They also impose monetary costs (χ) on individuals that vary across career status, insurance status and age. We estimate health transition matrices based on self-assessed health (individuals are in bad 'b' health if they report being in poor health in the data and good 'g' health

Table 2: Wage equation parameters: Career and non-career workers

	Men	Women
$cons$	2.78	2.62
	(.051)	(.043)
age	$6.37E^{-3}$	$9.29E^{-3}$
	$(3.09E^{-3})$	$(7.23E^{-3})$
age^2	$9.67E^{-4}$	$4.64E^{-4}$
	$(4.84E^{-4})$	$(4.55E^{-4})$
age^3	$1.11E^{-5}$	$5.54E^{-5}$
	$(6.21E^{-7})$	$(8.87E^{-6})$
$career$	.0715	.557
	(.069)	(.069)
$career \times age$	.002	-.006
	(.001)	(.001)
ρ	.962	.932
σ_ν	.157	.210

otherwise) reported in the HRS and PSID and data on out of pocket medical expenditures in the HRS. Individuals who their current health level as 'poor' are assumed to be in poor health and resultingly work-limited. Unhealthy individuals are subject to out of pocket medical expenditures that may be high (H) or low (L). For singles, we estimate three-state Markov transitions between states $\{g, bl, bh\}$ states by multivariate OLS of current on lagged health status. For marrieds, we estimate a seven-state transition matrix for the couple, with states $\{ggL, gbL, gbH, bgL, bgH, bbL, bbH\}$ which allows for arbitrary correlation between the husband and wife's physical conditions and level of insurance over time. Transition matrices are estimated for each household age j using all relevant observations in an eleven year interval with j at the center. We allow the size of the δ shock associated with each state to vary over the life cycle and is estimated to match observed labour supply by age and health status, with δ normalized to 1 in the good health, low-expenditure state. The medical expenditures in each state vary by whether the individual is in a career job; whether he is

65 or over, in either which case he is assumed to be covered by medical insurance. In each cell, 90% of poor health individuals are in the low cost category with the cell mean average medical expenditure payment and the remaining 10% are in the high cost category with corresponding mean cell payment.

3.1.3 Mortality risk

Following Rivas (2010), we calculate disabled and non-disabled mortality profiles from PSID data and the National Vital Statistics Report (2003) using Bayes Rule. The healthy (g) and unhealthy (b) profiles can be written as:

$$\begin{aligned} \varphi_{g,i}^{j} &= Pr(Surv|g) = \frac{\widehat{Pr}(Surv)\widehat{Pr}(g|Surv)}{Pr(g)} \\ \varphi_{b,i}^{j} &= Pr(Surv|b) = \frac{\widehat{Pr}(Surv)\widehat{Pr}(b|Surv)}{Pr(b)} \end{aligned} \tag{6}$$

where hats indicate that the information comes from the PSID micro data. Figure 6 plots the relative survival probabilities of healthy and disabled males and females using the subjective PSID measures of disability defined below. The figure, and corresponding regressions (not reported) show that (1) the importance of subjective good health increases with age and (2) that the effects of subjective good health on survival probabilities are significantly more important for men than for women at most ages.

Figure 6: Health-conditional survival indexes for males and females

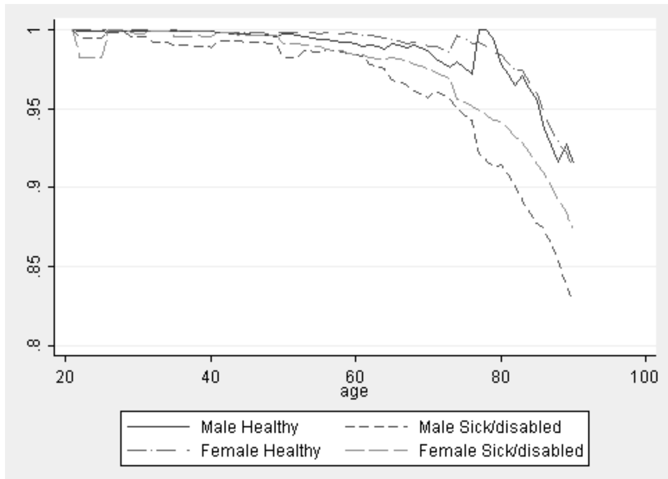

3.1.4 Preferences

We adopt Cobb-Douglas preferences in period-utility. Evidence from the vast literature on household consumption and labour supply has generally found strong evidence that consumption and leisure are inseparable in utility (see, e.g. Blundell and Walker (1986)). In the standard and non-cooperative models, individual utility of marrieds is given by:

$$u_i(c_i, l_i) = \frac{[\psi_C(\frac{c_i}{\tilde{\eta}_{ms}^{-1}})^{\gamma_i}(l_i)^{1-\gamma_i}]^{1-\omega_i}}{1-\omega_i} \tag{7}$$

If preferences demonstrate complementarity in spousal leisure, period utility is given by

$$u(c_i, l_i) = \frac{[(\frac{c_i}{\tilde{\eta}_{ms}^{-1}})^{\gamma_i}(l_i + I_{rr}\psi_L(ms)\min[l_i, l_{-i}])^{1-\gamma_i}]^{1-\omega_i}}{1-\omega_i} \tag{8}$$

In the above expressions, γ measures the consumption share in utility, ω is the coefficient of intertemporal risk aversion and $\tilde{M}$ consumption gain from marriage. Individuals benefit also from a bequest motive, denoted bq which, following the standard formation in the literature, and the specific form from De Nardi et al. (2010), takes the form $bq = \phi_{ms} \ln(a + \$500,000)$, whose curvature has been shown to be sufficient to replicate the distribution of bequests by wealth status in U.S. data. The motive is the same across genders but differs by whether an individual entered the model in a married household or single.

3.1.5 Model Estimation

Estimation of the models is performed through exactly-identified simulated method of moments (SMM). In addition to the parameterization of gender- and career-status specific wage growth, the program described above gives a total of fourteen parameters to estimate for the fully cooperative and non-cooperative models. These parameters are listed, along with the corresponding targets, in table 3. Tastes for consumption γ determine the average labor supply of healthy married men and women up to age 59. Time δ costs of poor health are allowed to vary with age, and are targeted to match, in conjunction with the out-of-pocket health costs taken from the data, the evolution of labor supply over the life cycle. Moving down the table, the gender-specific discount rates β are set to generate wealth-to-income ratios for all pre-retirement age households with a female and a male member respectively. Bequest motives bq are pinned down by the median asset holdings of retirees (roughly, individuals 62 and over). Acceptance rates into SSD are set to match the share of individuals between 40 and 65 in each health state who begin receiving benefits in a given year; we assume no difference in acceptance rates based on health insurance status. The rejection cost c_A controls the gross number of applicants. We target a number that, combined with the ps, leads to an overall rejection rate by disability spell[14]—of around 50%, which Gruber (2000) and Bound et al. (2002) separately cite as the prevailing post-appeal rejection rate of SSD applicants. The transitions from NC into C employment are targeted by gender to replicate the share of pre-retirement men and women 50-59 in career work. Finally, $\tilde{\eta}_M$, the degree to which couples "convert" their joint consumption into individual consumption determines the differential saving rates of marrieds and single households. $\tilde{\eta}_M = .6$ implies that married households do benefit from economies of scale in consumption.

[14]Because individuals in the simulated economy can continuously apply for benefits, while appeals in the U.S. are actually limited to three after initial application, the rejection rate in the model refers to all "spells" of applications—i.e. groups of one or more consecutive application periods. In practice, less than 1% of simulated SSD application spells last more than four consecutive periods.

Table 3: Parameterization of the benchmark (cooperative) and non-cooperative modesl

Parameter	Benchmark	Non-coop	Target
γ_f	.24	.32	Share of pre-retirement women in R
γ_m	.27	.30	Share of pre-retirement men in R
β_f	.905	.902	HH wealth-income ratio for homes w female member
β_m	.913	.898	HH wealth-income ratio for homes w male member
bq^S	$2.2e^{-7}$	$1.3e^{-7}$	HH wealth-income ratio for singles ≥ 62
bq^M	$4.5e^{-7}$	$5.6e^{-7}$	HH wealth-income ratio for marrieds ≥ 62
δ_j^f	$1.05(1.0 - .020j + .0023^2)$	$1.05(1.0 + .0046j + .00023^2)$	Share of 65/70 yr old women in R
δ_j^m	$1.05(1.0 - .013j + .0018j^2)$	$1.05(1.0 + .0011j + .00014j^2)$	Share of 65/70 yr old men in R
$p(\delta)$	.57	.59	Total share of SSD application ending in success
c_A	.0005	.0004	Share of poor-health individuals 49-60 applying for benefits
p_C^f	.91	.88	Share of working age women in C
p_C^m	.71	.78	Share of working age men in C
$\tilde{\eta}_M$	.60	.60	Saving of marrieds/Saving of singles

3.2 The policy environment

3.2.1 Social security

Under the rules of the U.S. Social Security system, individuals "save" for retirement by effectively paying a share of their per-period wage earnings to the Social Security Administration. Payments into the system by current workers finance the benefits of the current generation of retirees. The contributions of workers into the system function as IOUs from the government. When the current generation retires, their benefits will in turn be financed by the next generation of workers. Social Security is primarily a retirement program. However, it offers insurance to workers who become incapacitated in the form of disability benefits and to the spouses of workers who die prematurely in the form of spousal benefits.

We next outline the major features of Social Security that are captured in the model:

Benefit accrual and determination. SSR and SSD benefits for new applicants are determined as a function of previous Social Security contributions. The Average Indexed Monthly Earnings, or AIME, is the average monthly wage income from the worker's applicable work history, which comprises the 35 highest-earning years, or the highest 80% of earnings-years for applicants to SSD, up a maximum of five excluded years. From the AIME, the Primary Insurance Amount (PIA), and the actual retirement-age or disability benefit, b, are determined from the following well-known formula:

$$
\begin{aligned}
PIA &= 0.9\min\left(0.2\overline{w}, E\right) \\
&\quad +0.32\max\left(\min\left(E-0.2\overline{w}, 1.3\overline{w}-0.2\overline{w}\right), 0\right) \\
&\quad +0.15\max\left(E-1.3\overline{w}, 0\right) \\
a_t < \overline{a}^{mt}: &\quad b_{.} = \max(PIA, \underline{b}) \\
a_t > \overline{a}^{mt}: &\quad b_{.} = PIA
\end{aligned}
\tag{9}
$$

where $\overline{w}$ is the average per-capita wage earnings in the economy, E is shorthand for the AIME, and $\underline{b}$ is a floor benefit, equal to \$151 weekly for an individual (\$222 for a couple). This benefit is administered by the Social Security Administration as a separate program (Supplemental Security Income: SSI) for workers and non-workers who do not qualify for regular Social Security benefits or whose accrued earnings are too low.[15] Since SSI is means-tested as opposed to work-tested, $\overline{a}^{mt}$, equal to \$4000 for a single household head or \$6000 for a couple, is the maximum level of asset holding for which an individual or household can be guaranteed the floor benefit.[16]

State space limitations prevent an accurate accounting of which years should be included in benefit determination in the simulated economy.[17] Instead, we exclude earnings from all

[15] The SSA requires that individuals earn at least 20 "credits", where one credit is equal to wage earnings of \$4200 in current dollars, over 10 years prior to application. Some exceptions are made for younger workers.

[16] These figures are adjusted to exclude the average share of assets in housing wealth, which is not included in the determination of SSI resources.

[17] Huggett and Ventura (2000) discuss a similar technical difficulty in their calibrated analysis of the Boskin

years before age 26 and over 61 in the accumulated calculation, which gives five years of accumulation before individuals in the simulated economy first apply for disability benefits. Given the shape of the life cycle profile of earnings, these tend to be the lowest-earning years. The benefit accumulation calculation for workers is:

$$\begin{aligned} j < 60: \qquad E_t &= \frac{E_{t-1}(j_t - j_1) + w_t n_t}{j_t - j_2} \\ j \geq 60: \qquad E_t &= \frac{E_{t-1}(j_t - j_1) + w_t n_t}{j_t - j_2} \end{aligned} \tag{10}$$

When benefits are determined, the AIME is adjusted for growth up to the year the individual turns 60. In this stationary model, the growth rate is zero.

Disability benefit eligibility and receipt. In order to receive disability benefits under SSD, applicants must pass an eligibility test which insures that the disability is "total" in the sense that it precludes all "substantial gainful activity" and is expected to last at least 12 months. Slightly more lenient eligibility rules apply to individuals over 55. If rejected for benefits, an applicant may appeal up to four times, to four different levels of SSD adjudicators, a process that can take several years.[18] In separate studies, Bound et al. (2004) and Gruber (2000) report a final rejection rate of about 50% of initial applicants to the program, with a first-time applicant rejection rate around 67%.

In the model, these provisions are captured by acceptance rates that vary across age (pre- and post-55) and the disability/health index of the applicant. We use the shares of new SSD

Social Security reform proposal.

[18]An initial Request for Reconsideration goes to the SSA, after which the rejection may be appealed to an Administrative Law judge, to the Social Security Appeal Council, and finally to a federal court.

recipients in each self-assessed health/disability state in the PSID, and the share of benefit recipients older than 55 among the recipient population, to estimate these acceptance rates. Individuals older than 35 may apply for benefits in any period they want, from any d state, but at a cost of forgoing a year of work in order to demonstrate incapacity. Since the decision process for an application takes five months in the real world, successful applicants in the simulation receive only 60% of their benefit in the year they are accepted.

Early and delayed retirement. Under the SSA's rules, individuals may retire at any age between 62 and 70, subject to an adjustment of benefit size that roughly equates the expected discounted stream of benefits across retirement ages. For early retirees, benefits are reduced by 5/9 of a percent for every month before the full retirement age of 66. The factor of adjustment for later retirees is 8% of the PIA per year. As well, individuals can continue to replace lower-earning years in the calculation of their AIME until they formally retire. Both of these effects - adjustment and continued accumulation - are captured in the model.

Survivor and spousal benefits. Under current SSA rules, surviving spouses whose own Social Security benefit is smaller than their partner's receive their partner's benefit (unadjusted by the spouse's own benefit claiming date) until their own death. As well, secondary earners during the working life of a married household are entitled to the greater of either their own accumulated benefit or one half of their partner's PIA, adjusted to the second earner's own clam age, upon retirement. Both of these features of the program are captured in the simulated economy.[19]

[19]Under the rules for spousal benefits, spouses are not entitled to receive a half share in increases in the main earner benefit due to delayed retirement by the couple. Additional rules may provide benefits for the

Earnings test. All individuals become eligible for regular social security benefits at age 62. Individuals who continue to work while claiming benefits are subject to a benefit adjustment process that simultaneously "claws back" current benefits while increasing the individual's PIA in future claiming years due to the payroll contributions from current earnings. In 2000, the provisions of this benefit "earnings test" changed so that workers after the normal retirement age of 65 (now 66) do not face any benefit clawback while workers under 65 face a 50% clawback on earnings above a threshold level. Since the average individual in our sample reaches age 65 in 2002, we adopt the post-policy reform SSA rules. Benefit recipients 65 and over do not face a direct clawback,[20] though the taxation of benefits is likely to change (see Section 3.2.2). Recipients between 62 and 64, however, are subject to a deferral in their present average weekly benefit of $1 for every $2 earned above $250 (in 2006 USD).

3.2.2 Non-Social Security policy and institutions

We also incorporate the following non-Social Security features of the U.S. policy environment.

Taxes. Policy in the model is designed to reflect several features of the current U.S. policy environment in addition to Social Security. We model a progressive income tax with % rates of $\{10, 15, 25, 28, 33\}$, levied on (average weekly) income above $\{\$358, \$679, \$1660, \$2987, \$4364\}$ for marrieds and $\{\$179, \$340, \$830, \$1756, \$3470\}$ for singles. These numbers are based on

pre-majority age or disabled children of deceased parents. These rules are ignored in the simulated economy.

[20] In practice, individuals reaching normal retirement age *during the current year* are subject to a much smaller clawback of $1 for every $3 earned above a threshold roughly three times the pre-normal-retirement age threshold. We omit this detail.

the following assumptions: (1) all married individuals file jointly; (2) all filers claim the standard deduction and personal (but not dependent) exemptions; and (3) that only 2008 federal rates apply. Further, we follow a standard convention in the life cycle literature by assuming a 100% estate tax (no bequests), and a flat-rate consumption tax of 5.5% as in Imrohoroglu et al. (2003). We treat capital and labour income identically in the tax calculation, ignoring potentially favourable tax treatment of retirement savings or capital gains. The payroll tax is 15.3%, which has the combined employer-employee OASDI and Hospital Insurance (HI) payroll tax rate in the U.S. since 1990. It is levied on weekly average earned income up to $2040.

Social security benefits are taxed at a special rate modeled on 1993 federal legislation. Up to 50% of SSR/SSD benefits are taxable as income if total non-Social Security income plus 50% of benefits (called "adjusted income") fall above a certain threshold ($400 for singles; $640 for marrieds). In this case, taxable benefits are then the lesser of 50% of total benefits and the difference between adjusted income and the threshold. In 1993, a second threshold ($680 for singles; $880 for marrieds) with an associated rate of 85% was added. For individuals with post-retirement incomes higher than the second threshold, benefits subject to taxation equal the lesser of the amount calculated using the brackets and 85% of total SSD/SSR benefits. These features of benefit taxation are captured in the model. Revenues from taxation of Social Security benefits are added to the Social Security Trust Fund along with payroll taxes, which is a relevant detail only in the general equilibrium version of the model.

Because the focus is mainly on older couples (those over 45 years of age), we do not formally model the other main transfer programs: the earned income tax credit (EITC), food stamps, public housing, or Temporary Assistance for Needy Families (TANF). With the exception of housing provision until Title 8 and state programs, most of these policies favor households with children, which are necessarily not the focus of the model. We set a minimum level of income in each period equal to one fifth of the median household income (for marrieds and singles respectively) which is received and consumed if other sources of income are not sufficient to match it.

Minimum consumption floor Medicaid. Individuals (of any age) who are unable to meet their medical expense requirements out of their own wealth stocks qualify for Medicaid, which pays their expenses and leaves them with a base level of consumption, $\underline{c}$ ($\underline{c}$ for married households). Following De Nardi et al. (2010), we set $\underline{c}^M$ to \$300 for married couples and \$250 for singles, which is substantially less than most households' Social Security benefit entitlements.

4 Results

We now turn to assessing the implications of the model described in the previous section and considering its implications for retirement and benefit claiming. The following results are quite preliminary, but give a flavor of the model's predictions and implications for the retirement behavior of couples.

Figure 7: Cross sectional benefit claiming hazards: married men and women

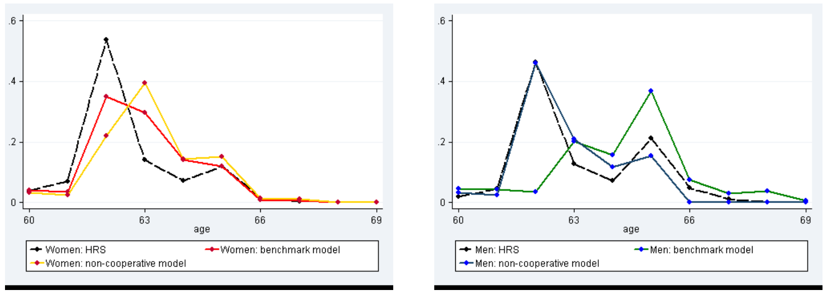

(a) Women (b) Men

Figure 8: Cumulative benefit claiming hazard for married men and women

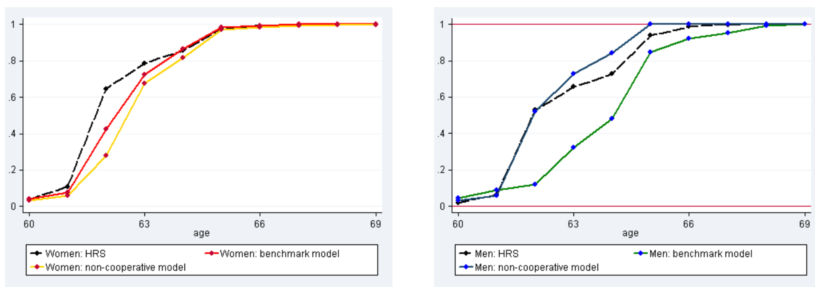

(a) Women (b) Men

4.1 Cross sectional benefit claiming patterns

Figures 7 and 8 show the cross-sectional benefit claiming hazards for married women (left) and men (right) in the standard cooperative model and the strategic model developed in this paper. Figure 2 shows the hazards by age between age 60 and 69 while figure 8 shows the corresponding cumulative hazards: the share of married non-db pension workers of each gender who have claimed benefits by the given age. In each figure, the black line shows the results from our HRS sample.

For married women, the fits of the two models are similar, which is not surprising. Most married women have relatively little incentive to delay benefit claiming since doing so does not increase household Social Security wealth so long as the higher-earning husband delays claiming. About 45% of wives expect to benefit from the spousal benefit while their husband lives; that is to receive a benefit equal to 50% of his. As in Rivas (2010), this incentive induces married women to claim benefits early. They also have no incentive to delay claiming in order to insure against long life, since most likely they will receive their husband's benefit in very old age. In fact, the "non-cooperative" model implies slightly *later* claiming by married women than the cooperative model or the data suggest, largely to counter the negative wealth effects of their husbands' much earlier benefit claiming. Nevertheless, both models predict that married women claim within two years of benefit availability, and effectively all married women have taken up benefits by age 66.

The story for married men is quite different. Married men in the fully cooperative model delay benefit claiming into their mid 60s, with the largest spike in claiming at age 65 and a positive share of men delaying their claims right up until age 69. This model provides a poor fit to the data on observed benefit claiming hazards among married men. In the non-cooperative retirement model, by contrast, the majority of men, around 52%, claim at age 62, as soon as benefits are available, with a small residual spike at age 65 and virtually no husbands claiming after that date.

What accounts for the difference between the models? When benefit claiming and retirement decisions are made separately, there is no reason for declining health to induce married men to change their optimal benefit claiming date from that which maximizes the couple's total expected Social Security wealthy, so long the household has accumulated sufficient assets and/or a large enough spousal PIA to allow for smoothing consumption through the 60s, following the husband's retirement from the labour force. A couple solving a stan-

Figure 9: Transitions into retirement for married men: Data and models

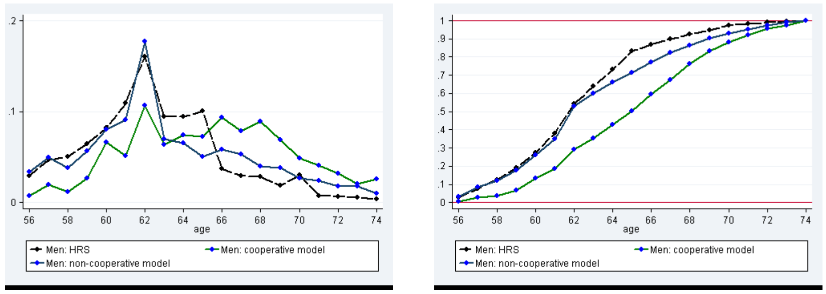

(a) PDF (b) CDF

dard joint optimization problem therefore delays the husband's age of benefit claiming. As well, a residual benefit of commitment to the jointly optimal solution within households is that households are also better able to save for retirement. In the cooperative benchmark model, only 23% of couples have fewer than $10,000 in liquid assets, compared to 43% in the non-cooperative model (and roughly 40% in the data.) When households are wealth constrained, husbands can effectively force the rate of resource consumption forward in time by substituting Social Security consumption for their own labor income by retiring from the labor force and claiming benefits.

To further motivate the retirement implications of non-cooperative behavior, figure 9 shows the observed hazard and cumulative hazard for final labor market exit for married men in the two models and in our HRS sample. The non-cooperative model (the blue line) replicates the spike in labor market exit at 62 among married men observed in the data, and the earlier cumulative retirement levels among married men, consistent with the story outlined above.

4.2 Joint retirement

We next turn to examining the models' implications for timing of benefit claiming, and final labor force exit, *within* households. Table 4 shows the same distributions from the two rows of table 1 and the corresponding distributions generated by the models. The patterns for benefit claiming are consistent with the cross-sectional individual evidence displayed above in which husbands' benefit claiming is delayed in the cooperative (benchmark) model. The cooperative model predicts that husbands typically claim benefits later than their wives; by contrast, the strategic model predicts that spouses claim roughly concurrently or, even, that wives claim slightly later than their husbands. This is because most couples jointly maximize their Social Security wealth when the husband delays claiming: the relatively high saving rates in cooperative households make this possible. Husbands in non-cooperative households, however, are not only interested in maximizing wealth, but it increasing their own share in the consumption of wealth, which implies moving a share of consumption into the present. The relatively low saving rates of low-income married households, due in large part to non-cooperative interactions between spouses earlier in the life cycle, means that husbands can increase current consumption by claiming early and concurrently leaving the labor force to consume more leisure.

Corresponding results for the intrahousehold timing of final labor market exit are less clear. Neither model performs very well, though the strategic model, which *overstating* the amount of joint labor market exit, performs better in the sense of absolute deviations from cross-sectional rates of intrahousehold claim date differences in the data. Figure 9 showed that the non-cooperative model does a good job of predicting the cross sectional dates of

Table 4: Husband benefit claiming age (retirement age) minus wife's in years

Difference in years	HRS 1998-2006		Benchmark model		Non-coop model	
−6	2.9	(14.7)	0.3	(12.5)	0.7	(13.9)
−4	14.9	(14.2)	2.2	(8.6)	6.9	(7.2)
−2	14.7	(8.0)	5.6	(4.6)	34.0	(5.1)
0	58.9	(37.5)	32.0	(17.8)	51.4	(50.9)
+2	4.8	(6.8)	39.3	(9.0)	5.7	(9.6)
+4	4.9	(7.2)	24.6	(17.5)	1.5	(4.9)
+6	1.3	(11.6)	6.5	(30.1)	0.5	(6.3)

final retirement from the labor force among men. By contrast, the optimal strategy for many cooperative households is for one spouse (usually but not always the husband, depending on the relative wage rates and health of the spouses) to work through his 60s while delaying benefit claiming. This failure of the cooperative model, however, is at least partly a red herring due to the absence of leisure complementarities, which have been shown by several authors to be important determinants of joint labour market exits within married couples, and which are addressed below [forthcoming.]

4.3 The roles of medical expenditure and bequests

Our models allow both for increasing medical expenditure risk and for bequest motives as described in section 4. In this subsection, we look at the implications of both these features of the model for explaining our results, focussing on the benefit claiming behaviour of married males in the two models. Figure 10 shows the predicted benefit claiming when the model is recalibrated excluding individuals' bequest motives. The top row of the figure shows the pdf and cdf of the benefit claiming hazard between ages 60 and 69 for all households in the sample. The second row shows the results decomposed between the 50% of higher-wealth and 50% of lower wealth households, where a household is categorized as high or low wealth based on whether it is above or below the median deflated wealth of all married households whose head is the same age. In the data, the categorization is made at either age 58 or the first age household is observed in the sample; for all the model is is made at age 58 for all households.

The results show three things: (1) both the benchmark and non-cooperative/strategic models perform worse when there is no bequest motive; (2) the non-cooperative model still outperforms the cooperative model even with no bequest motive; and (3) the failure of fit is concentrated among wealthier households for whom bequest motives are most important relative to other wealth and labour market considerations. For lower-wealth households, the non-cooperative model still performs quite well. The difference is less pronounced for the cooperative model, in large part because households in the lower 50% of households are wealthier in an absolute sense and so less constrained. In wealthy cooperative households, couples build up bequeathal wealth by having at least one of the partners continue to work which makes joint benefit claiming unattractive since little of the husband's income is received right away due to the Social Security earnings test.

Next, figure 11 repeats the analysis using a version of the model in which the variance in health costs is eliminated, by restricting health expenditures to be equal across h states and to be held to the level of a healthy 65 year old (for singles) or household with two healthy members and a 65 year old head (marrieds). While the role of health expenditure turns out to be less important than the bequest motive for explaining benefit claiming decisions of husbands in their 60s, the three main findings are similar to those listed for the model without

Figure 10: Benefit claiming by married men: No bequest motive

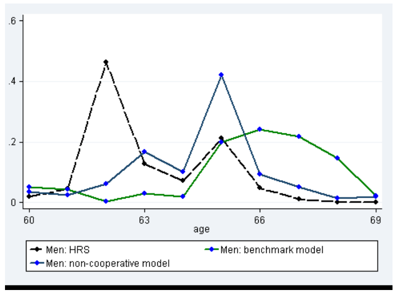

(a) Distribution by age

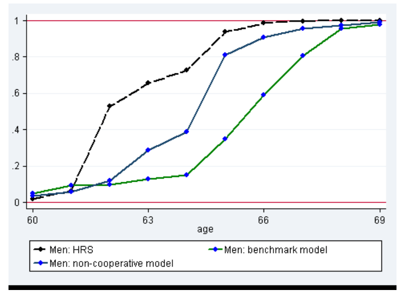

(b) Cumulative hazard

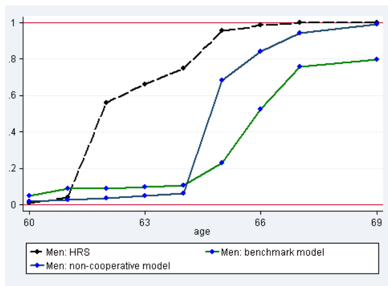

(c) High wealth households

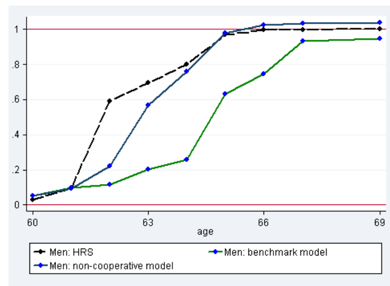

(d) Low wealth households

bequest motives. The difference is that, this time the model performs relatively worse for poorer than for richer households, in particular substantially underpredicting the spike in retirement at age 62. The reason is that that absence of very large health expenditures late in life makes liquid asset wealth a better substitute for Social Security wealth among households in their 60s. Households with modest asset holdings, in which the husband claimed very early, would simply substitute larger asset holdings for the lost expected Social Security wealth, leading to a small *reduction* in current consumption (due to the negative wealth effect), which is suboptimal for both partners. Thus only husbands in households that are actually wealth-constrained (in the sense of a binding borrowing limit) will actually see strategic retirements accompanied by benefit claims among husbands, most of which take place when the wife, rather than the husband, turns 62 and (often) exits the labor force.

Figure 11: Benefit claiming by married men: No health expenditure risk

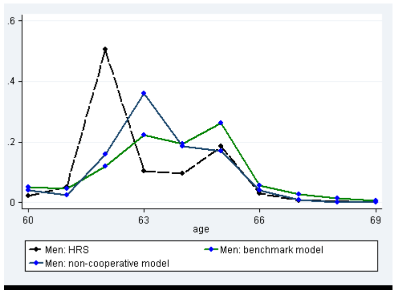

(a) Distribution by age

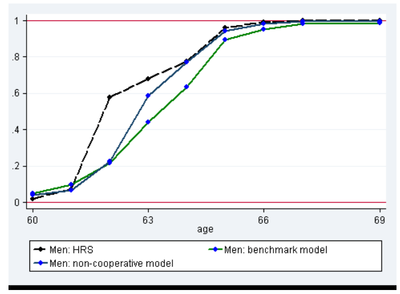

(b) Cumulative hazard

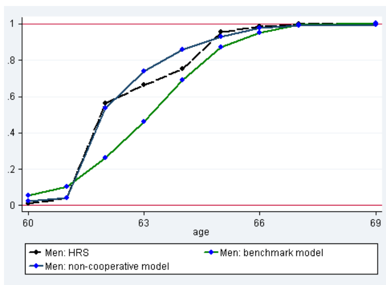

(c) High wealth households

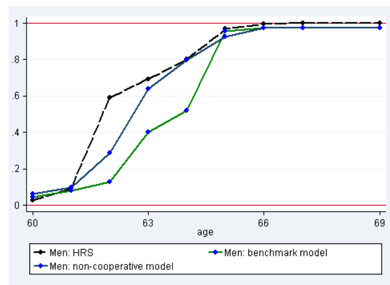

(d) Low wealth households

For higher-wealth households there is no effect. In fact, the lack of medical expenditure risk marginally improves the benefit claiming predictions of the benchmark (cooperative) model, mainly by reducing the need for precautionary saving and labor force participation by husbands in their 60s, and increasing the relative power of the bequest motive that induces couples to convert Social Security into asset wealth early.

4.4 Pre-retirement behaviour and wealth

[Forthcoming]

4.5 Complementarities in leisure

[Forthcoming]

4.6 Welfare and policy evaluation

[Forthcoming]

5 Conclusions

[Forthcoming]

References

Blundell, R., Walker, I., 1986. A life-cycle consistent empirical model of family labour supply using cross-section data. The Review of Economic Studies 53 (4), 539–558.

Bound, J., Burkhauser, R. V., Nichols, A., 2002. Tracking the household income of ssdi and ssi applicants. PSC Research Report 02-502.

Bound, J., Cullen, J. B., Nichols, A., Schmidt, L., 2004. The welfare implications of increasing disability insurance benefit generosity. Journal of Public Economics 88 (12), 2487–2514.

Chiappori, P.-A., Fortin, B., Lacroix, G., 2002. Household labor supply, divorce legislation, and the marriage market. Journal of Political Economy 110 (1), 37–72.

Coile, C., 2004. Health shocks and couples' labor supply decisions. Tech. rep., National Bureau of Economic Research.

De Nardi, M., French, E., Jones, J., 2010. Why do the elderly save? the role of medical expenses. Journal of Political Economy 118 (1), 39–75.

Gruber, J., 2000. Disability insurance benefits and labor supply. Journal of Political Economy 108 (6), 1162–1183.

Gustman, A., Steinmeier, T., 2004. Personal accounts and family retirement. Tech. rep., National Bureau of Economic Research.

Gustman, A., Steinmeier, T., 2009. Integrating retirement models. Tech. rep., National Bureau of Economic Research.

Gustman, L., Steinmeier, T., 2000. Retirement in dual-career families: A structural model. Journal of Labor Economics 18 (3), 503–545.

Huggett, M., Ventura, G., 2000. On the distributional effects of social security reform. Review of Economic Dynamics 2, 498–531.

Imrohoroglu, A., Imrohoroglu, S., Joines, D. H., 2003. Time inconsistent preferences and social security. Quarterly Journal of Economics 118 (2), 745–784.

Imrohoroglu, S., Kitao, S., 2010. Social security, benefit claiming and labor force participation: A quantitative general equilibrium approach. CRR (Boston College Center for Retirement Research) WP 2010-2.

Kambourov, G., Manovskii, I., 2009. Occupational specificity of human capital*. International Economic Review 50 (1), 63–115.

Kopecky, K., Koreshkova, T., 2009. The impact of medical and nursing home expenses and social insurance policies on savings and inequality.

Low, H., Pistaferri, L., 2009. Disability Risk, Disability Insurance and Life Cycle Behavior. Tech. rep., Stanford mimeo.

Lundberg, S., Pollak, R., 1991. Separate spheres bargaining and the marriage market. Journal of Political Economy 101 (6), 998–1010.

Lundberg, S., Startz, R., Stillman, S., 2003. The retirment-consumption puzzle: A marital bargaining approach. Journal of Public Economics 87 (5-6), 1199–1218.

Maestas, N., 2001. Labor, love and leisure: complementarity and the timing of retirement by working couples. UC Berkeley, Department of Economics working paper.

Manser, M., Brown, M., 1980. Marriage and household decision-making: A bargaining analysis. International Economic Review 21 (1), 31–45.

Mazzocco, M., Yamaguchi, S., 2007. Labor supply, wealth dynamics, and marriage decisions, university of Wisconsin Economics Department Working paper.

Rivas, M. C., 2010. Happy together: A structural model of couples' joint retirement decisions, working paper, UCLA.

Schirle, T., 2008. Why have the labor force participation rates of older men increased since the mid-1990s? Journal of Labor Economics 26 (4), 549–594.

Stock, J., Wise, D., 1990. 7 The Pension Inducement to Retire: An Option Value Analysis. Issues in the Economics of Aging, 205.

Van der Klaauw, W., Wolpin, K., 2008. Social security and the retirement and savings behavior of low-income households. Journal of Econometrics 145 (1-2), 21–42.