

Social Security, Endogenous Retirement and Intrahousehold Cooperation

Laura Turner
University of Toronto

Giovanni Gallipoli
Vancouver School of Economics

September 16, 2013

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 joint household utility, subject to the additional constraint that neither partner finds it optimal to deviate from the best constrained household allocation. We show that accounting for “non-cooperative” behavior through this additional constraint can rationalize various choices of older couples observed in the 1932-42 cohort of the Health and Retirement Study. In particular non-cooperative behavior helps with two recurring puzzles in the retirement literature: (i) the clustering of benefit claiming at the early age of 62 despite significant gains associated to delayed claiming by husbands; and (ii) the joint benefit claiming of couples. We contrast our findings to those from a standard unitary model of the household, extended to include a process for declining health, and show that the latter can rationalize neither early nor joint claiming behavior if individuals can independently make benefit and labor force participation decisions.

PRELIMINARY AND INCOMPLETE

1 Introduction

Economists have devoted much attention to retirement patterns of married couples in the US. 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 regardless of health status, and (2) the tendency of husbands and wives to retire and claim benefits concurrently, with 59% of married couples born between 1932 and 1942 in the Health and Retirement Study (HRS) claiming benefits within two years of 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, accounting for the fact that many single women (widows) become eligible at age 60 while all married women become eligible at 62. There is little difference in the timing of claiming by gender or marital status. Table 1 shows the distribution of time differences in the decisions of husbands and wives for both benefit claiming and age of final labour market exit.¹ Not surprisingly given the plots in figure 1 married couples of similar 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 at odds with the findings of the literature on optimal financial planning (e.g. Meyer (1996), Sun and Webb (2009)) which emphasizes the gains to married couples from having the wife claim her benefit at age 62 while the

¹The 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. Simple observation suggests that the graphs look very similar when wider age ranges are considered (see Rivas (2010) among others) or when db pension holders are included.

husband delays his benefit claim late into the retirement window. The main reason for married people postponing benefits, relative to singles, is the Social Security Administration (SSA)’s rule governing survivor (“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 if 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

²In 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.

³Specifically, using the Social Security rules that prevailed in 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.

⁴Young widows, those whose husbands die before early retirement age, can claim a benefit based on 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 if she claimed 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 “not penalized” survivor benefit at full retirement age. We do not model these mechanisms in the present version of the paper.

⁵This argument is complicated when either the husband, or the 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)). The Social Security Administration 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 Rivas (2010) by excluding

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

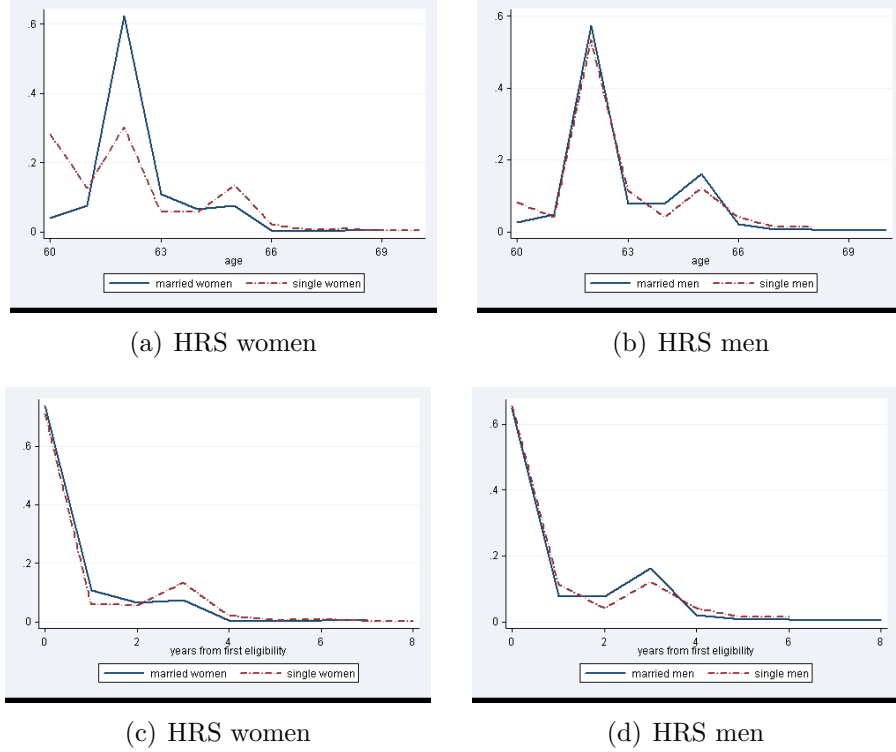


Table 1: HRS: Husband's minus wife's date of Social Security benefit claiming and final labour market exit in years

	-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

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 complementarities in spousal leisure (Maestas (2001), Coile (2004), Gustman and Steinmeier (2000), Schirle (2008) and Rivas (2010)).⁶ Rivas (2010) also shows that the Social Security spousal benefit generates modest incentives toward earlier retirement among con-

db pension holders from our analysis.

⁶Defined benefit pension incentives, which we ignore in this paper, also clearly play a role in explaining labor force withdrawal patterns.

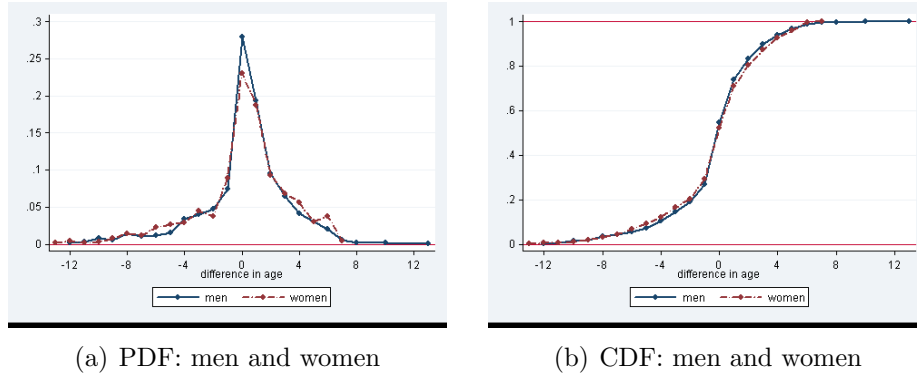
strained households in which the wife qualifies, since she cannot receive the benefit unless the husband has already filed. With the exception of spousal benefit incentives, however, the former explanations have more bite for explaining patterns of early retirement from the labor force than for explaining Social Security benefit claims. Moreover, health incentives cannot easily account for the lack of differences 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 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 p.d.f. and c.d.f. of the differences in age of benefit claiming and final labor market exit within our HRS sample over the range ‘-10 years’ up to ‘+10 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 should be treated as separate choice variables in the transition into formal retirement.

In this paper we show that observed patterns of labor market exit and benefit claiming among couples without defined benefit (db) pensions are much harder to reconcile with life cycle theory when labour force participation (or “career”) decisions and Social Security ben-

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

⁸The 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



efit 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 as 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) with the extension of household-level saving decisions. 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 labour supply or benefit status, given the current household choices of labor force status, assets and (public) consumption. Complicating the problem is the fact that any unilateral deviation in own career or benefit status is likely to trigger 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 crucial to explaining our results: 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 reinforce 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 harder to consumption-smooth in the absence of Social Security income. Second, most medical risk will be borne by the longer-living spouse, the wife. Since husbands, who have lower life expectancies, are typically interested in

⁹Among 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.

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 several cross-sectional patterns discussed 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 interesting welfare implications. Specifically, lack of cooperation over career choices has major welfare consequences for couples, reducing the efficiency of marriage. These welfare costs emerge before retirement since non-cooperation changes the career and saving incentives of couples earlier in the life cycle.

The rest of the paper proceeds as follows. We present the household problem and the nature of the strategic intra-couple 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 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 cross-sectional (individual and joint) retirement patterns of married households — particularly for husbands — between age 60 and 69.¹¹ We also discuss the welfare consequences of non-cooperation.

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

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

2 The household optimization problem

Households enter the model when the head is age 48 (if female) or 49 (otherwise). Females in the model are one year younger than males, 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. 90% of households 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 a time-invariant probability of surviving from age j to $j+1$ for an individual of gender i and health status h . When the member of a couple dies, the spouse lives as single until death. Households differ ex-ante in asset levels, as well as in their social security entitlements and wages. 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 some bundle of leisure and satisfaction derived from ls that may be positive or negative. They have time discounting factors β and face a time-invariant interest rate $r = .04$.

At each age j , males and females find themselves in one of a theoretically infinite number of seven-dimensional variable states. The dimension are, (1) accumulated assets a ¹⁴; (2)

¹²The sample is restricted to households in which the spouses are less than five years apart in age. In the full sample, the mean age difference between heads and spouses is roughly three years.

¹³This assumption simplifies the model and can be justified by the fact that US divorce rates among couples over 50 are much lower than among younger couples. we assume that marital “sorting” has taken place prior to age 50 and that, even when spouses do not cooperate, there are sufficient gains to marriage to avoid dissolution.

¹⁴We exclude housing wealth from a .

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.

Apart from mortality, there are three sources of life cycle risk: productivity risk, which affects w conditional on ls ; career risk which affects ls ; and risk of poor health, which imposes variable out-of-pocket medical costs and increases mortality. Marriage offers partial insurance against income risk since couples benefit from economies of scale in shared household-level consumption.

2.1 ls and ss states

Up to age 75, state ls (in set 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. 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 to transit into non-career work which pays a lower hourly wage on average, but also requires only 35 hours per week. Non-career workers flow into career jobs at an exogenous rate p_C , which may vary by gender. Social Security claim status, ss (in set SS) can also take 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 62.¹⁵ While applying

¹⁵We assume that once individuals reach 62 their applications are never turned down. More specifically, we assume that individuals who retire between 62 and 65 receive the regular retirement age benefit with probability $p_A(h)$ and the age-adjusted benefit with probability $1 - p_A(h)$.

for (disability) benefits before age 62, individuals must choose $ls = R$. Their applications are successful at rate $p_A(h)$, i.e. the probability of receiving disability benefits depends on health status. Applicants who apply from career jobs forfeit their job and fall off the career track into non-career work; they also face a disutility cost c_A from applying, expressed as a portion of average per-period utility. Given that we only follow individuals in their 50s and later, we assume that once an individual has left his career job he does not return to the “career” track. Individuals 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 as defined in our model receive Social Security benefits, which is seldom optimal due to the Social Security earnings test. This also reduces the computational burden.

2.2 Life cycle optimization

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

2.2.1 Singles

A single individual of 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:

$$V_j^{i,ls,ss}(x) = \max_{\{c,a'\}} u(c, ls) + \beta \varphi_i \mathbf{E}_j[V_{j+1}^i(x')] \quad (1)$$

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

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

where

$$I = w(ls)n(ls) + (1+r)a + b(a, e, ss, ls) - \chi(h) \quad (2)$$

I is household period income net of health costs $\chi(h)$: the sum of wage earnings from working n hours in labor force state ls , asset holdings, and benefits. \underline{c} is the minimum level of consumption guaranteed by the government. The benefit entitlement b depends generally on assets (due to means testing), 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. Households age 60 and over (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 $j \geq 60$. Finally,

$$V_{j+1}^i = \max_{\Omega_j(ls,ss)} \left\{ V_{j+1}^{ls,ss} \right\} \quad (3)$$

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

2.2.2 Married Individuals

In a cooperative 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_j^i . The married household has a global value function U_j , a weighted sum of the individual members' Υ s.

$$U_j(x_M) = \lambda \Upsilon_j^f(x_M) + (1 - \lambda) \Upsilon_j^m(x_M) \quad (4)$$

where

$$\begin{aligned} \Upsilon_j^i(x_M) &= u(\mathbf{c}, \mathbf{ls}) + \beta \mathbf{E}_j[\varphi_i \varphi_{-i} \Upsilon_{j+1}^i(x'_M) + \varphi_i(1 - \varphi_{-i}) V_{j+1}^i(x')] \\ \Upsilon_j^i(x_M) &\in \Upsilon_j^{i,LS,LS,SS,SS}(x_m) \\ \{ls_f, ls_m, ss_f, ss_m\} &= \arg \max_{\omega_M(j)} U_j(x_M) \end{aligned} \quad (5)$$

In this expression, consumption c and (dis/)utility from work are not chosen directly by the individual but come out of the household-level maximization of U . 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 determined at the individual-level, as 1-indexed states ($-i$ denoting the spouse). With probability $\varphi_i(1 - \varphi_{-i})$ each period one spouse dies leaving the remaining spouse a widow or widower, but there is no divorce. 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-pecuniary reasons. In practice, we choose $\lambda = .5$, which generates realistic mean labour supplies for husband and wives, and the model collapses back to the simple unitary set-up.

The household budget set for married individuals, B_M is:

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

where

$$I = w_f n(ls_f) + w_m n(ls_m) + (1+r)a + b(a, e_f, e_m, ss_f, ls_f, ss_m, ls_m) - \chi^M(h_m, h_f) \quad (6)$$

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 from the feasible set $\omega_M(j)$. 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 .

2.2.3 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 would be awarded if one could make decisions about own retirement and benefit claims unilaterally, or if one 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 the joint outcome, and they are fully able to commit to maintaining that joint outcome. We now consider the alternative scenario where this assumptions no longer hold, and in which members of the couple may deviate from the optimal 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 lack-of-commitment to the planner’s solution as $\hat{\Upsilon}$, the spouses, at (head’s) age j and planner’s allocation $Y = \{c, a', ls_m, ss_m, ls_f, ss_f\}$, jointly solve:

$$\begin{aligned}\hat{\Upsilon}_j^f &= \max_{\hat{\Omega}^f} \Upsilon_j^{f, \hat{ls}_f, \hat{ss}_f; \hat{ls}_m, \hat{ss}_m | \hat{ls}_f, \hat{ss}_f} \\ \hat{\Upsilon}_j^m &= \max_{\hat{\Omega}^m} \Upsilon_j^{m, \hat{ls}_m, \hat{ss}_m; \hat{ls}_f, \hat{ss}_f | \hat{ls}_m, \hat{ss}_m}\end{aligned}\tag{7}$$

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 associated 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 which 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 $\Upsilon_m = \hat{\Upsilon}_m$.¹⁶

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 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 tend to increase with age. This assumption allows us to model ss and ls decisions 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

¹⁶The 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 of *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.

higher than her payoff from non-claiming, with no Nash indeterminacies (see, for instance, the problem in Gustman and Steinmeier (2004)).

In this way we can therefore focus on the simple one-period Nash games that spouses play over $\hat{l}s$ (the solution to 7). We start by imposing 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 sub-game, the spouses revert to the planner's (jointly optimal) solution for that sub-game.
3. Assumption #3: If more than one (pure) Nash equilibrium exists in a sub-game, the players revert to the planner's solution for the sub-game.

In a couple consisting of non-career workers, or retirees under 75, the partners play a simple 2×2 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 l s (with its implied s s), and that the resulting Nash-equilibrium outcome differs from the planner's. Figure 3 illustrates four types of outcomes which can arise: the first two give rise to a potential deviation and the other two 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 $\{R, R\}$. Sub-figure 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 2×2 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 $\{NC, NC\}$.¹⁷

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 the implementation of the planner's solution, as assumptions (#1 – #3) prevent deviant outcomes in figures 1(c) and 1(d).

Assumptions (#1 – #3) are motivated by three simple ideas: (1) the single-shot Nash game played once in each period is actually an approximation to an uncoordinated repeated static game which 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 $ss = A$; (3) spouses find uncertainty over their ls status costly and therefore avoid it, converging instead on what they know to be the jointly optimal outcome.

The 2×2 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 which cannot be undone. The career worker therefore

¹⁷The 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

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

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

(a) Husband dominant retirement strategy (b) NE: Husband and wife dominant retirement strategies

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

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

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

controls the nature of the play, creating a sequential game in which she moves first. Figure 4 shows potential solutions to the game when the wife is a career worker and the husband has a bridge-job. The wife solves this game by backward induction, first resolving the 2×2 subgame that results if she retires, and then comparing this outcome to the outcome occurring if she does not retire. If the wife is better off playing the 2×2 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.

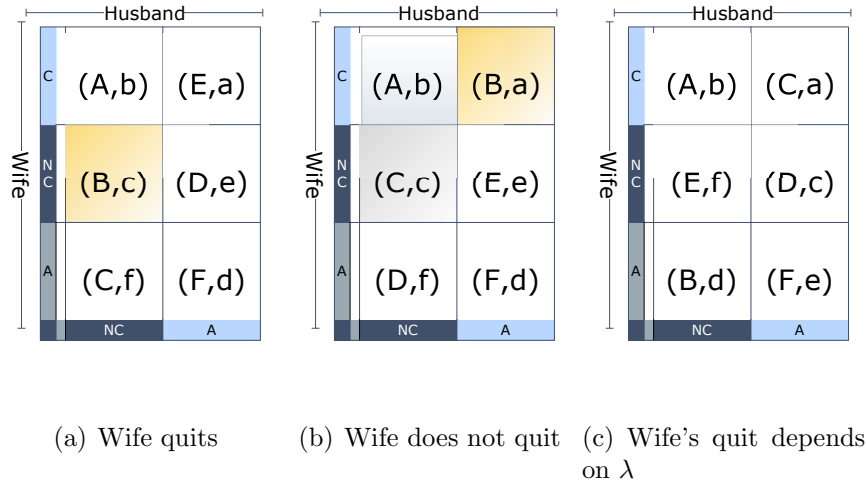
In the first panel of figure (4), the wife's payoff if she does not quit is D , since in this case her husband plays A . However, if she quits, the couple will converge on the Nash-equilibrium

$\{NC, NC\}$, which raises her payoff to B . She will therefore quit. Note that the outcome of the super-game is ex-post inefficient, even though the solution to the sub-game 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 to 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 2×2 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 2×2 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



The final game which arises in this model is played between a career husband and a career

Figure 5: Non-cooperative outcomes for career workers



wife. The resulting 3×3 Nash game can be conceived as a two-stage sequential game. At the first stage, 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 2×3 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 3×2 game described in the previous paragraph (playing NQ) and the 2×2 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 3×2 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 2×2 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 $\{\hat{l}s, \hat{s}s\}$ outcomes for different points in the couple's state space x_m and household inter-temporal allocation $\{c, a'\}$, given x_m . The planner's 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 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. They 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. They change jobs, but remain within occupation category¹⁸ 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. They 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 individual- or firm-specific. Roughly 55% of male workers, and 30% of female workers between 48 and 55, are identified as career workers; this share falls sharply in the late 50s and

¹⁸We use the 17-category occupation listing provided in the RAND HRS.

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:

$$\ln w_{it}^{g,ls} = \beta_0^{g,ls} + \beta_1^{g,ls} age_{it} + \beta_2^{g,ls} age_{it}^2 + \epsilon_{it}^{g,ls} \quad (8)$$

$$\epsilon_{it}^{g,ls} = \nu_{it}^{g,ls} + \eta_{it}^{g,ls} \quad (9)$$

$$\nu_{it}^{g,ls} = \rho^{g,ls} \nu_{it-1}^{g,ls} + \varsigma_{it}^{g,ls}; \quad \varsigma_{it} \sim N(0, \sigma_\varsigma^2), \quad \eta_{it} \sim N(0, \sigma_\eta^2) \quad (10)$$

for $g \in \{f, m\}$, $ls \in \{C, N\}$. Since these are reduced-form estimates, we search for structural parameters which 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 non-career jobs. The average career wage 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

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

	Men		Women	
	career	non-career	career	non-career
<i>cons</i>	1.53 (.419)	-4.04 (.800)	3.01 (.416)	6.09 (.642)
<i>age</i>	.0558 (.0140)	.174 (.0263)	$3.98E^{-3}$ (.0138)	-.135 (.0213)
<i>age</i> ²	$-5.16E^{-3}$ ($1.18E^{-4}$)	$-1.26E^{-3}$ ($2.16E^{-4}$)	$-1.15E^{-5}$ ($1.06E^{-4}$)	$1.02E^{-3}$ ($1.75E^{-4}$)
ρ	.944	.958	.906	.956
σ_{ζ}^2	.044	.105	.053	.071
σ_{η}^2	.152	.592	.128	.532
<i>mean</i>	\$28.4	\$18.4	\$19.4	\$10.4

age to about a 20% differential at age 60).

The second component of earnings risk is the probability of losing a career job, p_C^g for gender g . 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 in the HRS report leaving their jobs due to plant closures or lay offs, which suggest some involuntary separations. Poor health and burnout are much more common explanations.¹⁹

3.1.2 Health and medical costs

We assume that five percent of households experience health expenditure shocks (χ) at every age, with the estimated size of χ increasing with age. Mortality also rises when households experience the shock, at a rate calibrated to replicate the likelihood of dying following such

¹⁹Though we do not account for it in the model, the share of “involuntary” separations among all separations falls with age: for men, from 24.5% for under 55 to only 14% of men 55 and over, though this is likely due to selection.

a shock in the HRS. Equations 11- 12 show the transition matrices and evolution of costs associated with poor health (state 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.

$$\begin{bmatrix} pr(h_g|h_g) & pr(h_b|h_g) \\ pr(h_g|h_b) & pr(h_b|h_b) \end{bmatrix}_{married} = \begin{bmatrix} .977 & .023 \\ .929 & .071 \end{bmatrix}, \begin{bmatrix} pr(h_g|h_g) & pr(h_b|h_g) \\ pr(h_g|h_b) & pr(h_b|h_b) \end{bmatrix}_{single} = \begin{bmatrix} .965 & .335 \\ .818 & .182 \end{bmatrix} \quad (11)$$

$$\begin{aligned} \chi(age, h_b)_{single} &= -3266 + 171.8 * age - 2.85 * age^2 + .0160 * age^3 \\ \chi(age, h_b)_{married} &= -14602 + 634.4 * age - 8.79 * age^2 + .0405 * age^3 \end{aligned} \quad (12)$$

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 .

$$\begin{aligned}\varphi_{h_g,i}^j &= Pr(Surv|h_g) = \frac{\widehat{Pr}(Surv)\widehat{Pr}(h_g|Surv)}{Pr(h_g)} \\ \varphi_{h_b,i}^j &= Pr(Surv|h_b) = \frac{\widehat{Pr}(Surv)\widehat{Pr}(h_b|Surv)}{Pr(h_b)}\end{aligned}\tag{13}$$

where hats indicate that the information comes from the HRS.

3.1.4 Preferences

We adopt the following gender- and age-varying preference function u_g for gender g :

$$u_g(c_g, l_g) = \psi^g \frac{c_g^\alpha}{1 - \alpha} + \gamma_1^g 1_{career} + \kappa_2^g 1_{n-career}\tag{14}$$

where

$$\psi^g = 1.0 + \gamma_1^g \times age + \gamma_2^g \times age^2\tag{15}$$

$$\kappa^g = \kappa_0^g + \kappa_1^g \times age + \kappa_2^g \times age^2\tag{16}$$

$$\gamma^g = \gamma_0^g + \gamma_1^g \times age + \gamma_2^g \times age^2\tag{17}$$

In the above expressions, α is the coefficient of relative risk aversion in consumption, and ψ measures the age-dependent weight placed on consumption. 1_{career} is an indicator function for being in a career job and $1_{n-career}$ is an indicator function for being in a non-

career job. γ and κ measure the utility gains (or costs) associated to career and non-career work respectively at each age.

3.1.5 Model Estimation

Estimation of the models is performed through simulated method of moments (SMM). The program described above gives four sets of parameters to estimate for each of the fully cooperative and non-cooperative models. They are: (1) the path of wages for career- and non-career men and women that return the parameters in the first three rows of table 2; (2) valuation of consumption for men and women, ψ (3) utility gains from career and non-career γ and κ ; (4) disutility of applying for SSD benefits and likelihood of acceptance into the SSD program in the high health cost (poor health) state c_A and $p_A(h_b)$. These parameters and their estimated values are listed in table 3.

Tastes for consumption ψ determine the life cycle asset holdings of married and single households. γ and κ , along with career job loss probability p_C pin down gender-specific shares in each of the three ls states at each age up to age 70. c_A and p_A give the proportion of high-health cost and low-health cost households whose head is under 62 (in the top 5% and bottom 95% of health spenders) in which at least one member successfully applies for disability benefits under SSD or SSI and an overall rejection rate by disability spell²⁰—of around 50%, which Gruber (2000) and Bound et al. (2002) separately cite as the prevailing post-appeal rejection rate of SSD applicants.

²⁰Because 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 two consecutive periods.

Table 3: Parametrization of the jointly efficient (cooperative) and non-cooperative (strategic) models

Parameter	Benchmark	Non-cooperative
$\ln w^{f,nc}$	$3.47 - .0173 \times age - 4.29E^{-5} \times age^2$	$3.43 - .0178 \times age - 5.05E^{-5} \times age^2$
$\ln w^{f,c}$	$3.14 - .0119 \times age - 4.96E^{-5}$	$3.23 + 2.59E^{-4} \times age - 7.68E^{-5} \times age^2$
$\ln w^{m,nc}$	$1.91 + 0.014 \times age + 6.54e^{-5} \times age^2$	$1.94 + 0.010 \times age + 6.54e^{-5} \times age^2$
$\ln w^{m,c}$	$4.67 - 0.069 \times age + 6.10e^{-4} \times age^2$	$5.29 - 0.067 \times age + 5.64e^{-4} \times age^2$
γ^f	$1.21 + .153age - 4.43E^{-4}age^2$	$2.33 + .0988age - 2.16E^{-4}age^2$
γ^m	$5.90 + .0230age - 3.01E^{-5}age^2$	$9.12 - .0803age - 1.15E^{-4}age^2$
κ^f	$12.1 - .0294age + 4.07E^{-5}age^2$	$9.09 - .0174age + 3.81E^{-5}age^2$
κ^m	$.341 + .0719age - 3.05E^{-4}age^2$	$-2.45 + .0227age - 3.24E^{-4}age^2$
ψ^f	$1.0 + .00180age - 4.40E^{-5}age^2$	$1.0 + .00151age - 4.05E^{-5}age^2$
ψ^m	$1.0 + .00189age - 6.06E^{-5}age^2$	$1.0 + .00152age - 5.04E^{-5}age^2$
$p_A(h_b)$.495	.479
c_A	.998	1.06

Finally, in both models we set some parameters directly. Specifically, we set the process for the innovation in wages ρ and σ_ζ^2 equal to the values in table 2; we set the coefficient of risk aversion $\alpha = 2.0$ and discount factor $\beta = .96$; we assume that married couples benefit from economies of scale in consumption such that one unit of household consumption translates into .6 units of individual consumption for each spouse, and we set the likelihood of losing a career job equal to 1% per year for both men and women. Individuals are not accepted into SSD in the good health state: $p_A(h_g) = 0$.

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:

$$\begin{aligned}
PIA &= 0.9 \min(0.2\bar{w}, E) \\
&\quad + 0.32 \max(\min(E - 0.2\bar{w}, 1.3\bar{w} - 0.2\bar{w}), 0) \\
&\quad + 0.15 \max(E - 1.3\bar{w}, 0) \\
a_t < \bar{a}^{mt} : \quad &b = \max(PIA, SSI) \\
a_t > \bar{a}^{mt} : \quad &b = PIA
\end{aligned} \tag{18}$$

where \bar{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.²¹ Since SSI is means-tested as opposed to work-tested, \bar{a}^{mt} ,

²¹The 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.

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.²² 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. 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:

$$E_{age} = \frac{E_{age-1}(age - 30) + w_{age}^{ls}n_{age}^{ls}}{age - 31}$$

and for workers over 60 is:

$$E_{age} = \max\left(\frac{E_{age-1}(age - 30) + w_{age}^{ls}n_{age}^{ls}}{age - 31}, E_{age}\right)$$

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

²²Huggett and Ventura (2000) and others discuss a similar technical difficulty in their calibrated analysis of the Boskin Social Security reform proposal.

appeal up to four times, to four different levels of SSD adjudicators, a process that can take several years.²³ 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%, which we replicate in the model.

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 claim age, although only once the main earner has claimed his own benefit, or at age 65.²⁴ 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 of her life.

Unlike spousal benefits, that are based on the husband's PIA, survivor benefits encompass

²³An 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.

²⁴At 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.

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 claw-back,²⁵ though the taxation of benefits is likely to change (see Appendix A). 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 higher than their average Social Security earnings to date).

²⁵In 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.

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 A.

3.2.2 Optimal claiming by singles and couples

Table 3.2.2 shows the optimal claiming age of males, and associated wealth gains relative to claiming at 62, for men in our HRS cohort based on the Social Security rules described above. The calculations assume that wife claims her benefit exactly at age 62. Focusing on 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 at age 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. Overall, the wealth results from the table are 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.

It is apparent from these calculations that, regardless of the partners’ respective health status, it is jointly optimal for the husband to delay his benefit claim into his late 60s.

	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	

4 Results

We now turn to assessing the implications of the model described in the previous section, and to considering its implications for retirement and benefit claiming. The following results give a flavor of the model's implications for the retirement behavior of couples and the welfare implications of non-cooperation in marriage.

4.1 Parameterization

We begin by discussing the parametrization in table 3. Individuals' disutility from career work increases fairly rapidly as individuals age while the disutility from non-career work increases more slowly. Figures 6 and 7 show the paths of workers in career and non-career

Figure 6: Shares of career and non-career workers by age: benchmark model

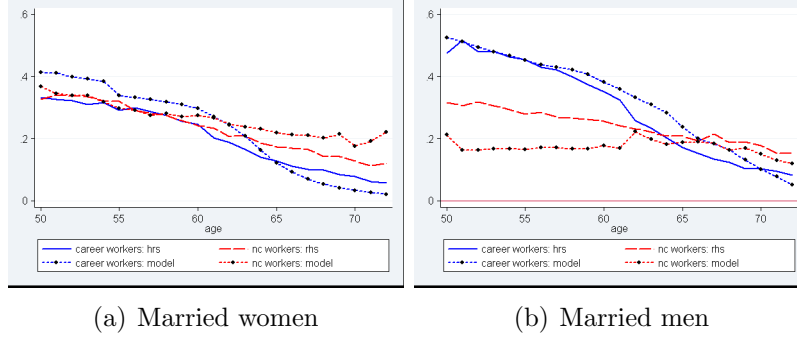
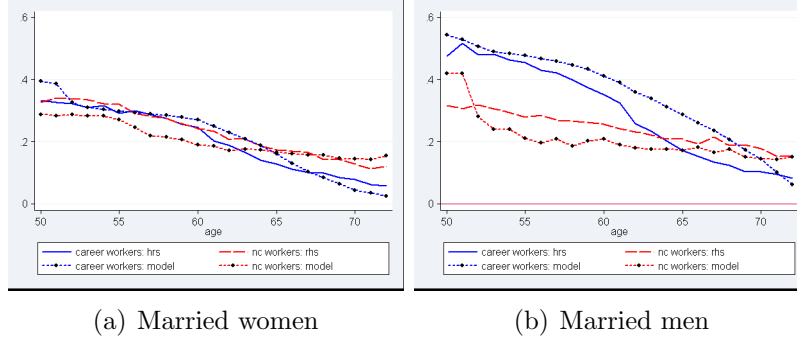


Figure 7: Shares of career and non-career workers by age: non-cooperative model




work. Figure 8 shows the path of assets in the models.

4.2 Cross sectional benefit claiming patterns

Figures 9 and 10 show the (cross-sectional) benefit claiming hazards for single and married women (left) and men (right) in the standard cooperative model and the strategic model developed in this paper. In each figure, the black line shows the results from our HRS sample of non-db workers.

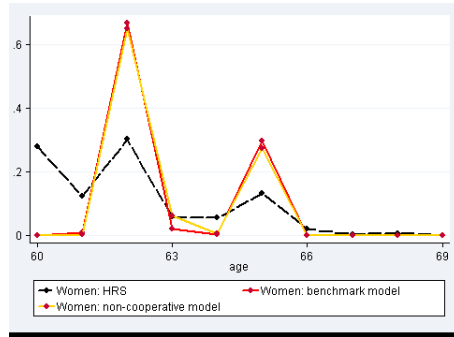
For married women, the fit of each of the two models is similar, which is not surprising. Most married women have relatively little incentive to delay benefit claiming since doing so does not increase household Social Security wealth so long as the higher-earning husband

Figure 8: Median assets by age for married couples




(a) Married couples

Figure 9: Cross sectional benefit claiming hazards: single men and women




(a) Women

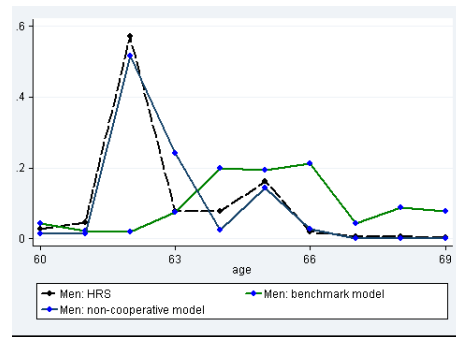


(b) Men

Figure 10: Cross sectional benefit claiming hazards: married men and women



(a) Women



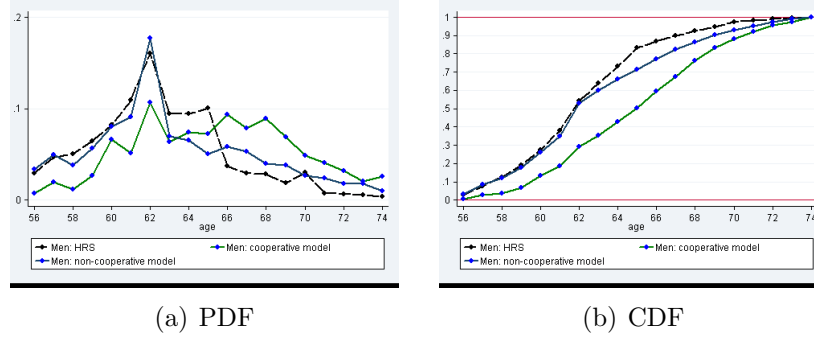
(b) Men

delays claiming. About 45% of wives expect to benefit from the spousal benefit while their husband lives; that is, they expect to receive a benefit equal to 50% of the husband's. 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 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 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 the decisions to exit the labor force and claim benefits 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 wealth, 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

Figure 11: Transitions into retirement for married men: Data and models



a standard 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 11 shows the hazard rate and the cumulative hazard for final labor market exit of married men in the two models, as well as in the 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.

4.3 Joint retirement

We next turn to examining the models' implications for timing of benefit claims, 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 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 in 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 decrease disutilities from work.

Corresponding results for the intra-household timing of final labor market exit are less clear. Neither model performs very well, though the strategic model, while *overstating* the amount of joint labor market exit, performs better in the sense of matching intra-household "claim date differences" in the data. Figure 11 shows that the non-cooperative model does a good job of predicting the dates of final retirement from the labor force among men. By

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)

contrast, the optimal strategy for many cooperative households is for one spouse (usually the husband, depending on relative wages and health) to work through his 60s while delaying benefit claiming. This failure of the cooperative model, however, is likely to be due partly to the absence of leisure complementarities, which have been shown to be important determinants of joint labour market exits within married couples, and which we will address in future work.

5 Conclusions

This paper investigates the empirical relevance of possible inefficiencies due to non-cooperation within the household. We offer a potential resolution to the puzzle of both early and joint retirement of couples discussed in the literature: this is achieved in a framework in which labor market exit and benefit claims’ decisions are separately made.

Our work is motivated by the empirical observation that early retirement and Social Security benefit claiming patterns appear at odds with standard household inter-temporal theory. We show that removing the common assumption of full cooperation between spouses

can crucially change the model’s predictions about households’ retirement behavior.

In a non-cooperative model spouses deviate from household-level optimization rules in deciding when to unilaterally retire or transition out of career work into less demanding jobs, because they cannot internalize the entire benefit from their own productivity (i.e. their career status). This mechanism produces early retirement and benefit claiming patterns, as well as substantial synchronicity in such choices within couples.

The model implies that non-cooperative behavior within households can potentially generate major inefficiencies and welfare costs. Future work will involve studying and quantifying the inefficiencies of non-cooperation within households and how deviations from optimal behavior can strongly influence the way couples respond to policy incentives.

References

- 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, 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 retirement-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.
- Meyer, M., 1996. Making claims as workers or wives: The distribution of social security benefits. *American sociological review*, 449–465.
- 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.

Sun, W., Webb, A., 2009. How much do households really lose by claiming social security at age 62? Available at SSRN 1408729.

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.

A Other (Non-Social Security) policies and institutions

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 married individuals 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.

Other benefits. 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.