Social Security, Endogenous Retirement and Intrahousehold Cooperation

Giovanni Gallipoli
University of British Columbia

Laura Turner
University of Toronto

June 12, 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 despite fairly significant wealth gains associated to later claiming by husbands; and (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.

PRELIMINARY AND INCOMPLETE

1 Introduction

Economists have devoted substantial attention to the retirement patterns of married couples in the U.S. The most robust, and well-known facts regarding retirement and Social Security benefit claiming behavior 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. Figure 1 shows the proportion of single and married women (left) and men (right) claiming benefits at each age from 60 to 69. The lower panel shows the cumulative benefit claiming hazards by year from initial eligibility for benefits, taking account of the fact that many single women become eligible at age 60 while all married women become eligible at 62. There is very little difference in the timing of claiming across either gender or marital status. Table 1 shows the distribution of the difference in the benefit claiming age of the husband and the wife, for both benefit claiming and age of final labour market exit.[1] Not surprisingly given the plots in figure 1, married couples of approximately the same age claim benefits together and as early as possible. Unattached men and women also claim their benefits as early as possible.

For married couples, these patterns are strongly at odds with the advice of financial planning literature (e.g. Meyer and Reichenstein (2010), Sun and Webb (2009)), which emphasizes the gains to married couples from having the wife claim her benefit at age 62 while

[1]The sample includes households in which the wife is between zero and three years younger than the husband, and neither spouse holds 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 Casanova-Rivas (2010) among others) and/or when db pension holders are included.

the husband delays his benefit claim late into the retirement window. The main reason for married people postponing benefits, relative to singles, is the U.S. Social Security Administration (SSA)'s rules governing survivor (traditionally called "widow") benefits. For the majority of married households in which the husband's Primary Insurance Amount (PIA), the benefit he is entitled to claim at the "regular retirement age" of 65 (as of 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 7% a year for each year of delayed claiming (a rate that is nearly actuarially fair averaged across the entire population)[3]. Under most scenarios, a main earner bequeaths his entire Social Security *benefit* to a surviving previously-retired spouse at his death, regardless of whether the head has yet claimed the benefit or not and at what age the spouse initially claimed her own benefit.[4] This has the effect of making the household's combined mortality profile the relevant schedule for maximization of Social Security wealth, at least if the household is otherwise unconstrained and can use assets, along with the wife's personal Social Security benefit, to smooth consumption during the period before claiming full benefits.[5]

[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 7% per year between the ages of 62 and 65, and by 6% per year between ages 65 and 69.

[4]Young widows, those whose husbands die before early retirement age, can claim a benefit based on the husband's PIA as early as age 60. In this case, regular actuarial adjustments are made depending on the age she claims the benefit just as they would be should she claim Social Security on her own earnings history. It is also possible for young widows to "switch" between one benefit and the other by, for instance, claiming the benefit based on their own record at age 62, subject to an actuarial penalty, and to convert to the unpenalized survivor benefit at full retirement age. We do not model these mechanisms in the present version of the paper.

[5]This argument is complicated when either the husband or his spouse holds a defined benefit (db) pension or spent part of his/her career in federal employment. Pensions often have built-in Social Security offset rules that nullify much of the incentive to delay claiming (Stock and Wise (1990)) while the Social Security Administration public offsets the benefits of some former federal employees since they receive a government pension and did not pay Social Security taxes on their earnings. Since most government workers are db pension holders, and db pension incentives are in general quite difficult to model, we follow Casanova-Rivas

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

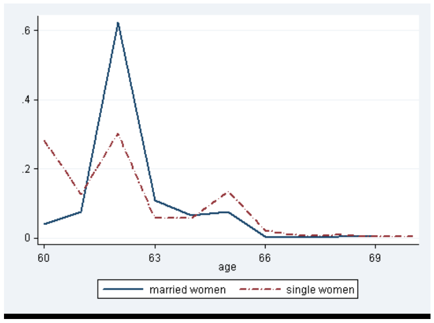

(a) HRS women

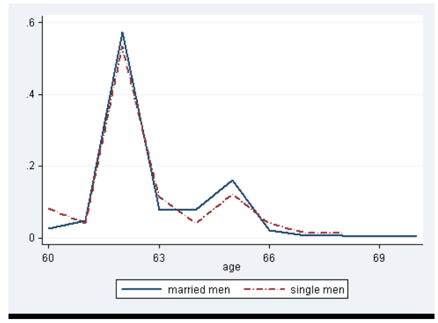

(b) HRS men

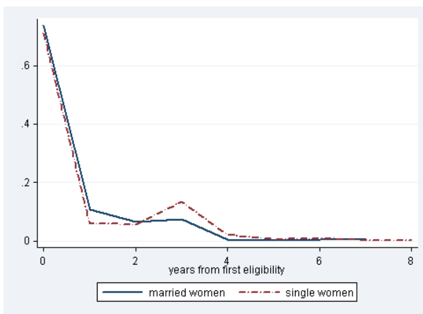

(c) HRS women

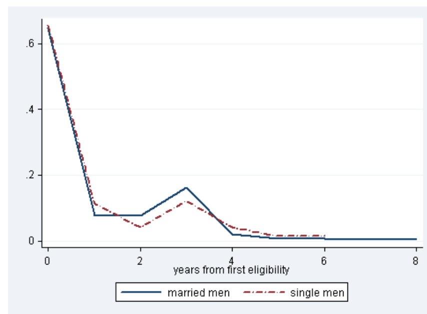

(d) HRS men

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 of total)	2.9	14.9	14.7	58.9	4.9	4.8	1.3
labor force exit (share of total)	14.7	14.2	8.0	37.5	6.8	7.2	11.6

In recent years, economists have proposed several explanations to rationalize both early and joint retirement, of which the most important appear to be the roles of declining health (Imrohoroglu and Kitao (2010) for single agents; Van der Klaauw and Wolpin (2008) for couples) and complementarties in spousal leisure (Maestas (2001), Coile (2004), Gustman and Steinmeier (2000), Schirle (2008) and Casanova-Rivas (2010)).[6] Casanova-Rivas (2010) also shows that the Social Security spousal benefit generates modest incentives toward earlier

(2010) by excluding bd pension holders from our analysis.

[6]Defined benefit pension incentives, which we ignore in this paper, also clearly play an important role in explaining labor force withdrawal patterns. See, e.g. Maestas et al. (2006).

retirement among constrained households in which the wife qualifies, since she cannot receive the benefit unless the husband has already filed. With the exception of spousal benefit incentive, however, the former explanations have more bite for explaining patterns of early retirement for the labor force than for explaining Social Security benefit claiming. As well, health incentives cannot easily account for the lack of difference in the claiming patterns of married and single households since single males are typically less healthy than married men and have fewer liquid assets, making them more likely to need to claim Social Security benefits early. This difficulty can be circumvented by "bundling" the benefit claiming and labor market exit decisions of each spouse.[7] 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.[8] 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 choice variables in the transition from middle into old age.

In our paper, 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

[7]For example, Van der Klaauw and Wolpin (2008), Casanova-Rivas (2010), Gustman and Steinmeier (2004) and an earlier draft of this paper take this approach.

[8]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

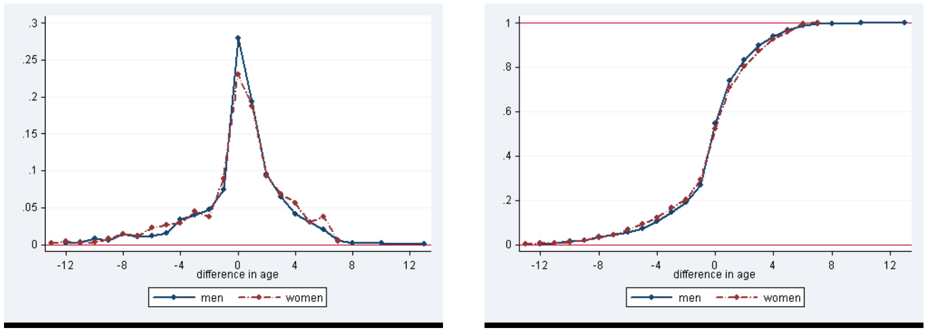

(a) PDF: men and women

(b) CDF: men and women

with life cycle theory when labour force participation (or "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 Casanova-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 the model to include a household-level saving decision. Effectively, in each period, the household jointly optimizes its forward-looking utility 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 of labour supply, assets, and (public) consumption. 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 intuitive restrictions imposed on the nature of deviations, an incentive compatible allocation must be a pure-strategy Nash equilibrium for both partners where, as in Gustman and

Steinmeier (2009), the payoffs are the continuation value functions for each spouse.

One additional feature of the model is very important for explaining our result: the riskiness of assets due to health expenditure shocks. As argued by De Nardi et al. (2010) and Kopecky and Koreshkova (2009), individuals face an increasing, and increasingly variable, process for health expenditures as they age. These health expenditures—which for our sample up to 2004 included prescription drug costs and continue to include nursing home care as well as "perks" such as private hospital rooms which are not covered under Medicare—may be paid for either out of pocket or, subject to a wealth means-test, through public insurance, specifically Medicaid, which leaves beneficiaries with a minimum level of consumption.[9] While much recent research has been devoted to the implications of late-life health expenditure risk on wealth accumulation and inequality, in our model health expenditure risk also has implications for the timing of benefit claiming. In a fully cooperative joint-optimization framework, the riskiness of assets provides an additional incentive for delaying, and thereby increasing, Social Security benefits. Approximately 70% of households survive a health-expenditure (and associated health) shock after age 70: in 50% of cases, both spouses survive. If assets are depleted, having a Social Security annuity becomes very important. In the non-cooperative context, however, health expenditure risk can further induce early claiming by husbands. For one thing, households are relatively constrained since they need to maintain a buffer stock of assets to deal with health risk, making it difficult to consumption-smooth in the absence of Social Security income. Second, most medical risk

[9]Among the very elderly, the share receiving Medicaid support is in fact quite large. For instance, 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.

will be borne by the longer-living spouse, the wife. Since husbands, who have lower life expectancies, are typically interested in shifting consumption into the present, this imperfect substitutability increases the incentive of husbands to retire early.[10]

Our "strategic retirement" model does a good job of predicting the cross-sectional patterns presented in this section: the spike in retirement at 62; the distribution of joint benefit claiming and joint retirement within couples; and the individual timing of benefit claiming and final labor market exit. The model also has substantial welfare implications. Specifically, not cooperation over career decisions has major welfare consequences for couples, reducing the efficiency of marriage.

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.[11]. We also discuss the welfare consequences of non-cooperation for households.

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

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

2 The household optimization problem

We follow households from when the head is age 48 if female (49 if male) until death. Men are always one year older than women in the model. This is the mean difference in age between husbands and wives in our HRS sample.[12] Men (women) survive to a maximum age of 101 (100) years. 80% of individuals are married at age 50, and there is no divorce thereafter.[13] A model period is one year. Lifespan of individuals and households is uncertain ex-ante. $\varphi_{h,i}^{j}$ (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-expenditure 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 so as to match the distribution of HRS households at 48-49. 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). 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 expenditure state h, (4) accumulated Social Security credit from previous work e; (5) marital status $ms \in \{single, married\}$, (6) labour force status ls, and (7) benefit claim status ss. The last two states are described in

[12] The sample is restricted to households in which the wife is between zero and three years younger than the husband for a median difference of -1 year.

[13] Introducing exogenous separations would not complicate the model; we discuss the possibility of endogenous separations in a forthcoming robustness check.

more detail below.

There are three major sources of life cycle risk: productivity (w) risk and job loss (ls) risk, health expenditure risk which imposes out-of-pocket medical costs and increases mortality risk; and mortality risk itself. Married households pool income and enjoy an economies of scale in shared household-level public consumption.

Up to age 75, ls can take values of C for "career-job", NC for "non-career job" or R, not working ("retired"). 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. A career worker exogenously loses his job with probability p_{NC}, which differs by gender and age. Social Security claim status, ss, also takes two values, NB for not receiving or claiming benefits, B for receiving benefits. There are both retirement benefits (for individuals over 62) and disability benefits for younger individuals. While applying for and holding disability benefits before age 62, individuals must choose $ls = R$. Their applications are successful at rate p_δ which we take to be constant.[14] Applicants who apply from the career job state forfeit their job and fall off the career track into non-career work. Individuals applying for regular benefits after age 62 can apply from any ls state and continue to work either kind of job while claiming benefits. However, individuals cannot reverse their benefit claims.[15]Acceptance of the regular retirement benefit application is

[14]The Social Security Adminstration uses a more lenient rule for evaluating claims for individuals over 55 (Low and Pistaferri (2009)). We omit this rule for simplicity and because disability application behavior is only a very secondary concern in our model.

[15]As pointed out by, for instance, de Silva (2006), this is unrealistic. Individuals can and do occasionally suspend benefits having already claimed them for one or more years. Unfortunately, allowing this possibility is computationally intractable.

automatic: we assume, with very little cost in generality, that everybody qualifies for Social Security. Those whose earnings are very low qualify for an alternative minimum benefit, SSI, described in section 3.2.1.

2.1 Life cycle optimization when spouses are non-cooperative

Below, we sketch out a very standard dynamic optimization problem for single and married households, beginning with single households. We then explain how the problem changes and outline the computational approach when spouses behave strategically over ls and ss decisions.

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 $x_s = \{a, e, w, h\}$. 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 $x_s \equiv x$ is:

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

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:

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

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

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

where I is household period income net of health costs $\chi(h_b)$ and $\underline{c}$ is the minimum level of consumption provided for high (bad) health-cost individuals who cannot cover their own costs. (The consumption floor is zero for individuals in ok health but never binds due to the provision of benefits.) $\tilde{T}$ 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 62 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 standard unitary 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 $x_M = \{a, h, e_m, e_f, w_m, w_f, \lambda\}$. Assets are accumulated and health expenditure shocks experienced at the household level, but social security accumulation and wages 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 set $\lambda = .5$ for all couples, though we conduct robustness checks of the results using other "reasonable" values of λ.

The household budget set for marrieds, B_M is:

$$(\tilde{T} - l_f^{ls} w_f^{ls} + \tilde{T} - l_m^{ls} w_m^{ls} + (1+r)a$$
$$+ b(a, e_f, e_m, ss_f, ls_f, ss_m, ls_m) - \chi^M(h) = I$$

$$I = c + a^{'} \quad a' \geq \underline{a} \qquad if \quad I - \underline{a} > \underline{c}^M(h)$$
$$c = \underline{c}^M(h) \quad a = \underline{a} \qquad if \quad I - \underline{a} \leq \underline{c}^M(h)$$

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 and, as for singles, depends on health expenditure status.

2.1.3 Non-cooperation and strategic retirement

We now consider the case in which the spouses are not bound to respect the jointly optimal allocation but can make their own choices over whether or not to claim benefits (ss) or to work in the allocated ls state. 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 evaluating his payoffs. In general, with three ls and two ss states, this is a very complicated problem, especially given that the spouses also make a choice over *household*-level variable a. It turns out, however, that the problem can be rendered very simply under fairly minor assumptions about how indeterminant games are resolved.

Define $\hat{\Upsilon}$ as each couple's joint payoff after ss and ls status are chosen optimally by each spouse. The couple, or household planner's value function $\hat{U}$ is then:

$$\hat{U}(x_M) = \max_{a'} \lambda \hat{\Upsilon}_f^{BR}(x_M, a') + (1-\lambda)\hat{\Upsilon}_m^{BR}(x_M, a') \tag{4}$$

$$s.t. \tag{5}$$

$$BR^f(\cdot) = BR^M(\cdot) \tag{6}$$

where BR stands for "best response" and $BR \equiv \{LS^i, SS^i | LS^*_{-i}, SS^*_{-i}\} \equiv \{\mathcal{L}^i | \mathcal{L}^*_{-i}\}$ for spouse i with partner $-i$. and the planner can only implement allocations in which neither spouse has an incentive to deviate from the given allocation. Equation 4 implies the following assumptions:

1. Assumption #1: If more than one (pure) Nash equilibrium exists, the spouse's always choose the jointly optimal one.

2. Assumption #2: If no (pure) Nash equilibrium exists given a' and x_M, there is no solution and the allocation cannot be implemented. In principle it is possible that there could be states in which no implementable solution exists because the spouses never agree. In practice, however, the "planner" can impose a level of saving sufficiently high that the spouses must agree on the income-maximizing choice. Obviously, these states generate low joint utility and are avoided.

Computationally, the solution works as follows: for each point in the state space x_M at age j, the planner considers a vector of discreet carry-forward assets $A'(a) = \{1, 2, 3, NA'\}$ where a is the current level of assets, a component of x_M. The planner solves for the best incentive compatible allocation $\{\mathcal{L}_f, \mathcal{L}_m\}$ at a', assuming one exists (otherwise a sufficiently low utility is assigned to prevent that a' from being chosen in the last stage). In the final stage, the planner chooses the best overall allocation $\{a', \mathcal{L}_f(x_M, a'), \mathcal{L}_m(x_M, a')\}$.

The spouses' individual optimizations at each a' can be written:

$$\hat{\Upsilon}_i^{C,NA} = \max \begin{Bmatrix} \Upsilon_i^{R,NA,\mathcal{L}^*_{-i}}, & \Upsilon_i^{NC,NA,\mathcal{L}^*_{-i}}, & \Upsilon_i^{C,NA,\mathcal{L}^*_{-i}} \\ \Upsilon_i^{R,A,\mathcal{L}^*_{-i}}, & \Upsilon_i^{NC,A,\mathcal{L}^*_{-i}}, & \Upsilon_i^{C,A,\mathcal{L}^*_{-i}} \end{Bmatrix} \tag{7}$$

for workers currently in a career job and not holding benefits;

$$\hat{\Upsilon}_i^{NC,A} = \hat{\Upsilon}_i^{R,A} = \max \left\{ \begin{matrix} \Upsilon_i^{R,NA,\mathcal{L}^*_{-i}}, & \Upsilon_i^{NC,NA,\mathcal{L}^*_{-i}} \\ \Upsilon_i^{R,A,\mathcal{L}^*_{-i}}, & \Upsilon_i^{NC,A,\mathcal{L}^*_{-i}} \end{matrix} \right\} \tag{8}$$

for workers in ls state either NC or R and not holding benefits;

$$\hat{\Upsilon}_i^{C,A} = \hat{\Upsilon}_i^{R,A} = \max \left\{ \Upsilon_i^{C,A,\mathcal{L}^*_{-i}}, \quad \Upsilon_i^{R,A,\mathcal{L}^*_{-i}}, \quad \Upsilon_i^{NC,A,\mathcal{L}^*_{-i}} \right\} \tag{9}$$

for workers in ls state C who are already holding benefits; or

$$\hat{\Upsilon}_i^{C,A} = \hat{\Upsilon}_i^{R,A} = \max \left\{ \Upsilon_i^{R,A,\mathcal{L}^*_{-i}}, \quad \Upsilon_i^{NC,A,\mathcal{L}^*_{-i}} \right\} \tag{10}$$

for workers in ls state NC or R who are already holding benefits.

The latter situation is clearly the simplest—as well as one commonly generated in the model economy—implying a 2x2 static Nash game than can be represented in the standard way. Figure 3 shows the potential outcomes of this game, and the situations that generate deviations from the planner's given allocation or that cannot be implemented. Note that the need to avoid intractable areas of the state space is one way in which the incentive compatibility constraints bind on the household. Figure 4(b) shows a situation in which the household may actually implement a Pareto inefficient outcome at a': a form of the prisoner's dilemma. In practice, these sorts of inefficiencies are unlikely to occur as the planner would choose a different level of carry-forward assets. However, the scope for inefficiency is clearly quite large in the model.

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

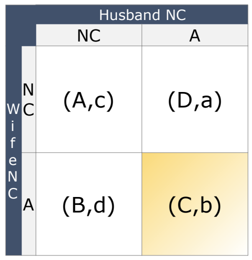

(a) Husband dominant R strategy

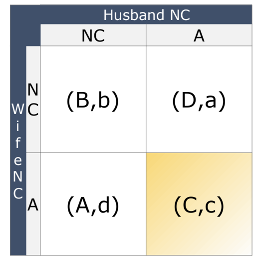

(b) NE: Husband and wife dominant R strategies

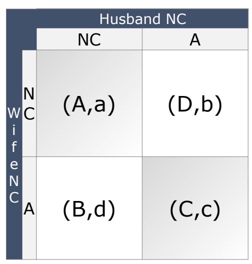

(c) Multiple NE: Implements planner's solution

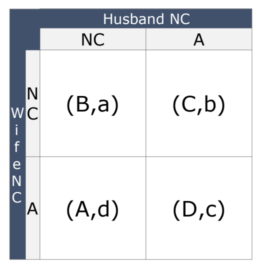

(d) No pure NE: Infeasible choice of a'

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 Sources of life cycle risk

There are three major sources of risk in the model: (1) productivity risk, which takes the form both of wage risk conditional on ls status, and job-loss risk, a forced downgrade of ls status from career to non-career work; (2) health expenditure risk; and (3) mortality risk.

3.1.1 Wage and career risk

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

1. are first observed with a tenure of eight or more years in the current job, or attain eight years of employment on the job within the first three survey waves before leaving that job.

2. change jobs, but remain within occupation category[16] and experience an average wage over the first two survey waves in the new jobs that is greater than or equal to the wage in the last year of the old job.

3. Work a minimum of 20 hours a week on average in the job over the year.

The second condition allows us to identify as career workers those who may have participated in on-the-job-search and upgraded their full-time job. The restriction that workers move into new jobs within their own occupation is consistent with the evidence in Kambourov and Manovskii (2009), who show that human capital tends to be occupation- rather than

[16] We use the 17-category occupation listing provided in the RAND HRS.

individual- or firm-specific. Roughly 55% of male workers, and 30% of female workers in their between 48 and 55 are identified as career workers; this share falls sharply in the late 50s and 60s and eventually reaches around 2% by age 70. Married men and unmarried women are more likely to be career workers at every age, which we account for in the initial distributions.

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 system of equations separately for non-career, career men, non-career women, and career women:

$$\ln w_{it}^{g,ls} = \beta_0^{g,ls} + \beta_1^{g,ls} age_{it} + \beta_2^{g,ls} age_{it}^2 + \epsilon_{it}^{g,ls} \tag{11}$$

$$\epsilon_{it}^{g,ls} = \nu_{it}^{g,ls} + \eta_{it}^{g,ls} \tag{12}$$

$$\nu_{it}^{g,ls} = \rho^{g,ls} \nu_{it-1}^{g,ls} + \varsigma_{it}^{g,ls}; \qquad \varsigma_{it} \sim N(0, \sigma_\varsigma^2), \qquad \eta_{it} \sim N(0, \sigma_\eta^2) \tag{13}$$

for $g \in \{f, m\}$, $ls \in \{C, N\}$. Since these are reduced-form estimates, we search for structural parameters that generate them in cross-section regressions on the simulated data. We assume that η is simple classical measurement error and do not include it explicitly as a transitory shock in agents' optimization problems.

The resulting wage targets, and the average wage by gender and career status in our HRS sample, are given in table 2. Two, fairly unsurprising, features are of note. First, shocks to career jobs are moderately less persistent and have much smaller variance, than shocks to non-career jobs. Second, career jobs are much more remunerative (unconditionally) than

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

	Men		Women	
	career	non-career	career	non-career
cons	1.53	-4.04	3.01	6.09
	(.419)	(.800)	(.416)	(.642)
age	.0558	.174	$3.98E^{-3}$	-.135
	(.0140)	(.0263)	(.0138)	(.0213)
age^2	-5.16E^{-3}	-1.26E^{-3}	-1.15E^{-5}	1.02$E{-}3$
	(1.18E^{-4})	(2.16E^{-4})	(1.06E^{-4})	(1.75E^{-4})
ρ	.944	.958	.906	.956
σ^2_ς	.044	.105	.053	.071
σ^2_η	.152	.592	.128	.532
[6pt] *mean*	$28.4	$18.4	$19.4	$10.4

non-career jobs. The average career job is 45% higher than the average non-career wage for women, and 35% higher for men (although for men, the career wage effect dissipates with age to about a 20% differential at age 60).

The second component of earnings risk is the probability of losing a career job. We model this probability as $p^g_{NC} = p^g_0 + p^g_1 \times age$, allowing for layoff risk to vary by age and gender. In the data it is hard to fully distinguish voluntary and involuntary job separations, but respondents are asked the reason they left their previous employer. About 16% of both women and men leave their jobs due to plant closures or lay off, which suggest involuntary separations. Poor health and burnout are much more common explanations. The share of "involuntary" separations among all separations falls with age, for men from 24.5% of men under 55 to only 14% of men 55 and over. We can therefore test to see if the share of involuntary departures from career jobs line up with what we see in data. We return to this issues in section 4.

3.1.2 Health and medical costs

We assume that five percent of households experience health expenditure shocks (χ) by at every age with the estimated size of χ increasing with age. Mortality also rises when households experience the shock by a rate calibrated to replicate the likelihood of dying following such a shock in the HRS. Equations 14- 15 show the transition matrices and evolution of costs associated with the poor health (state h_b, with the good state denoted h_{-b}) by age, estimated separately for married and single households. These data are based on "out of pocket" medical expenditures from the HRS, where the cost in the bad state is the difference of the median out-of- pocket expenditure of households in the top 5% of payers and households in the bottom 95%. The persistence of the poor health state h_b is not very high, especially for married households, though the matrices understate the persistence slightly because of the increased likelihood of dying in the bad state. The costs of poor health, however, increase substantially over the late life cycle from \$250 at age 50 to \$700 at age 80 and above.

$$\begin{bmatrix} pr(h_{-b}|h_{-b}) & pr(h_b|h_{-b}) \\ pr(h_{-b}|h_b) & pr(h_b|h_b) \end{bmatrix}_{married} = \begin{bmatrix} .977 & .023 \\ .929 & .071 \end{bmatrix}, \begin{bmatrix} pr(h_{-b}|h_{-b}) & pr(h_b|h_{-b}) \\ pr(h_{-b}|h_b) & pr(h_b|h_b) \end{bmatrix}_{single} = \begin{bmatrix} .965 & .335 \\ .818 & .182 \end{bmatrix} \tag{14}$$

$$\chi(age, h_b)_single = -3266 + 171.8 * age - 2.85 * age^2 + .0160 * age^3$$

$$\chi(age, h_b)_married = -14602 + 634.4 * age - 8.79 * age^2 + .0405 * age^3 \tag{15}$$

3.1.3 Mortality risk

We take mortality profiles from the National Vital Statistics Reports (1995, 2000 and 2010 reports corresponding to the average of our cohort in each year) for men and women. We use HRS data on health expenditure to calculate adjusted mortality profiles for households in health expenditure state h_b using the identity that, at each age and for each gender,

$$\varphi_g = \widehat{Pr}(Surv|h_b)\widehat{Pr(h_b)} + \widehat{Pr}(Surv|h_{-b})\widehat{Pr(h_{-b})}$$

$$\varphi_g = \widehat{Pr}(Surv|h_b) \times .05 + \widehat{Pr}(Surv|h_{-b}) \times .95$$

where hats indicate that the information comes from the HRS.

3.1.4 Preferences

We adopt Cobb-Douglas preferences in period-utility. Economists have generally found strong evidence that consumption and leisure are inseparable in utility (see, e.g. Blundell and Walker (1986)), and the ability to substitute efficiently between consumption and leisure is likely to be especially important in the context of household retirement. Utility in the model is given by:

$$u_g(c_g, l_g) = \frac{[(\frac{c_i}{\psi(ms)})^{\alpha} \times l_(g, age)^{1-\alpha_g}]^{1-\omega}}{1-\omega} \tag{16}$$

and "effective leisure" $l_g = T - \gamma^g n^{ls} = T - (\gamma_0^g + \gamma_1^g age n^{ls} - \gamma_1^g age^2 n^{ls}$.

In the above expressions, α measures the consumption share in utility, ω is the coefficient of relative risk aversion and ψ is a cohabitation equivalence scale measuring each individual's "take" out of household-level consumption (=.6 for marrieds; =1 for singles). T is an individuals's total weekly time endowment and γ measures the leisure cost of work. Specifically, younger workers may derive some leisure enjoyment from their jobs so that time at work does not offset leisure one for one. As individuals age and "burn out" or become less physically healthy, the total leisure cost of work increases and may exceed one if work imposes mental or physical stresses from which the individual must recover during his non-work hours. Because the γs and *alpha* are not jointly well identified, we set $\alpha = .5$ and estimate the γs that best replicate the life cycle profiles of career and non-career work and retirement by gender.

3.1.5 Model Estimation

Estimation of the models is performed through simulated method of moments (SMM). There are 38 parameters in total, recapped in table **??**. The targets are:

- Wage profiles and variance by gender and C/NC type
- Liquid asset holdings for married and single households before and after age 62
- The life cycle profiles (in five-year age intervals) of ls status by gender
- Application and application success rates for disability insurance
- The benefit claiming hazards of single men and women from 62 to 69.

The model is over-identified and therefore does not provide a perfect fit to the targeted profiles. Note, however, that we do not target the benefit claiming hazards of married couples which are left as an out-of-sample test of the respective models.

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

Parameter	Benchmark	Strategic
γ_f	$2.37 - .0682age + 4.07E^{-3}age^2$	$1.48 - .0227age + 5.36E^{-3}age^2$
γ_m	$1.54 + .109age + 9.91E^{-4}age^2$	$1.23 + .104age + 5.05E^{-4}age^2$
β_f	.905	.902
β_m	.913	.898
ω	2.02	2.49
c_A	1.59	2.23
p_δ	.48	.53
p^f_{NC}	$P_{NC} = .001 + .0003 \times age$	$P_{NC} = .003 + .00001 \times age$
p^m_{NC}	$P_{NC} = .006 - .0001 \times age$	$P_{NC} = .002 + .0002 \times age$
$\ln w^{f,nc}$	$3.47 - .0173 \times age - 4.29E^{-5} \times age^2$	$3.43 - .0178 \times age - 5.05E^{-5} \times age^2$
$\ln w^{f,c}$	$3.14 - .0119 \times age - 4.96E^{-5}$	$3.23 + 2.59E^{-4} \times age - 7.68E^{-5} \times age^2$
$\ln w^{m,nc}$	$1.91 + 0.014 \times age + 6.54e^{-5} \times age^2$	$1.94 + 0.010 \times age + 6.54e^{-5} \times age^2$
$\ln w^{m,c}$	$4.67 - 0.069 \times age + 6.10e^{-4} \times age^2$	$5.29 - 0.067 \times age + 5.64e^{-4} \times age^2$

3.2 Social security and optimal claiming

3.2.1 Social Security in the model

We model the following major features of Social Security:

Benefit accrual and determination. Social Security 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 disability insurance, up a maximum of five excluded years. From the AIME, the Primary Insurance Amount (PIA), and the actual retirement-age or disability benefit, b is determined according to:

$$
\begin{aligned}
PIA &= 0.9 \min(0.2\overline{w}, E) \\
&\quad +0.32 \max(\min(E - 0.2\overline{w}, 1.3\overline{w} - 0.2\overline{w}), 0) \\
&\quad +0.15 \max(E - 1.3\overline{w}, 0) \\
a_t < \overline{a}^{mt} : &\qquad b_{.} = \max(PIA, SSI) \\
a_t > \overline{a}^{mt} : &\qquad b_{.} = PIA \qquad (17)
\end{aligned}
$$

where $\overline{w}$ is the average per-capita wage earnings in the economy, E is shorthand for the AIME, and SSI is a floor benefit, equal to \$151 weekly for an individual (\$222 for a couple). This benefit, the Supplemental Security Income 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.[17] 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 receive the benefit.

State space limitations prevent an accurate accounting of which years should be included in benefit determination in the simulated economy.[18] Instead, we exclude earnings from all 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.

[17] 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.

[18] Huggett and Ventura (2000) and others discuss a similar technical difficulty in their calibrated analysis of the Boskin Social Security reform proposal.

Given the shape of the life cycle profile of earnings, these tend to be the lowest-earning years. The benefit accumulation calculation for workers under 60:

$$E_{age} = \frac{E_{age-1}(age-30) + w^{ls}_{age} n^{ls}_{age}}{age-31}$$

and for workers over 60 is:

$$E_{age} = max(\frac{E_{age-1}(age-30) + w^{ls}_{age} n^{ls}_{age}}{age-31}, E_{age})$$

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 practice, to receive Social Security benefits before age 62 under the disability insurance program (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, with 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.[19] 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 practice, disability benefits may be awarded to people who do not actually face a work-limiting condition. In the model, we assume that individuals can apply for disability

[19]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.

benefits in any state. As described in the estimation section, we calibrate the application cost c_A (as a share of period utility of a representative individual) and the likelihood of acceptance p_δ so that roughly .8% of working age individuals apply per period the overall acceptance rate (allowing for multiple subsequent attempts) is 50%. While applying for and receiving disability benefits, individuals must remain in state R, though they may begin working again at age 62.

Early and delayed retirement. 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 (6.7% per year) before the full retirement age of 65. The factor of adjustment for later retirees is 6% 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 accounted for in the model.

Survivor and spousal benefits. Second earners, typically wives, 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, although only once the main earner has claimed his own benefit, or at age 65.[20] Casanova-Rivas (2010) has shown that spousal benefit provisions induced some early joint benefit claiming since they allow constrained couples to access between 120% and 140% of the main earner's PIA when the wife reaches age 62, depending on the age difference of the spouses. Survivor benefits entitle the surviving spouse, again typically the wife, to receive the greater of her own or her spouse's benefit for the remainder

[20]At full retirement age, the main earner can claim and suspend his own benefits allowing his wife to claim the spousal benefit while his benefits continue to appreciate over the retirement window.

of her life. Unlike spousal benefits, that are based on the husband's PIA, survivor benefits encompass for all delayed credits, which makes delayed claiming a very good deal for some households, as described further below. Young widows who have not yet claimed benefits on their own earners record are able to retire as early as 60 under the survivor benefit, but subject to a standard early-claiming penalty. More complicated rules that allow young widows to claim on their own benefits and then "switch" to a non-adjusted survivor entitlement at full retirement age are omitted from the current version of the model due to computation complexity.

Earnings test. Individuals who continue to work between ages 62 and 64, while receiving benefits, are subject to a benefit adjustment process that "claws back" and defers current benefits. 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,[21] though the taxation of benefits is likely to change (see Appendix 6). Recipients between 62 and 64 remain subject to a deferral in their present average weekly benefit of $1 for every $2 earned above $250 (in 2006 USD). The lost benefits are treated as partially delayed claims and appreciate at the same rate as standard delayed benefits. As well, individuals can further raise their PIA by replacing low-earning with high-earning years (in the model, if their current earners are

[21] 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.

higher than their average Social Security earnings to date).

Minimum consumption floor and Medicaid. Individuals who are unable to meet their medical expense requirements in state h_b out of their own wealth stocks qualify for Medicaid, which pays their expenses and leaves them with a base level of consumption, $\underline{c}$. Following De Nardi et al. (2010), we set $\underline{c}^M$ to \$300 for married couples and \$200 for singles, which is approximately equal to the income received from SSI.

Additional, non-Social Security, aspects of the policy environment, including taxation, are described in Appendix 6.

3.2.2 Optimal claiming by singles and couples

Table 3.2.2 shows the optimal claiming age of males, and associated wealth gains—as a share of total Social Security wealth—relative to claiming at 62, for men in our HRS cohort based on the Social Security rules described above. The calculations assumes that the wife claims her benefit exactly at age 62. Focussing on the claim date of men, we report results for men and women by the couple's health status and the wife's PIA relative to the husband's when the husband is 62. We assume an interest rate of 4%, which is a typical value used in household life cycle literature as well as the value adopted in the model. We calculate health-adjusted mortality profiles by following individuals who reported themselves in less than "very good" health at age 62 and adjusting the Vital Statistics data similarly as in section 3.1.2. Overall, the wealth results from the table are quite similar in magnitude to figures reported by Sun and Webb (2009) in their analysis of optimal claiming and fairly typical of the wealth gains and optimal claiming dates suggested in the financial planning literature.

Two features of the table are important in motivating the results in the next section. First, regardless of the partners' respective health status, it is jointly optimal for the husband to delay his benefit claim into his late 60s. If the husband is unhealthy, he is less likely to live to average mortality, but still leaves his benefit to his (one year younger) spouse in the form of a survivor benefit. This effect holds regardless if whether or not the wife will claim spousal benefits ($PIA_f < .5PIA_m$). However, even if both spouses are healthy, it is *still* optimal to delay claiming to age 68, since the couple's joint mortality profile still compares favourably to a single individual's mortality profile.

Second, the delaying incentives for the single man are lower, for the reasons explained already explained. Indeed, if a single male (assuming he is indifferent to potential claimants just as an ex wife who qualifies for his survivor benefit) is self-assesses as unhealthy in the HRS at 62, he maximizes his wealth by claiming his benefit immediately.[22] A healthy single male should ideally postpone claiming to age 67, but the associated wealth cost for failing to do so is only about half as large as the average cost for a married couple.

Finally, it is worth emphasizing that these figures probably understate the return to delayed claiming. As Meyer and Reichenstein (2010), Sun and Webb (2009) and other authors have noted, the figures in the table only capture differences in the present value of social security wealth, and do not capture the gains from increased insurance against longevity risk. As emphasized in this paper, neither do they capture the value of the Social

[22]For single women, who have longer-than-average longevity, the optimal claiming ages are later. However, women's lower earnings means that they may qualify for SSI regardless of when they retire, and may also depend on a former spouse's benefit.

	Wife's PIA relative to husband	
Wife's health / Husband's health	.4	.75
Healthy / Healthy	68 **(6.1%)**	69 **(5.4%)**
Healthy / Unhealthy	69 **(6.3%)**	69 **(5.3%)**
Unhealthy / Healthy	68 **(5.4%)**	69 **(5.3%)**
Unhealthy / Unhealthy	68 **(5.2%)**	68 **(4.6%)**
Single male, healthy	67 **(3.3%)**	
Single male, unhealthy	62	

Security annuity against medical expenditure shocks that can wipe out liquid asset stocks.

4 Results

We now turn to assessing the implications of the model described in the previous section and considering its implications for benefit claiming and the joint retirement of spouses. The following results are incomplete, but give a flavor of the model's implications for the retirement behavior of couples and the welfare implications of non-cooperation in marriage. We refer to the cooperative model as the "benchmark" and the non-coooperative model as "strategic".

4.1 Parametrization

We begin by discussing the implications of the parametrization reported in table 3. The cooperative and strategic models differ in three main ways:

1. Men and women both have lower preference for leisure (more exactly, face less disutility from work) in the strategic as compared to the benchmark model.

2. Men and women are more risk-averse, and less able to intertemporally substitute utility, in the strategic as compared to the benchmark model ($\omega = 2.49$ vs. $\omega = 2.01$).

3. In addition to their lower mortality risk at every age, women have higher discount rates than men in both the models.

These features of the estimated models are important but unsurprising. The strategic model of the household is not efficient. Both members of the couple have an incentive to try to offload main earner status onto the spouse since the benefits of work are shared jointly through the public consumption good but leisure is enjoyed individually. This gives spouses an incentive to quit a career job. Relatively low leisure preference is therefore required to induce married men and women to remain in the labor force, and especially in career jobs.

The greater risk-aversion estimated in the strategic model is also intuitive. Since we target a distribution of assets for married households, the model implies that, when couples are not fully cooperative in household decision making, asset accumulation is too low relative to the utility-maximizing level. This is in part because spouses have disagreement over the timing of consumption: men, with lower discount factors and lower expected longevity, prefer to consume in the present. At higher saving rates, they are likely to quit their jobs so as to shift leisure rather than consumption into the present.

A secondary finding, that is consistent with the above discussion, is that the disability application cost c_A is higher in the strategic than the benchmark model. Men and women who operating with respect to a personal rather than a joint utility function have a greater likelihood of applying for benefits since receiving disability insurance necessitates that the individual not work, thereby transfering work obligations to the spouse.

4.2 Cross sectional benefit claiming patterns

We now turn to examining the implications of the model for the timing of Social Security benefit claiming within married households. Figure 4 plots the benefit claiming hazards, by year from initial eligibility, for single men and women in the models. The black lines show the same hazards estimated from the HRS. We show the results by year from initial eligibility for benefits (based on the marital history of single women in the HRS) rather than by age since many of these women (roughly 20%) are young widows claiming survivor benefits at age 60. The model replicates the early claiming behaviour of singles. Single men demonstrate slightly too much retirement at 65, an over-responsiveness to the Social Security earnings test. (When the earnings test is eliminated in the model, the men retiring at 65 shift their claiming date forward.)

Overall, the wealth gains to singles from delayed claiming are fairly low. Though they are higher for single women who have higher expected longevity than for single men, women's benefit entitlements are also lower on average, giving them access to the income floor provided

Figure 4: Single men and women: Benefit claiming hazard by age of first eligibility

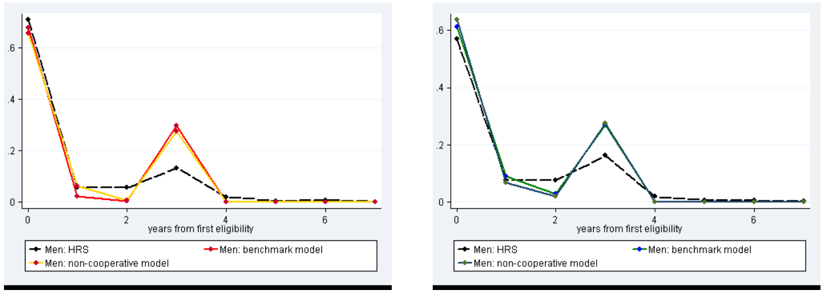

(a) Single women (b) Single men

by SSI. Additionally, the low liquid asset holdings of singles makes it difficult to consumption smooth while benefits appreciate, especially in the strategic model in which individuals have a relatively low elasticity of intertemporal substitution.

Next, Figures 5 shows the cross-sectional benefit claiming hazards for married women (left) and men (right) in the benchmark and strategic models. The top row of Figure 5 shows the hazards by age between age 60 and 69 the bottom row 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 corresponding result from our HRS sample.

For married women, both the benchmark and the strategic models fit the data very well. This is unsurprising. Most married women have relatively little incentive to delay benefit claiming since doing so does not increase total household Social Security wealth since, with relatively high probability, she will switch from her own to her partner's benefit before the couple fully amortizes the forgone income. As well, by claiming benefits, she provides the household with additional income during the retirement window. Nevertheless, the strategic model implies slightly *earlier* claiming by married women than the benchmark model, which

Figure 5: Married men and women: Benefit claiming hazards and cumulative hazards by age

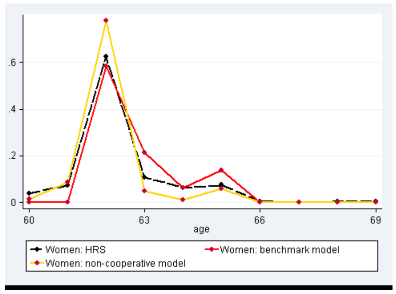

(a) Women

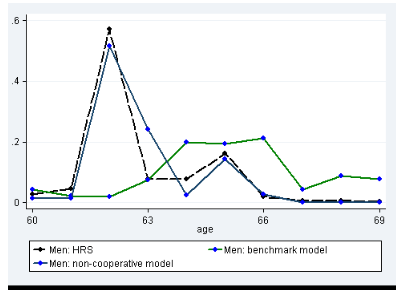

(b) Men

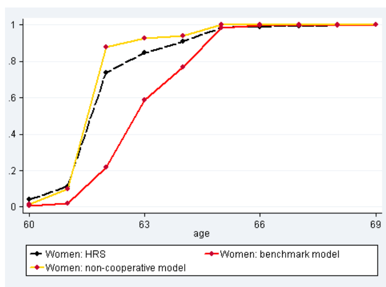

(c) Women

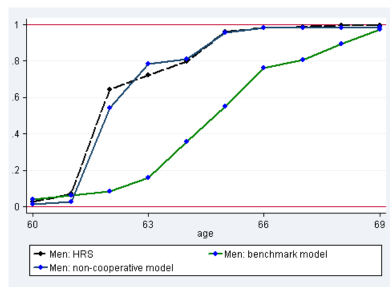

(d) Men

very modestly overpredicts claims at 65. This is because—as emphasized in the financial planning literature—it is not necessarily optimal for the wife to retire at 62 when she is not eligible for the spousal benefit. Although unusual in this cohort, in about 10% of cases the wife is the main earner and it is in the best interest of the couple for her, rather than her husband, to delay claiming.

The story for married men is quite different. Married men in the benchmark model delay benefit claiming into their mid 60s, with the largest share of claims between 64 and 66, 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. The strategic model, by contrast, provides a very good fit to the data, with around 52%, claiming at age 62, as soon as benefits are available, with a small residual spike at age 65 and virtually no

Figure 6: Transitions into retirement for married men: Data and models

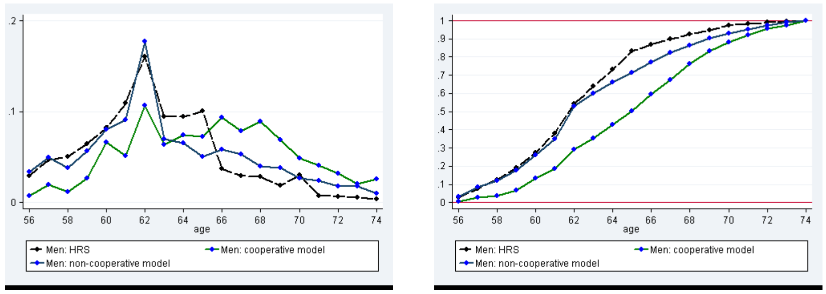

(a) PDF (b) CDF

husbands claiming after that date.

What accounts for the difference between the models? As already discussed, in the strategic model households are much more constrained, while their relatively high risk aversion induces them to maintain a buffer stock of assets against a health expenditure shock, and thereby prevents the drawing down of wealth stocks to finance delayed claiming. More directly, men have less incentive to delay their claiming date since the main beneficiary of the higher long-term benefit is the wife. This is especially true if the jointly optimal strategy involves the husband continuing to work. In the game theoretic context of the strategic model, the husband may have $\{R, A\}$ as a dominant strategy for any jointly optimal allocation that involves acceptable (for both spouses) household-level saving.

To further motivate the retirement implications of non-cooperative behavior, figure 6 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.

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

Difference in years	HRS 1998-2006		Benchmark model		Strategic model	
−6	2.9	(14.7)	0.3	(12.5)	0.7	(13.9)
−4	14.9	(14.2)	2.2	(8.6)	6.5	(7.2)
−2	14.7	(8.0)	5.6	(4.6)	24.0	(5.1)
0	58.9	(37.5)	18.0	(17.8)	61.8	(50.9)
+2	4.8	(6.8)	41.1	(9.0)	5.7	(9.6)
+4	4.9	(7.2)	29.8	(17.5)	1.5	(4.9)
+6	1.3	(11.6)	5.5	(30.1)	0.5	(6.3)

4.3 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 benchmark model. The benchmark model predicts that husbands typically claim benefits later than their wives; by contrast, the strategic model predicts that spouses claim roughly concurrently. It also captures the fact the joint benefit claiming is in fact *more prevalent* than joint labour market exit.

4.4 Robustness checks

[Forthcoming]

5 Conclusions

[Forthcoming]

6 Appendix: 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 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.

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., 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.

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

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., 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.

Maestas, N., Zissimopoulos, J., Karoly, L., 2006. The effect of retirement incentives on retirement behavior: Evidence from the self-employed in the united states and england, mRRC Working Paper 2007-155 and RAND Working Paper WR-528.

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

Meyer, W., Reichenstein, W., 2010. Social security: When to start benefits and how to minimize longevity risk. Journal of Financial Planning 10.

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.

Sun, W., Webb, A., 2009. How much do households really lose by claiming social security at age 62, center for Retirement Research Working Paper.

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.