Social Security, Endogenous Retirement and Intrahousehold Cooperation

Giovanni Gallipoli
University of British Columbia

Laura Turner
University of Toronto

February 15, 2011

Abstract

This paper studies the retirement incentives generated by the U.S. Social Security system in a framework which allows for different degrees of cooperation and strategic interaction between spouses. We examine the relative empirical performance of models specified under different assumptions about preferences (specifically over complementarities in leisure during retirement) and intrahousehold commitment to a jointly optimal retirement path. We assess the models' relative ability to replicate a large set of economic choices observed at or around the time of retirement. These choices include the tendency of individuals to retire close to the first availability of regular Social Security retirement benefits; for spouses to retire approximately at the same time; and for transitions into retirement to be made either through "bridge jobs"—exit from a full-time career job several years before full retirement—or through application for Social Security disability benefits. Our results suggest a role for non-cooperative behavior in households in which main-earners are subject to large transfer liabilities towards second-earners, which effectively introduce a 'tax wedge' on earned income and affect both labor supply and retirement decisions. The different models also suggest different welfare implications for current Social Security policy in partial equilibrium, with the utility value of most features of the current system increasing with complementarities in leisure, but decreasing in the degree non-cooperation between household members approaching retirement.

1 Introduction

A significant body of research has studied the retirement incentives of public pension systems, especially for the U.S. Social Security system.[1] This literature often highlights the difficulty of explaining observed patterns of retirement: the large peak in regular retirement at age 62, the first year at which Social Security retirement benefits are available to most workers; and the tendency of husbands and wives to retire together when the husband reaches 62.[2] The peak in retirement at 62 is puzzling in the context of a standard life cycle model with heterogeneous agents, because individuals can increase the size of the (riskless) claimed benefit at a rate of around 6.5% a year (under 2000 SSA rules) by delaying retirement up to age 70, which is, on average, approximately actuarially fair and should therefore be an attractive option for the subset of individuals who expect to live a long time. Joint retirement is difficult to reconcile with standard preferences since men and women have different incentives to delay retirement: women live longer on average and are more likely to rely primarily on household-level savings or on a partner's accumulated benefit income both before and after widowhood.

This paper builds on the standard life cycle model of a two-person household with a retirement decision[3], and examines two modifications to this benchmark model that can potentially reconcile observed retirement behaviour. The first allows married couples to

[1]See, for instance, Coyle and Gruber (2000) on incentives for early or delayed retirement for workers in their 60s, or Autor and Duggan (2003) on incentives for younger workers to apply for disability benefits.

[2]Gustman and Steinmeier (2005) provide a recent overview of the structural and empirical literature on retirement patterns, and discuss a solution to the early retirement puzzle based on heterogeneous discount rates.

[3]See Kreider (1999) or Burkhauser et al. (2004) for examples of dynamic models with a retirement decision, or Rivas (2007) for a married household.

benefit from complementarities in leisure during retirement. Several authors have previously advanced leisure complementarities as a potential explanation for the early joint retirement patterns of couples.[4] The mechanism driving this result is straightforward: the utility boost from joint retirement induces couples to forego the economic benefits of delaying retirement, and raises their incentive to retire together or within one or two years of each other. Indeed, when calibrated to replicate the share of individuals leaving work by age 64, we find that the 'complementarities' model can easily explain the observed retirement patterns. However, it does a poor job of explaining some other aspects of life cycle retirement behaviour: in particular, it predicts too much delayed claiming of Social Security benefits; too little use of 'bridge jobs' as transitions into retirement; and much shorter durations on disability insurance (SSDI) than are observed in the data.

An alternative model places the retirement problem, and related career decisions, in a household model in which retirement choices are individual-level decisions that may be made "non-cooperatively". In all models we adopt a standard approach from the dynamic collective household literature by assuming that couples make choices according to a fixed 'power-sharing' rule or 'marriage contract' that assigns a weight λ in the household problem to one member's private continuation function and (1-λ) to the other. The twist to this standard framework is that, while household members follow the rules of their marriage contract conditional on their career or retirement status, they may find it optimal to violate the contract in choosing when to retire from work or quit a career job. Effectively, this is because the marriage contract imposes a positive externality on the less productive spouse

[4]See Maestas (2001) for the U.S. case, and Jiménez-Martín et al. (1999) for Europe.

and, conversely, prevents individuals from internalizing all the benefits of career work. Since the household planner imposes the efficient solution to the period-by-period problem, a high-productivity ('career') worker has obligations to provide for his spouse that he can forego by changing his labour force state (lfs) and the parameters of the household planner's problem.

The theoretical framework we develop to examine the implications of non-cooperation over career decisions is similar to Gustman and Steinmeier (2004), with the addition of productivity and health risk that increases over the life cycle and a more flexible model of retirement in which individuals can move between lfs states—i.e. can 'un-retire'—up until age 75. It turns out that this model can potentially capture several features of observed retirement patterns: step retirement by the use of non-career bridge jobs, early retirement under the Social Security Disability Insurance (SSDI) program; and return to work by older workers who have previously claimed their social security benefits. The introduction of 'un-retirement' and bridge jobs produces additional testable implications for the model. In the non-cooperative model, in contrast to the complementarities model, individuals who enter SSDI rarely exit again before death or reaching regular retirement age as they benefit from the shift of primary earner status to their spouse. They are also much more likely to shift into lower-paid bridge work prior to retirement in order to escape primary earner responsibilities. The resulting lower earning potential in the later part of life in turn reduces the incentives of the spouses to delay retirement leading to a large spike in retirement at age 62.

The models developed below—the benchmark, complementarities- in-leisure, and non-cooperative model—are designed to capture the salient features of the U.S. social safety net, SSDI, early and delayed retirement (SSR) benefits, spousal and widow benefits. Individuals are subject to productivity and health risk processes that intensify over the life cycle. Within this framework, we examine the effects of retirement policy and resulting patterns of retirement in a world in which households are fully cooperative over all aspects of married life, and a model in which they are cooperative only over continuous variables—consumption, labour supply and asset accumulation—that are chosen within, and conditional on, household members' labour force status. Both the complementarities and the non-cooperative models perform reasonably well at replicating rates of retirement from the labour force among married couples. However, the non-cooperative models do a better job at matching patterns of SSR claiming, SSDI application and exit patterns and a somewhat better job of explaining bridge job behaviour. The different models also have non-trivial implications for policy. In the cooperative models, with and without complementarities, the status quo produces a net welfare gain for essentially all individuals born into the economy, relative to either an across-the-board benefit cut or elimination of either SSD or early retirement benefits under SSR. This is not true for the non-cooperative case. In particular, husbands are better off in the non-cooperative world if the spousal benefit under Social Security were eliminated. This is because the existence of the spousal benefit reduces the wife's attachment to the labour force and her likelihood of being in a career job in late middle age, and forcing the husband into primary earner status under the marriage contract.

The rest of the paper proceeds as follows. We present the household problem and the nature of the non-cooperative retirement game 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. In section 4.1 we compare the performance of the three models in replicating several features of household behaviour among the 1998-2006 panel of the Health and Retirement Study—in particular, patterns of social security benefit (SSR) take-up during the retirement window ages 62-69 within and across households; bridge job behaviour of married men and women; and use of the SSDI program for individuals with poor health. In section 4.2, we compare the welfare results from the two models under the existing Social Security program (the 'benchmark') and three policy experiments described above. Section 5 concludes.

2 Individuals and households

Life in the model is as follows: households, either one-member or two-member, form at age 25 and survive to a maximum age of 90 years. A model period is one year. Lifespan of individuals and households is uncertain ex-ante. $\varphi^j_{\delta,i}$ (hereafter φ_i for brevity) is the time invariant probability of living from age j to $j+1$ for an individual of gender i and disability/health status δ. 15% of households in the model are single, and the remaining 85% are married in the first period of life.[5] Individuals of each gender are identical ex-ante, but households form with different levels of assets: 80% of individuals are born into the economy with zero assets. 20% are born with assets equivalent to $40,000, which gives a reasonable approximation to the wealth distribution of the lower 98% young households observed in the 1989-2001 wealth files of the PSID. Individuals value consumption c and leisure l. They have gender-specific time discounting factors β_i and face a time-invariant interest rate $r = .042$. Disposable time, $\tilde{T}$, is the amount of time individuals can divide between labour (n) and

[5] Actual average ages of first-marriage which are 25 for women and 27 for men.

leisure, net of sleep, time lost to sickness or disability, and mandatory non-labour, non-leisure activities (see section 3.1 for further discussion of these.)

At every age j, males and females find themselves in one of a theoretically infinite number of six-dimensional variable states. The dimension are (1) accumulated non-human wealth holding a; (2) wage/productivity level w; (3) health/disability status δ; (4) accumulated Social Security credit from previous work E; (5) marital status $ms \in \{0, 1, 2\}$ (0 for single; 1 for low-quality marriage; 2 for high-quality marriage), and (6) labour force status lfs, which takes three values: "career-job", "non-career job" or "retired". There are two major sources of life cycle risk: productivity risk, which affects w, and disability/health risk, which imposes a cost in terms of disposable time, $\chi(\delta)$, increases mortality risk, and imposes medical costs on individuals under 65 in who are currently in non-career jobs. The ms state has three potential direct effects on welfare, the first two of which are common to all three models developed. The first arises through fertility, which generates dependency costs of children $\widetilde{\eta}$ early in the life cycle. The second arises due to the economies of scale in shared household-level consumption.[6] Third, married individuals may receive utility from leisure enjoyed jointly.

Current lfs is endogenous for individuals under a maximum retirement age of 75. Individuals are retired when they are not working and are receiving benefits under the Social Security or Supplemental Security Income (SSI) programs. Non-retired individuals may be in either career- or in non-career jobs. Wages are exogenously lower in non-career jobs, and

[6] The "consumption boost" from marriage is calibrated to generate reasonable spousal labour supplies. In future work, the extent to which marriage generates consumption economies of scale can relate to the quality of the match.

individuals do not have access to health insurance in these jobs. Individuals can choose to move from a career to a non-career job by quitting the career job; while in a career job individuals work a minimum of 40 hours per week, but may choose to work part time, full time or not at all while doing non-career work. Non-career jobs become career jobs at rate p_C when individuals work at least 40 hours a week and at rate 0 otherwise.

In terms of the model state space, $lfs \in \{C, NC, R, A\}$. C (NC) denotes an individual who is currently working in a career (non-career) job. R denotes an individual who is fully retired and receiving Social Security retirement benefits.[7] The final state, A, denotes an individual who is currently not working while applying for benefits. Unsuccessful SSD applicants who apply from the career job state forfeit their job and fall off the career track into non-career work.

2.1 Life cycle optimization

Below, we sketch out the dynamic optimization problem for single and married households. We begin with single households whose problem is the simplest, and also is the same in all the models, regardless of my assumptions about intrahousehold cooperation and preferences.

2.1.1 Singles

In the standard cooperative, the non-cooperative, and the complementarities in leisure regimes, a single individual with gender i and age j has a generalized value function V_j^i,

[7] Once individuals claim retirement benefits, they do not work. Under SSA rules, individuals do not have to retire from their career job to claim SSR benefits. However, they are subject to a reduction in their benefit of \$13,560 for every dollar earned before regular retirement age. After regular retirement age, many career jobs contain incentives to force or induce career workers to retire. These incentives are also captured in the model.

which can be carried across lfs states. The state vector is $x = \{a, E, w, \delta, 0\}$. The dynamic optimization problem of the individual in a career job (with corresponding value function $V_j^{i,C}$) and state vector x is:

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

subject to budget set B_S:

$$\begin{aligned}
\xi_1 : &\quad (\tilde{T} - \chi(\delta) - l)w + (1+r)a + b(.) = c + a' \\
\xi_2 : &\quad l < \tilde{T} - \chi(\delta) \\
\xi_3 : &\quad a' \geq 0
\end{aligned} \tag{2}$$

where

$$V_{j+1}^i = \max\left\{V_{j+1}^{i,C}, V_{j+1}^{i,NC}, V_{j+1}^{i,A}\right\} \tag{3}$$

The expectation $\mathtt{E}$ at age j is taken with respect to δ and w. The problem for a non-career worker is similar:

$$V_j^{i,NC}(x) = \max_{\{c,l,a'\}} u(c,l) + \beta_i \varphi_i \mathtt{E}_j[V_{j+1}^i(x')] \tag{4}$$

subject to B_S, and

$$V^i_{j+1} = \max\left\{p_C(l^*)V^{i,C}_{j+1} + (1-p_C(l^*))V^{i,NC}_{j+1},\;\; V^{i,A}_{j+1}\right\} \tag{5}$$

The dynamic program for a retiree is:

$$V^{i,R}_j(x) = \max_{\{c,a'\}} u(c, l^{max}) + \beta_i \varphi_i \mathtt{E}_j[V^i_{j+1}(x')] \tag{6}$$

subject to B_S, and

$$\begin{aligned} V^i_{j+1} &= \max\left\{V^{i,NC}_{j+1}, V^{i,R}_{j+1}\right\} && \text{if} \quad j < 75 \\ V^i_{j+1} &= V^{i,R}_{j+1} && \text{if} \quad j \geq 75 \end{aligned} \tag{7}$$

Finally, the dynamic program for an applicant is:

$$\begin{aligned} V^{i,A}_j(x) &= p_\delta V^{i,R}_{j+1}(x) + (1-p_\delta)(u(c, l^{max}) + && \text{if} \quad j < 62 \\ &\qquad \beta_i \varphi_i \mathtt{E}_j[V^{i,NC}_{j+1}(x')] - c_A) && \\ V^{i,A}_j(x) &= V^{i,R}_j(x) && \text{if} \quad j \geq 62 \end{aligned} \tag{8}$$

The solution to the optimization problem gives optimal intratemporal choices c^* and l^*. The term $l^{max} = T - \xi(\delta)$ is the leisure the individual obtains when he or she does not work. The parameter $b(.)$ corresponds to the net transfer the individual receives from the government, after taxes are paid and transfers, including Social Security benefits, are received. It is described in detail in Section 3. Career workers (under the maximum retirement age) have three lfs options each period: remaining in their full-time career job, quitting to a lower-paid, flexible-hours non-career job, or applying for retirement benefits. Non-career workers have two options: to continue in their non-career jobs (with potential to move back into career jobs conditional on l^*) or to apply for retirement benefits. Individuals who have not yet reached the early retirement age of 62 cannot access their regular social security benefits, but may apply for disability benefits. Application requires a full year not working, regardless of whether the application is successful. The probability of acceptance p_δ depends on current health status and whether the worker is over or under age 55. As well, there is a one-shot non-pecuniary cost associated with rejection, c_A, which we take to be a stigma or hassle cost. This cost does not apply to individuals who are eligible for regular retirement benefits; for age-eligible individuals, applications are always successful.

At 62, individuals become eligible for regular retirement (SSR) benefits, and SSD recipients have their benefits converted to the regular retirement-age SSR benefit.[8] Retirees, under SSD or SSR, may leave the rolls and re-enter the labour force in a non-career job to work in any year before they reach the maximum retirement age of 75.

2.1.2 Marrieds: The household-level optimization problem

The dynamic problem for married couples at the household level is conceptually similar to 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

[8]Elimination of SSD application after 61 is a simplification given that, empirically, we observe individuals continuing to apply for SSD benefits into their 60s. This is because of the regular benefit reduction before age 66, described in Section 3, does not apply to disability benefits claimed after age 66. In the model, however, we assume that individuals who retire between 62 and 65 receive the regular retirement age benefit with probability $p_\delta(x)$ and the age-adjusted benefit with probability $1 - p_\delta$.

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.

$$U_j(x_M) = \lambda \Upsilon_j^f(x_M) + (1-\lambda)\Upsilon_j^m(x_M) \tag{9}$$

The household's state vector is $x_M = \{a, E_m, E_f, w_m, w_f, \delta_m, \delta_f, \lambda\}$. Assets are accumulated at the household level, but social security benefit accumulation, wages and health are individual-level, $\imath-$ indexed states. The utility weight λ is chosen in the initial period of the couple household's existence. A typical assumption, dating to Manser and Brown (1980) is that λ is chosen through cooperative (Nash) bargaining, in which both household members' threat points are V_1^i 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 do not parameterize psychic gains from marriage, but choose a constant λ for all couples that generates realistic mean labour supplies for husband and wives.

The household budget set for marrieds, B_M is:

$$
\begin{aligned}
\xi_1 : \quad & (\tilde{T}_m - \chi(\delta_m) - l_m)w_m + (\tilde{T}_f - \chi(\delta_f) - l_f)w_f + (1+r)a + b(.) = \\
& c + a^{'} \\
\xi_2 a : \quad & l_m < \tilde{T}_m - \chi(\delta_m) \\
\xi_2 b : \quad & l_f < \tilde{T}_f - \chi(\delta_f) \\
\xi_3 : \quad & a^{'} \geq 0
\end{aligned} \tag{10}
$$

B_M is analogous to B_S for a two-member household, and consumption c is household-level, rather than individual-level, consumption.

The household planner's dynamic program and choice over lfs states of the household members is similar to the problem solved by singles, except that the planner chooses the optimal l and lfs for *both* members concurrently. There are sixteen possible lfs states for a married household, the product of the vectors $\Psi\Psi^{'}$ where $\Psi = [C_f, NC_f, A_f, R_f]$. Each member of the household faces the same transition possibilities between lfs states for given age and individual state vector x as his or her single counterpart defined above, though the *payoffs* within each lfs state depend not on x but on x_m.

For expositional purposes, the value function for the husband and wife of a couple both in career jobs is:

$$\begin{aligned}
\Upsilon_j^{m,CC}(x_M) &= u^m(c^*, l_m^*, l_f^*) \\
&+ \beta_m \varphi_m \mathtt{E}_j[\varphi_f \Upsilon_{j+1}^{m,*}(x_M^{*'}) + (1-\varphi_f)V^m(x_m^{*'})] \\
\Upsilon_j^{f,CC}(x_M) &= u^m(c^*, l_m^*, l_f^*) \\
&+ \beta_f \varphi_f \mathtt{E}_j[\varphi_m \mathtt{E}\Upsilon_{j+1}^{f,*}(x_M^{*'}) + (1-\varphi_m)V^f(x_f^{*'})]
\end{aligned} \tag{11}$$

Here, the $*$'d choice parameters are the values chosen by the planner given $\{j, x_m\}$. These choices include the next-period household lfs. So long as the marriage continues, both partners passively accept the planner's choices over household level $\tilde{c}$ and a, and individual l and lfs, which are chosen to maximize a weighted sum of the partners' individual Υs. Upon the death of one or the other household member, which occurs with probability $1 - \varphi_i$ for spouse i, the surviving spouse i reverts to the single optimizing state (with value function $V^i(x)$), and inherits the household's stock of assets (that is, the planner's choice of carry-forward assets). $\Upsilon_j^i, i \in \{m, f\}$ may, in theory, be greater or less than the payoffs to autonomous optimization in the single state, V, depending on whether the economic and psychic gains from marriage outweigh the loss of autonomy. Note that period utility u depends on household consumption, own leisure and spousal leisure. In the benchmark model and non-cooperative models, spousal leisure carries a weight of zero in u; in the complementarities-in-leisure model, however, own and spousal leisure enter period utility as complements.

2.1.3 Marrieds: non-cooperation

Introducing non-cooperation into the model can be best thought of as adding an additional step to the household optimization problem outlined above. The household social planner first chooses the optimal allocation and lfs of members according to the marriage contract as described above. Before accepting the planner's decisions, individual household members determine whether or not they can individually benefit from *deviating* from this allocation by choosing a different lfs—for example, quitting their career job or putting in for social security benefits. In a final stage of the problem within any age and realization of state space, individual i (with spouse $-i$) finds his optimal individual payoff from across the feasible choices of lfs. For a career worker, this final-stage problem is:

$$\hat{\Upsilon}j^i = \max \begin{Bmatrix} p(\delta)\hat{\Upsilon}_j^{i,R,lfs_{-i}|A} + (1-p(\delta))\hat{\Upsilon}_j^{i,A,lfs_{-i}|A}, \\ \hat{\Upsilon}_j^{i,C,lfs_{-i}|C}, \quad \hat{\Upsilon}_j^{i,NC,lfs_{-i}|NC} \end{Bmatrix}$$

For a non-career worker, the problem is:

$$\hat{\Upsilon}j^i = \max \left\{ p(\delta)\hat{\Upsilon}_j^{i,R,lfs_{-i}|A} + (1-p(\delta))\hat{\Upsilon}_j^{i,A,lfs_{-i}|A}, \quad \hat{\Upsilon}_j^{i,NC,lfs_{-i}|NC} \right\} \tag{12}$$

For a retiree under 62, the problem is:

$$\hat{\Upsilon}j^i = \max \left\{ \hat{\Upsilon}_j^{i,R,lfs_{-i}|R}, \ \hat{\Upsilon}_j^{i,NC,lfs_{-i}|NC} \right\} \tag{13}$$

In these equations, $\hat{\Upsilon}$ is the within-marriage value function for individuals who accept the household planner's choices within any given lfs but who deviate from the planner's choice over lfs—in other words, it is the analogue of Υ within the non-cooperative regime. Note that, in the non-cooperative world, $\hat{\Upsilon}$ is also the value function that enters the planner's problem U, which can lead to different household outcomes even when individuals end up following retirement paths that do not deviate from the planner's. In this regard, the problem presented in (12-13) 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, non-cooperative decisions over retirement are examples of *internal* threat points (see Lundberg and Pollak (1991)). In this case, 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.

If neither spouse is in a career job, and the household is under the maximum retirement age of 75, the spouses play the relevant game defined by 12-13 simultaneously. Therefore, there is not a unique outcome for all realizations of state space. We impose the following assumptions on non-cooperative retirement behaviour:

1. Assumption #1: Spouses never play mixed strategies.

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

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

4. Assumption #4: If a Nash equilibrium exists that makes both spouses worse off, the players revert to the planner's solution.

For a couple made up of non-career workers, or retirees under 75, the couple plays a simple 2x2 static Nash game, and assumptions #1-#4 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 lfs, and that the Nash equilibrium outcome that results differs from the planner's. Figure 1 shows four types of outcome that can arise: the first of which give rise to a potential deviation and the other three which 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 applies for benefits regardless of the wife's decision. The wife applies for benefits only if the husband also does so: the Nash equilibrium is $\{A, A\}$. Subfigure 1(b) shows a case in which both partners have dominant strategies to retire, but that playing these strategies leaves both members of the couple worse off than following the jointly optimal decision. Here, the couple finds itself in a "prisoner's dilemma" and cooperates to avoid this outcome (Assumption 4), landing on {NC,NC} with payoffs {B,b}. In figure 1(c), there is no pure strategy Nash equilibrium. In figure 1(d), both $\{A, A\}$ and $\{NC, NC\}$ are Nash equilibria.

Assumptions #1-#3, that prevent deviant outcomes in figures 1(c) and 1(d),are motivated by three ideas: (1) the single-shot Nash game played once per period is actually an approximation to a repeated static game that is played within a period until a stable outcome is reached; (1) retirement (or benefit application) is not a binding state for individuals under 75; at any point, an individual can opt back into non-career work; (3) spouses find uncertainty over their lfs status within a model period costly and therefore avoid it (a form of risk aversion not formally captured by preferences), converging instead on the focal point provided by the planner's solution.

Figure 1: Non-cooperative outcomes for non-career workers

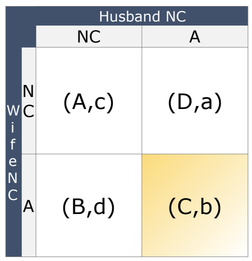

(a) Husband dominant retirement strategy

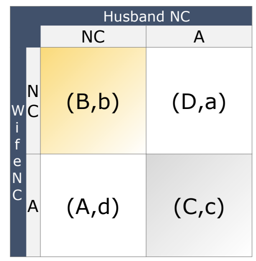

(b) NE: Husband and wife dominant retirement strategies

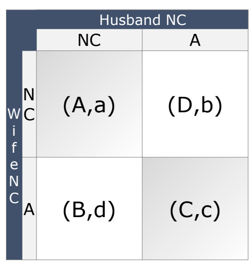

(c) Multiple NE: Revert to planner's solution

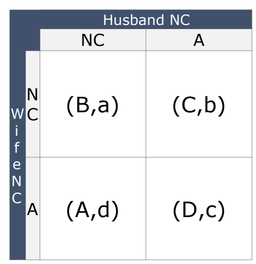

(d) No pure NE: Revert to planner's solution

The 2x2 case holds for pairs both of whom are either currently retired or in non-career ("bridge") jobs. The problem becomes slightly more complicated, with greater potential for deviant behaviour, if at least one spouse is in a career job, $lfs^i = C, i \in \{f, m\}$. This is because, unlike applying for benefits, retiring from a career job *is* a binding decision that cannot be undone. Figure 2 shows potential solutions to the game that arises when the wife is a career worker and the husband a bridge worker. The wife solves this game by backward induction, first resolving the 2x2 game that results if she retires, and then comparing this outcome to that which occurs if she does not retire. If the wife is better off playing the 2x2 game with her husband, she quits her career job. Otherwise she remains in her career job and accepts the outcome of the 3x2 Nash game. In the first panel of figure (2), no Nash equilibrium exists in the 3x2 game, but a Nash equilibrium ($\{NC, NC\}$) does exist in the 2x2 game induced by the wife's quitting her career job. However, the wife will not retire so long as the planner's solution is $\{C, NC\}$ since otherwise both spouses are worse off. (The wife would deviate if the planner's solution is $\{C, A\}$, that is, if the λ governing the cooperative decision were very small.) In the situation represented in figure 2(b), there are two Nash equilibria ($\{C, A\}$ and $\{A, NC\}$), but the wife can eliminate the $\{C, A\}$ equilibrium by quitting her career job. Since this raises her payoff from C to A, she quits. Finally, even if no pure Nash solution exists to the derived 2x2 game, the wife may still quit since the planner's solution to the 2x2 game might still improve her outcome. This possibility is depicted in figure 2(c). There is no Nash equilibrium in either the 3x2 or the 2x2 game induced by the wife's quit. However, if the planner chooses $\{C, NC\}$ as the cooperative solution to the 3x2 game and either $\{A, NC\}$ or $\{NC, NC\}$ ($\{A, NC\}$ as shown) as the cooperative solution to the 2x2 game, then the wife quits to receive the higher payoff associated with this reduced game. Again, the wife's decision to deviate clearly depends on λ: how much the planner's solution favours her. A symmetric game, and similar logic, governs the optimal retirement of a career husband with a bridge-working or currently retired wife.

The final game that arises in the model is played between a career husband and a career wife. The resulting 3x3 Nash game can be conceived as a pair of sequential games. At the

Figure 2: Non-cooperative outcomes for a career and a non-career worker

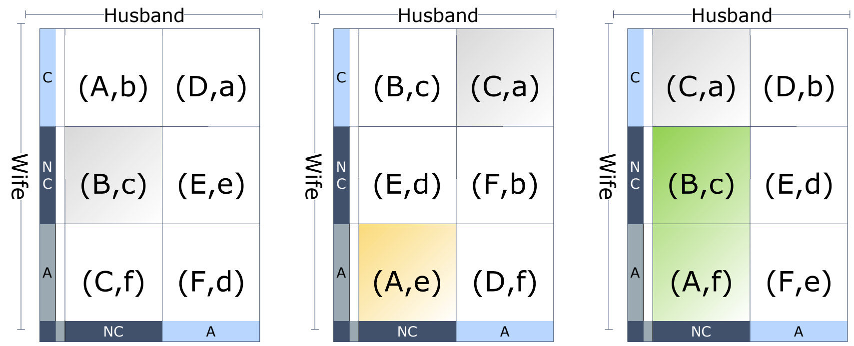

(a) Single NE (b) NE by wife quit 1 (c) NE by wife quit 2

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 2x3 game induced by quitting (the area under the thick dotted line passing to the south of the hatched $\{C, C\}$ square in figure 3) or remaining in $\{C, C\}$. The husband has a symmetric choice set when the wife plays NQ at the first stage. If the husband quits, the wife has a choice between the 3x2 game described in the previous paragraph (playing NQ) and the 2x2 game described at the beginning of this section (playing Q). If she is indifferent between these choices, that must mean that quitting is the outcome of the 3x2 game between a career wife and non-career/applicant husband given x_M. In this case, $\{Q, Q\}$ is a Nash solution to the first stage game and the couple locates and plays the 2x2 game defined in the square in the right bottom of figure 3. Thus, indifference over outcomes in the first stage game is resolved in favour of quitting. Otherwise, we assume the solution to the first-stage game follows the same assumptions that govern the final-stage 2x2 Nash game: if no pure strategy Nash equilibrium exists, or if multiple pure strategy

Nash equilibria exist (say $\{NQ, Q\}$ and $\{Q, NQ\}$), then the spouses revert to the planner's solution.

The different games, and deviations from the cooperative outcome, outlined above arise because neither member of the couple receives the full benefit of being in a 'good' (i.e. highly-productive) lfs state. The resulting externalities clearly lead to potential for non-cooperative deviations, though the importance of such deviations for explaining observed retirement and benefit-claiming behaviour is an empirical question that we address with the numerical model developed below.

Before turning to the numerical model, however, it is worth discussing the main idea about family interaction that underlies the non-cooperative model: that lfs-decisions, in contrast to saving or consumption decisions, are personal and can be taken unilaterally. A spouse may be able to observe her partner's general state of health, but may not be able to judge exactly, or effectively argue, how much that health affects his ability to work, leaving scope for deviation into state A. As well, an individual can quit a career job by resigning; since work decisions are made in the workplace and not in the home, the scope of spousal interference may be limited, at least in the short run. (Of course, the threat of divorce could play a role; however, divorce is relatively rare among couples approaching retirement age.) More generally, conflict over which spouse plays the primary earner role, and when it is acceptable to retire, is likely to be reflected in social norms. It turns out in the numerical analysis that it is *wives*' deviations from the lfs-specific marriage contract through early retirement and avoidance of career job is the main driver of the results, with husbands playing a fairly reactionary role. This is intuitively more consistent with the "war babies"

Figure 3: Non-cooperative outcomes for career workers

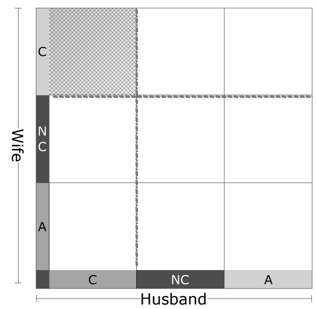

cohort of the HRS, which represents the first generation in which the majority of married women participated in the labour force, albeit mainly as second earners. However, we might expect a different result for later cohorts in which women face smaller earnings differentials compared to their husbands and and in which "career wives" are both more common and considered necessary for economic survival of many marriages.

3 The numerical model

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 two main types of risk: productivity (wage) risk and health or disability risk, both of which increase over the life cycle. As well, individuals face mortality risk that depends on health and age. Finally, individuals face variations in the amount of disposable time that can be devoted to labour and leisure, due to changing family and work obligations over the life cycle.

3.0.4 Life cycle effects and sources of risk

3.1 Disposable time endowments

Disposable time in the model is time that can be devoted to either labour or leisure. It is equal to an individual's temporal endowment of 168 hours a week minus required sleep time of 42 hours[9] and time lost to obligatory non-labour, non-leisure (nll) activities, such as personal care and housework, that are requirements of daily life. We calculate nll time from information available in the 1981-1997 panels of the PSID and the 2005 and 2006 American Time Use Surveys (ATUS) for single and married men and women explicitly, following the criterion used in Aguiar and Hurst (2007)'s Leisure Measure II to distinguish nll and leisure activities. Disposable time rises in early life as individuals' gain back the nll time devoted to education. It falls again, especially for women, during childbearing years and when younger children (under 6) are in the house. Finally, disposable time rises gradually through middle age as children leave, but declines slightly for married women through the retirement window 62-70, while continuing to rise slightly for singles and divorcees of both genders.

3.1.1 Career and non-career jobs and productivity shocks

Workers who are not currently retired may be in either a career- or a non-career job. In the data, we define career workers as those who report having worked at least 1600 hours in the previous year and also report being employed at the time of the survey. Conditional on lfs, wages evolve exogenously according to the following process:

[9] Six hours a night is a standard estimate from the psychology literature of the mean hours per night required for normal functionality. See, for instance, Taber and Hurley (2006).

$$
\begin{aligned}
\ln w_{it} = &\quad \ln w_t + \varphi^C X_{it} + \mu_{it}^C \qquad if \quad Z_{it}\varsigma + \nu_{it} \geq 0 \\
\ln w_{it} = &\quad \ln w_t + \varphi^{NC} X_{it} + \eta_{it}^{NC} \quad if \quad Z_{it}\varsigma + \nu_{it} < 0
\end{aligned}
\tag{14}
$$

where $\{\nu, \mu, \eta\}$ X, are distributed normally, conditional on X:

$$
\Omega = \begin{bmatrix} 1, & . & . \\ \sigma_{C1}^2 & \sigma_C^2 & . \\ \sigma_{NC1}^2 & . & \sigma_{NC}^2 \end{bmatrix}
\tag{15}
$$

Levine (1983) shows that the MLE estimator is consistent when the residuals are serially correlated and/or heteroskedastic. We assume that both serial correlation and heteroskedasticity in age are present in the error process for wages:

$$
\begin{array}{ll}
\mu_{it} = \mu_{it}^{\rho} + \mu_{it}^{me} & \mu_{it} = \eta_{it}^{\rho} + \eta_{it}^{me} \\
\mu_{it}^{\rho} = \rho^C \mu_{it-1}^{\rho} + \upsilon_{it}^C & \eta_{it}^{\rho} = \rho^{NC} \eta_{it-1}^{\rho} + \upsilon_{it} \\
\upsilon_{it} \sim N(0, \sigma_{\rho,C}^2(j)) & \upsilon_{it}^{NC} \sim N(0, \sigma_{\rho,NC}^2(j)) \\
\mu_{it}^{me} \sim N(0, \sigma_{me,C}^2) & \mu_{it}^{me} \sim N(0, \sigma_{me,NC}^2)
\end{array}
$$

and

$$X \equiv \{age, age^2, age^3, educ, race, health\}$$

$$Z_{PSID} \equiv \{children, marital\ status, religion\}$$

$$Z_{HRS} \equiv \{Z_{PSID},\ pension\ coverage, employer\ health\ insurance\}$$

In the system described above, w_t is the market price of labour in year t (a constant term shifted by a year dummy) and X contains a cubic in age, years of education, indicators for race and (for workers observed in years 1984 and later) an index for self-assessed health. The vector of first-stage regressors, Z, depends on whether the individual is drawn from the PSID or the HRS. Z_{PSID} contains all the variables in X, plus an indicator for marital status, number of children in the household, and a dummy for Catholic. Z_{HRS} contains these variables plus an indicator for being covered by a employer pension, for receiving health insurance through work, through the spouse's work, and whether own employer health insurance also covers the spouse. Following Wooldridge (1995), we also include in X and Z time averages of the covariates in X for each individual to control for a possible fixed effect affecting both wages and selection into career jobs.

To run the regressions, we first pool individuals from the 1976-1993 waves of the PSID with the first eight waves of the HRS (1992-2006), and keep all individuals born between 1932 and 1942. This allows me to follow this cohort back in time to a minimum age of 34. Since the PSID and HRS represent different populations (the PSID represents the whole US population circa 1969), we use data from the overlapping year, 1992, to estimate the likelihood of any individual in the PSID sample also appearing in the HRS, and use the resulting propensity scores to weight the PSID observations in the switching regression. At the third stage, we estimate the process for wage shocks, comprised of an autocorrelated and a white noise component, the latter which we assume to be pure measurement error. We use a piece of the generalized non-linear estimation process described by Baker and Solon (2003), in allowing the variance of the autocorrelated part of wages to be heteroskedastic in age. In the final stage, we use the residuals to compute the variance and first eight autocorrelations in the error. We then take the sample average of each moment by age and stack them. Finally, we use nonlinear least squares to compute the population model parameters that best match the sample moments, according to (16):

$$
\begin{aligned}
& ar_i = \sigma^2_{fx} + I_{ar0}\sigma^2_{me} + \sum_{k=0}^{8} \rho^{2k+i}(\gamma_0 + \dot{\gamma_1}(j-k) + \dot{\gamma_2}(j-k)^2) \\
& i = \{1, ..., 9\} \qquad (16)
\end{aligned}
$$

where I_{ar0} indicates that measurement error only shows up in the variances. Because the Wooldridge panel estimator controls for fixed effects by including sample averages of the independent variables at both stages, there is, by assumption, no fixed-effects component in the third-stage regression. The resulting second- and third-stage results that parameterize the model are given in the bottom panel of table 1. As expected, non-career work provides lower compensation than career work, and the relative wage loss becomes larger with age. For men, especially, predicted non-career wages decline over the working life while predicted career wages rise to around age 45 and then remain fairly constant. The coefficients on $\dot{\gamma_1}$ and $\dot{\gamma_2}$ suggest that, for men and women, the variance of shocks first decreases and then increases over the working life, similar to the pattern found by Baker and Solon (2003).[10][11]

Table 1: Wage equation parameters: Career and non-career workers

	Career workers		Non-career workers	
	Women	Men	Women	Men
age	.340	.587	-.678	-1.631
	(.780)	(.092)	(.0163)	(.808)
age^2	-.0066	-.0118	.0149	.0328
	(.0159)	(1.033)	(.808)	(.021)
age^3	.00004	.00007	-.00011	-.00022
	(.00011)	(.00004)	(.00011)	(.002)
$cons$	-4.30	-7.32	10.9	29.0
	(12.7)	(10.2)	(13.3)	(16.9)
σ_{me}	.030	.028	.134	.052
	(.049)	(.092)	(.016)	(.047)
ρ	.923	.891	.984	.803
	(.042)	(1.033)	(.016)	(.026)
$\dot{\gamma}_0$	.532	.737	.057	1.849
	(.158)	(.00004)	(.0010)	(.448)
$\dot{\gamma}_1$	-.023	-.031	-.0016	-.076
	(.0065)	(1.033)	(.0010)	(.0174)
$\dot{\gamma}_2$	.000297	.000337	.000015	.000799
	(.000068)	(.000040)	(.00015)	(.00017)
N	5236	5593		
$Wald(9)$	451.3	268.0		

3.1.2 Health/disability shocks

Health/disability shocks (δ) are similar to productivity shocks in that they increase the leisure cost of work shocks—are "time stealing" in the sense imagined by Batavia and Beaulaurier (2001). We estimate health transition matrices based on self-assessed health and subjective

[10]Note that, for non-career workers, we do not always observe the wage, as some non-career workers are out of the labour force. Thus, the non-career parameters are likely to be somewhat inconsistent and require additional calibration in the model. In the current results, the parameters reported below are treated as the true parameters in the model.

[11]The data also allow me to estimate p_C, the probability of transitioning from a non-career to a career job. This can potentially vary by age and sex. In this paper, however, we set $p_C = 15\%$, which is roughly the unconditional empirical probability of transitioning into a career job for workers over 40.

reports of work-limiting disability.[12] Individuals are disabled if they report that a 'physical or nervous condition' limits their ability to perform work 'sometimes', 'often' or if they can do no work due to the condition; or if they report their current health level as 'poor'. we estimate Markov transitions between disability states by multivariate OLS of current disability status on lagged disability status. There are three states: not disabled, currently disabled and recently disabled (that is, disabled in th previous period), the latter which helps to capture the fact that the likelihood of being disabled decreases for those who have been disabled before. For marrieds, we estimate a nine-state transition matrix for the couple, which allows for arbitrary correlation between the husband and wife's physical conditions over time. Transition matrices are estimated for each household age j using all relevant observations in an eleven year interval with j at the center. The size of the δ shock associated with each state is assumed to be constant over the life cycle and is calibrated to match observed labour supplies in each state, with δ normalized to 1 in the non-disabled state.

3.1.3 Mortality risk

Following Rivas (2007), we calculate disabled and non-disabled mortality profiles from PSID data and the National Vital Statistics Report (2003) using Bayes Rule. The healthy and

[12] The PSID contains only subjective global measures of disability and health. These measures have come under scrutiny in the empirical disability literature. (See, for instance, Coyle (2004) for a brief overview.) The case against this type of measure is that it may be susceptible to justification bias if workers with high valuation of leisure or exogenously low returns to work are more likely to report themselves disabled, which would bias estimates of disability effects on labour market activity and outcomes away from zero. Working in the opposite direction, subjective global assessments may be measured with error since the researcher captures cross-sectional variation in the interpretation of the question as well as in health states. Using Canadian data, Campolieti (2002) finds that that second effect dominates when using self-reported disability status to predict the participation decision of older men, leading to an underestimate of the role of disability. Au et al. (2005) find a similar result for the effect of self-assessed health status on participation, also using Canadian data.

Figure 4: Health-conditional survival indexes for males and females

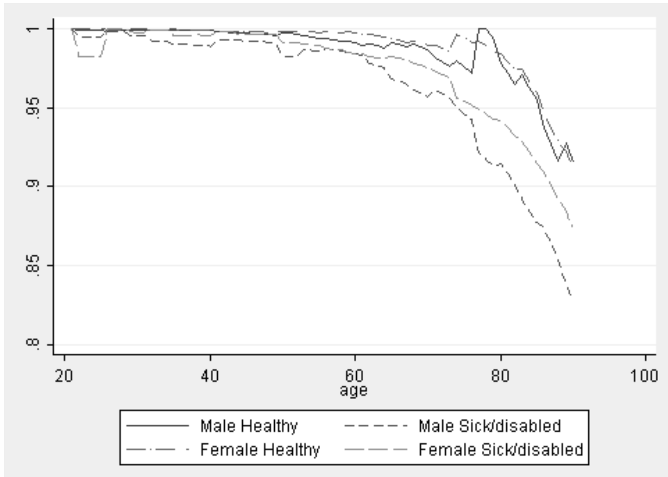

disabled profiles can be written as:

$$
\begin{aligned}
\varphi^{j}_{\delta=1,i} &= Pr(Surv|\delta=1) = \frac{\widehat{Pr}(Surv)\widehat{Pr}(\delta=1|Surv)}{Pr(\delta=1)} \\
\varphi^{j}_{\delta>1,i} &= Pr(Surv|\delta>1) = \frac{\widehat{Pr}(Surv)\widehat{Pr}(\delta>1|Surv)}{Pr(\delta>1)}
\end{aligned} \tag{17}
$$

where hats indicate that the information comes from the PSID micro data. Figure 4 plots the relative survival probabilities of healthy and disabled males and females using the subjective PSID measures of disability defined below. The figure, and corresponding regressions (not reported) show that (1) the importance of subjective good health increases with age and (2) that the effects of subjective good health on survival probabilities are significantly more important for men than for women at most ages.

3.1.4 Preferences

We adopt Cobb-Douglas preferences in period-utility. Evidence from the vast literature on household consumption and labour supply has for the most part found that consumption and leisure are best modeled as complementary in utility (see, e.g. Blundell and Walker (1986)). In the standard and non-cooperative models, utility is given by:

$$u_i(c_i, l_i) = \frac{[\psi_C(ms)(\frac{c_i}{\tilde{\eta}_{j,ms}})^{\gamma_i}(l_i)^{1-\gamma_i}]^{1-\omega_i}}{1-\omega_i} \tag{18}$$

For the complementaries-in-leisure model, period utility is given by

$$u(c_i, l_i) = \frac{[(\frac{c_i}{\tilde{n}_{j,ms}})^{\gamma_i}(l_i + I_{rr}\psi_L(ms)\min[l_i, l_{-i}])^{1-\gamma_i}]^{1-\omega_i}}{1-\omega_i} \tag{19}$$

In the above expressions, γ measures the consumption share in utility, ω is the coefficient of intertemporal risk aversion and $\tilde{\eta}_{j,i,ms}$ is the age-, gender-, marital-status-specific weighted measure of dependents supported by the individual. Fertility in the model is exogenous. For dependency costs of children, we use the modified OECD scale, assigning a weight of 0.3 to children under 17 and 0.5 to children over 16. The average number and timing of children are calculated from the 1994-2005 PSID sample by head's gender and age. Half of single women have one child, at age 19, whom they support for 24 years; the rest have no children. Single men do not have any children for tax and transfer or purposes but share the consumption-costs of the children of single women proportionately to their income (i.e this is a world of perfect paternity tests), and also devote the associated nll time for fathers. Married couples also have two children at 23 and 34 and support them for 23 and

19 years respectively. Married couples consume household resources jointly, with parameter $\psi_C(ms = M) \geq\leq .5$ governing the relative individual-level consumption gain from joint consumption (with $\psi_C > .5$ indicating the presence of a household-level public good). In the complementarities model, ψ_L captures the additional utility gain the individual receives from leisure enjoyed with her spouse with neither is working ($I_{RR} = 1$); in theory, both of these parameters are likely to vary with the quality of the marriage.

The main drawback of Cobb-Douglas preferences is that they do not allow us to disentangle relative risk aversion in consumption from risk aversion in leisure (or intertemporal elasticity of labour supply). There is some evidence from the micro literature that women are more risk averse in consumption than men (Mazzocco (2004)), and also that they have higher Frisch elasticities of labour supply (Blundell and MacCurdy (1999)). These facts are difficult to reconcile with Cobb-Douglas preferences because, holding the ratio of labour to leisure constant, the Frisch elasticity is decreasing in ω.[13] However, Mroz (1987) argues that *uncompensated* elasticities of female labour supply are similar to those for men in a well-specified model. For my benchmark models, we set γ to capture to gender-specific prime-age (25-59) labour supplies observed in U.S. data. For men, the resulting γ is around .45 in both the cooperative and non-cooperative models. By setting a value for ω of 2.5 for men, we arrive at an average Frisch elasticity of about .6, which is just at the high end of the Frisch elasticities estimated in the empirical life cycle literature on male labour supply under the

[13]Frisch elasticity of labour under Cobb-Douglas preferences is:

$$\frac{1 - \alpha(1-\omega)}{\omega}\frac{l}{n}.$$

The derivative with respect to ω is $(\frac{\gamma}{\omega} - \frac{1-\gamma(1-\omega)}{\omega^2})\frac{l}{n}$ which is decreasing so long as $\gamma < 1$.

assumption of separable preferences[14] (while $\omega - 2.5$ is just at the high end of estimated coefficients of intertemporal risk aversion for males). Women have a slightly smaller calibrated γ of around .35, corresponding to their lower labour supplies. Setting ω_f to 3.0 gives me roughly the same Frisch elasticity of about .6 for women as for men.

3.1.5 Calibration

The program described above gives me a total of fourteen parameters to calibrate (fifteen in the complementarities model). They are listed, along with the corresponding targets, in table 2. All calibration is done using PSID data. We calibrate the cooperative and partial/fully non-cooperative model separately and report the resulting parameter values for each. However, many of the results, especially utility and welfare results, are always given in reference to the calibration of the cooperative model.

Table 2 gives an indication of the size of disability costs, $\chi(\delta)$, the values of which have a straightforward interpretation as weekly hours lost to disability. The resulting costs are very large, in the order of 46 hours for men in the worst health state and 33 hours for women in the worst health state. Moving down the table, the gender-specific discount rates are set to reflect the fact that single households headed by women hold proportionally more wealth as a ratio of income than single households headed by men. Acceptance rates into SSD are set to match the share of individuals between 30 and 65 in each health state who begin receiving benefits in a given year. The rejection cost c_A controls the gross number of applicants. We target a number that, combined with the ps, leads to an overall rejection rate[15]—of around

[14] Domeij and Floden (2005) estimate a Frisch elasticity for men in the absence of borrowing constraints of .5.

[15] Because individuals in the simulated economy can continuously apply for benefits, while appeals in the

50%, which Gruber (2000) and Bound et al. (2002) separately cite as the prevailing post-appeal rejection rate of SSD applicants. Finally, planner's utility weight for the wife λ and the consumption gain to marriage $\psi_C(M)$ are set to generate an accurate joint labour supply of (healthy) married houseeholds conditional on married males supplying the correct number of hours. In the complementarities model, the leisure boost to joint retirement ψ_L is set to replicate the correct share of individuals who are working zero hours by age 64. Intuitively, the higher is the leisure gain to joint retirement, the fewer individuals should be working into their late 60s.

U.S. are actually limited to three after initial application, the rejection rate in the model refers to all "spells" of applications—i.e. groups of one or more consecutive application periods. In practice, less than 1% of simulated SSD application spells last more than four consecutive periods.

Table 2: Calibrated parameters and targets

Parameter	Benchmark	Non-coop	Compl.	Target	Target value
γ_f	.45	.41	.44	Unmarried healthy female labour supply 25-60	32.0 hrs
γ_m	.47	.47	.47	Unmarried healthy male labour supply 25-60	39.6 hrs
ω_f	3.0	3.0	3.0	Assigned	n/a
ω_m	2.5	2.5	2.5	Assigned	n/a
β_f	.985	.982	.984	HH wealth-income ratio for homes w female member	2.50
β_m	.983	.977	.980	HH wealth-income ratio for homes w male member	2.35
$\chi(\delta_{2f})$	15	22	14	Unmarried female labour supply $\|\delta_f = 2$	22.5 hrs
$\chi(\delta_{3f})$	45	40	45	Unmarried female labour supply $\|\delta_f = 3$	9.6 hrs
$\chi(\delta_{2m})$	12	15	11	Unmarried male labour supply $\|\delta_m = 2$	26.4 hrs
$\chi(\delta_{3m})$	53	50	52	Unmarried male labour supply $\|\delta_m = 3$	9.0 hrs
$p(\delta_1)$	.16	.18	.19	Share of new SSD recipients among pop with $\delta = 1$	.2%
$p(\delta_2)$	.27	.21	.24	Share of new SSD recipients among pop with $\delta = 2$	1.5%
$p(\delta_3)$	.36	.63	.33	Share of new SSD recipients among pop with $\delta = 3$	5.7%
c_A	1.01	1.02		Total share of SSD application ending in success	50%
λ	.57	.55	.58	Married male labour supply	44.8 hrs
$\psi_C(M)$	.71	.72	.71	Joint labour supply of married households	70.7 hrs
$\psi_L(M)$	n/a	n/a	.254	Share of individuals not working by age 64	.81

3.2 The policy environment

3.2.1 Social security

Under the rules of the U.S. Social Security system, individuals "save" for retirement by effectively paying a share of their per-period wage earnings to the Social Security Administration. Payments into the system by current workers finance the benefits of the current generation of retirees. The contributions of workers into the system function as IOUs from the government. When the current generation retires, their benefits will in turn be financed by the next generation of workers. Social Security is primarily a retirement program. However, it offers insurance to workers who become incapacitated in the form of disability benefits and to the spouses of workers who die prematurely in the form of spousal benefits. Social Security currently accounts for about 37% of total federal expenditure in the U.S.

Here, we outline the major features of Social Security that are captured in the model:

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

$$
\begin{aligned}
PIA \quad &= \quad 0.9 \min\left(0.2\overline{w}, E\right) \\
&\quad +0.32 \max\left(\min\left(E - 0.2\overline{w}, 1.3\overline{w} - 0.2\overline{w}\right), 0\right) \\
&\quad +0.15 \max\left(E - 1.3\overline{w}, 0\right) \\
a_t < \overline{a}^{mt}: &\quad b_{\cdot} = \max(PIA, \underline{b}) \\
a_t > \overline{a}^{mt}: &\quad b_{\cdot} = PIA
\end{aligned}
\tag{20}
$$

where $\overline{w}$ is the average per-capita wage earnings in the economy, E is shorthand for the AIME, and $\underline{b}$ is a floor benefit, equal to \$151 weekly for an individual (\$222 for a couple). This benefit is administered by the Social Security Administration as a separate program (Supplemental Security Income: SSI) for workers and non-workers who do not qualify for regular Social Security benefits or whose accrued earnings are too low.[16] Since SSI is means-tested as opposed to work-tested, $\overline{a}^{mt}$, equal to \$4000 for a single household head or \$6000 for a couple, is the maximum level of asset holding for which an individual or household can be guaranteed the floor benefit.[17]

State space limitations prevent an accurate accounting of which years should be included

[16]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.

[17]These figures are adjusted to exclude the average share of assets in housing wealth, which is not included in the determination of SSI resources. See 3.2.2 for further discussion.

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

$$
\begin{aligned}
j < 60: \qquad E_t &= \frac{E_{t-1}(j_t - j_1) + w_t n_t}{j_t - j_2} \\
j \geq 60: \qquad E_t &= \frac{E_{t-1}(j_t - j_1) + w_t n_t}{j_t - j_2}
\end{aligned}
\tag{21}
$$

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

Disability benefit eligibility and receipt. In order to receive disability benefits under SSD, applicants must pass an eligibility test which insures that the disability is "total" in the sense that it precludes all "substantial gainful activity" and is expected to last at least 12 months. Slightly more lenient eligibility rules apply to individuals over 55. If rejected for benefits, an applicant may appeal up to four times, to four different levels of SSD adjudicators, a process that can take several years.[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%.

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

[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.

In the model, these provisions are captured by acceptance rates that vary across age (pre- and post-55) and the disability/health index of the applicant. We use the shares of new SSD recipients in each self-assessed health/disability state in the PSID, and the share of benefit recipients older than 55 among the recipient population, to calibrate these acceptance rates. Individuals older than 35 may apply for benefits in any period they want, from any d state, but at a cost of forgoing a year of work in order to demonstrate incapacity. Since the decision process for an application takes five months in the real world, successful applicants in the simulation receive only 60% of their benefit in the year they are accepted.

Early and delayed retirement. Under the SSA's rules, individuals may retire at any age between 62 and 70, subject to an adjustment of benefit size that roughly equates the expected discounted stream of benefits across retirement ages. For early retirees, benefits are reduced by 5/9 of a percent for every month before the full retirement age of 66. As of 2008, the factor of adjustment for later retirees is 8% of the PIA per year. As well, individuals can continue to replace lower-earning years in the calculation of their AIME until they formally retire. Both of these effects - adjustment and continued accumulation - are captured in the model.

Survivor and spousal benefits. Under current SSA rules, surviving spouses whose own Social Security benefit is smaller than their partner's receive their partner's benefit until their own death. As well, secondary earners during the working life of a married household are entitled to the greater of either their own accumulated benefit or one half of their partner's benefit upon retirement. Both of these features of the program are captured in the simulated economy.[20]

Post-retirement work and "unretirement". Finally, we take seriously the argument made by Cahill et al. (2005) and others that retirement is not a once-for-all decision that spells the end of work. SSR recipients in the model can continue to work after taking up Social Security benefits. In keeping with current SSA policy, workers 66 and older suffer no reduction in benefits regardless of how much they work, though the taxation of benefits is likely to change (see Section 3.2.2). Retirees aged 62-65 are subject to a reduction in their present benefits and an increase (due to additional accumulation) in future benefits. SSD recipients younger than 62 cannot work while receiving benefits,[21] but may leave SSD at any age prior to 62. The SSA has run a series of programs designed to help SSD beneficiaries transition back to work, allowing recipients to maintain their benefits during a trial work period of up to nine months. To capture this effect, individuals may receive a three-quarters benefit during their first year of "unretirement". After returning to work, individuals retain their pre-retirement accumulation and continue to accumulate benefits like normal, healthy workers. Relatively few American SSD recipients - on the order of 12% in 2004 - leave the rolls and return to work (Autor and Duggan (2006)). This figure is not targeted in the model, providing a test of the model's ability to adequate capture the interaction of Americans with their full Social Security system.

[20]Under the rules for spousal benefits, spouses are not entitled to receive a half share in increases in the main earner benefit due to delayed retirement by the couple. Additional rules may provide benefits for the pre-majority age or disabled children of deceased parents. These rules are ignored in the simulated economy.

[21]Under the SSA's rules, SSI and SSD beneficiaries may not earn more than an amount indicating "minimal gainful activity", equal to $940 USD in 2008. For nearly all workers in the model, this amount is less than half what would be earned at the lowest positive labour grid point.

3.2.2 Non-Social Security policy

The model also incorporates 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.

Food stamps. Besides Social Security, the Food Stamp program is the most significant U.S. federal transfer and the only universal means-tested program. In 2006, food stamps were paid out to an average of 10.3% of households during any month at a total benefit cost of $31 billion USD. Eligibility for food stamps requires that applicants meet a gross income, a net income and an asset test. We model only the latter explcitly. The asset test varies across type of household. Households containing at least one member over 60 or who is disabled (which we take to mean receiving SSD benefits) may hold $700 in (weekly-adjusted) current liquid assets while younger and healthy households may hold $470. SSI recipients do not face an asset test for food stamps. Since some assets—in particular, housing—are exempt from the means test, we adjust the assets subject to the test in the model by calculating the average share of wealth in housing by age, gender, marital status and total stock of wealth using the 1999 and 2001 PSID wealth files.[22]

Conditional on passing the asset test, a formula gives food stamps as a maximum monthly

[22]The calculation is performed only on household with non-negative non-housing wealth. The resulting mean share of housing wealth is 61.5%, which is slightly higher than a similar estimate by Wolff (2007), who finds that Americans hold about 50% of wealth in housing. However, he includes pension wealth in his calculation while we exclude it. For simplicity, we also assume that all defined-contribution retirement savings is subject to the means test, though fact 401(k) wealth is ineligible for food stamps means-testing under federal law. Four states currently also exclude wealth held in IRAs.

allotment minus 30% of "net income", where net income under the USDA food stamp rules is equal to gross income minus a standard deduction of about \$30 and a further deduction of 20% of current earned income.[23] Actual foodstamp payments are then equal to the maximum of zero and this number. The maximum allotments vary by household size. To approximate how the program impacts various types of household and individual, we use the number of children by age estimated from the PSID and used to derive equivalence scales in 3.1.4 above. Two final simplifying assumptions about food stamps follow Low et al. (2007). First, unlike other public transfers and private transfers between households, we assume that access is universal to qualifying households; that is, the takeup rate is 100%. Second, we assume that food stamps are interchangable with cash for recipient households.

Earned income tax credit. The EIC (formerly EITC) is a wage supplement to low- and medium-income workers that is paid annually as a refundable tax credit. Similarly to food stamps, the income limits for which a refund can be attained vary by marital status and number of children and on having income from investments falling below a threshold of (weekly adjusted) \$700. The maximum eligible incomes and associated refunds are calculated from tables given in the 2007 edition of the IRA's Publication 596 and reach a maximum of \$800 in weekly earnings for married couples with two children. The maximum benefit in the model is about \$94, also for a married couple with two children earning between \$250 and \$350 a week.

Medicare, Health Insurance and Private Pensions. Disabled workers ($\delta \in \{2, 3\}$) are subject to out-of-pocket medical costs of \$30 and \$101 respectively, calculated from the

[23]Other small deductions may apply that we do not consider.

2006 HRS as differentials from average out-of-pocket expenses for individuals in excellent or very good health. These costs are reduced to \$24 and \$73 for individuals in households in which at least one spouse is in a career job, since these individuals are covered by employer-based health insurance. They are also reduced for individuals receiving SSDI benefits. Starting from the cost level of a career worker at age 64, medical expenses for individuals 65 and over increases linearly with age at a rate of .61 per year for healthy, .68 per year for recovering (poor health in a recent period) and .78 per year for currently unhealthy (poor health) individuals. Note that these costs understate the risk associated with health status especially for younger workers. Unhealthy individuals under 65 who are in career jobs pay only about 25% less out of pocket than those in non-career jobs, but the variance of their expenditure is cut by over a third, from \$211 to \$139.

Lastly, individuals (only single workers and married males for simplicity) in career jobs are eligible for pensions. The pensions themselves are subsumed into assets; however, several studies have demonstrated that defined benefit pensions create incentives for early retirement and penalize delayed retirement. This is due to several features of the accrual formula (see Stock and Wise (1990)); a simple reduced-form way to capture these penalties is by assuming that pension payouts obey the formula:

$$pn = k w_r y_s$$

where pn is the current pension annuity entitlement, k is a constant replacement rate, w_r is the wage in the last year worked, and y_s is a measure of years of service in the career job. In general, since wages begin to decline around age 60 for career workers, the rate of pension accumulation also begins to fall. The decline in accrual from age 60 to age 75 acts like a reduction in the interest rate faced by the household with a career worker since the pension is accumulating slower than if the worker could lock in his pension entitlement in the peak accrual year and receive compensation for the remainder of his career entirely in the form of spot wages which could then be invested. Based on the estimates in table 1, and assuming that retirement with replacement rate k is not possible until 60, the pension-penalty-adjusted interest rate $\hat{r}_p$ falls from 97% of r to -25% of r for career males between 61 and 74, with the effect only becoming large after age 67. Non-career workers are not affected.

4 Results and discussion

We now turn to the results from the simulated economies described in Sections 2 and 3. Section 4.1 examines the respective ability of the benchmark, non-cooperative and complementarities models to replicate certain aspects of U.S. household behaviour, in particular the retirement and benefit claiming patterns of older married men and women. Section 4.2 examines the implications of the models for current retirement policy and for three basic modifications to this policy.

4.1 Retirement and program usage

4.1.1 Retirement patterns

The HRS provides two useful measures of retirement transition behaviour: the age at which individuals first claim Social Security benefits, and the year when they first self-assess as "fully retired"—that is, no longer working. In figures 5 - 8 we report the simulated patterns of retirement using both measures for married men and women, and plot these against the same patterns estimated from the 1998-2006 waves of the Health and Retirement Survey corresponding to the 1932-1942 birth cohort. Because members of households are the same age in the model, we restrict the HRS comparison sample to households in which the head and wife are less than two years apart in age. In the simulations, the year of benefit claiming corresponds to the *first* year benefits are claimed, even if the individual returns to work in a bridge job in a later period. In the data, many bridge workers continue to receive their benefits; for computational reasons, this is not the case in the model: Individuals must be either working in bridge jobs or retired and claiming benefits, but can move between the two states costlessly, and receive benefit adjustment to compensate for the benefits forgone during the working year.

The figures report the main results of the paper. Unsurprisingly, the benchmark cooperative model performs the worst of all the models. Given that returns to delayed retirement are close to actuarially fair for the whose retiring population, some individuals—particularly those who are healthy and who expect to live, or expect their dependent spouse to live, a long time—should delay retirement. For men, a small "bump" in retirement at age 62 accounts for individuals who would have preferred to retire earlier due to poor health, low labour market returns or high accumulated wealth but have waited for the availability (and costlessness) of regular retirement benefits. Afterward, there is a relatively constant exit hazard from work into retirement as individuals reach their optimal retirement point over the retirement window. In all the models, a large share of women retire at 62 because their labour market returns are relatively low, having waited until after their children left the home to try to move into a career, and they expect to receive a spousal (and later a survivor) benefit. The "two peaked" retirement for women—at age 62 and age 69 when returns to delayed retirement end—is also a function of women's second earner status. A share of married women do find careers and even become equal or primary earners in their home. These women delay retirement to increase their own benefit because they do not plan to rely on a spousal benefit. This is evidenced by the fact that the average wage of wives retiring at 69 is $17.94, while the average wage of wives retiring at 62 is only $13.04.

Note that the large peak in retirement at 62 in the benchmark model is somewhat special to the cohort under consideration. Later cohorts of women, who are typically more active contributors to household income and receive higher wages over the working life, are likely to provide an even worse fit in the benchmark (and also the complementarities) model due to a greater incentive to postpone retirement to increase their benefits.

Next we consider the complementarities model. The model provides a good fit for women with the main peak in benefit claiming at 62, no peak in retirement at age 69, and also a reasonable departure hazard from the labour force in figure (7(c)). However, the model performs only slightly better for men than the benchmark in predicting the empirical pattern of benefit claiming, with the largest peak in retirement occurring at age 66 as pension penalties become large. This is because, while spouses have an incentive to *quit work* together, there is still a joint incentive for the main earner to delay claiming. The wife claims her benefit early and the couple consumes her benefits and their asset holding and enjoy each other's leisure. Later in the retirement window, the husband claims his appreciated benefit, which in many cases also serves as the spousal, and later the widow benefit. There is empirical evidence that delayed claiming of benefits matters (see Bound et al. (2002)) but the average delay is typically on the order of six months, not multiple years as predicted by the complementarities-in-leisure model.

Finally, the non-cooperative model gives overall the best fit for benefit claiming, replicating the peak in benefit claiming for males at 62 and almost no claiming for either spouse after age 66. The model overpredicts the rate of women's final exit from the labour force at age 62 relative to the complementarities model, but gives the best fit for male exit hazards. The reason for early joint benefit claiming rests on the fact that the much of the benefit of delayed claiming accrues to wives rather than husbands. In the non-cooperative model husbands have a heightened incentive to deviate from the planner's solution and retire right away. As discussed below, the effect is strengthened by the incentive of husbands to leave their career track before the retirement window which has the effect of forcing some of the earnings burden onto the wife and reduces the value of continued work to the husband (see below). Wives may respond to their husbands' early retirement by delaying their own retirement; however, they typically do not if their own retirement benefit is still likely to be smaller than the survivor benefit they receive upon the husband's death. The wife also cares about asset accumulation. However, when the husband retires, his action raises the marginal utility of consumption of the household, which reduces the optimal amount of saving. (Since non-cooperation can be predicted, there should not be a major offsetting wealth effect.) This also reduces the wife's incentive to keep working.

Table 3 provides a somewhat more formal measure of fit for the respective models, re-

Figure 5: Age of benefit take-up for husbands: Data and models

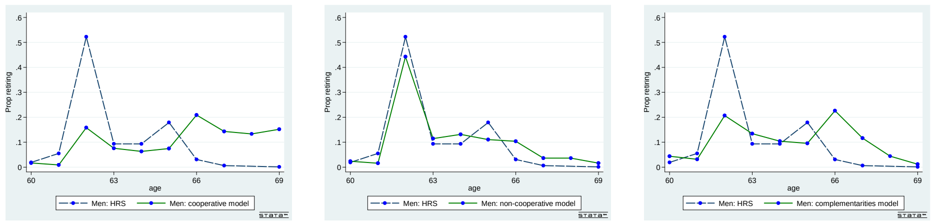

(a) Benchmark model (b) Non-coop model (c) Compl-in-leis model

Figure 6: Age of benefit take-up for wives: Data and models

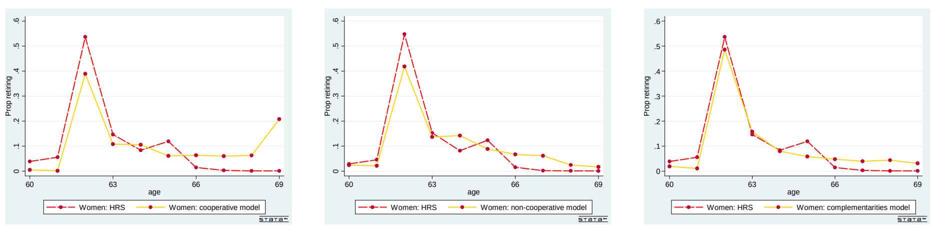

(a) Benchmark model (b) Non-coop model (c) Compl-in-leis model

Figure 7: Final year working for husbands: Data and models

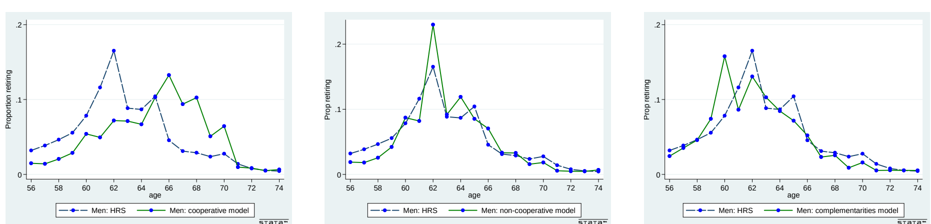

(a) Benchmark model (b) Non-coop model (c) Compl-in-leis model

Figure 8: Final year working for wives: Data and models

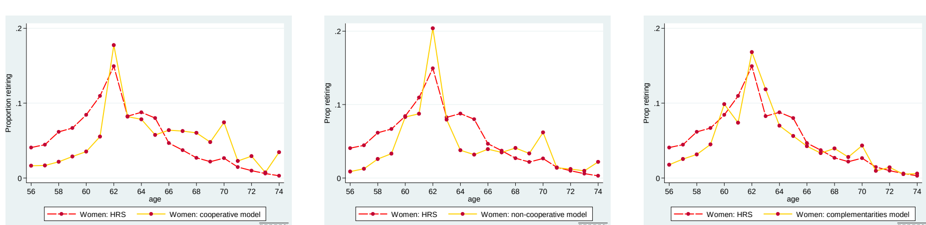

(a) Benchmark model (b) Non-coop model (c) Compl-in-leis model

Table 3: Differences in cumulative retirement across the models

	Benchmark	Non-coop	Compl-in-leis
Benefit claiming: Wives	1.029	.369	.815
Benefit claiming: Husbands	.722	.362	.634
Labour market exit: Wives	.610	.273	.288
Labour market exit: Husbands	.504	.328	.297

porting the average cumulative difference between the observed and predicted results in each of the graphs—i.e. the CDF of the HRS retirement hazards minus the CDF of the model retirement hazards. The non-cooperative model performs best for both men and women at predicting the cumulative distribution of benefit claiming. The complementarities and non-cooperative model perform nearly identically at predicting labour market exit. The complementarities model does approximate the actual distribution slightly better for women and the non-cooperative model approximates the empirical distribution slightly better for men.

Table 4 shows the performance of the models vis-a-vis the data in replicating the distribution of the proximity within which husbands and wives retire by claiming their Social Security benefits. Negative numbers in the table indicate that the husband retires ahead of the wife while positive numbers indicate that the wife retires first. We again find the counterintuitive result that couples in the non-cooperative model are *more* likely (in fact, too likely, relative to the data) to retire within two years of each other than in either of the other models. The share of husbands and wives retiring together is lowest in the cooperative model since husbands and wives have different returns to delayed claiming and continued work. The complementarities model underpredicts joint claiming, but does a better job or

Table 4: Husband retirement minus wife retirement in years

Difference in years	HRS 1998-2006	Benchmark model	Non-coop model	Compl-in-leis model
-6	0.2	13.9	1.0	5.1
-4	4.6	10.0	0.5	3.5
-2	3.7	4.8	0.4	2.1
0	72.8	39.1	80.1	51.5
$+2$	10.1	5.8	12.1	12.1
$+4$	8.4	12.0	8.1	20.5
$+6$	0.2	14.3	0.5	5.3

replicating the distribution of joint self-assessed retirement (not shown).

4.1.2 Disability benefits and program use

Another important comparison point for the models is their predictions for patterns of SSD use, specifically the models' predictions for the rate at which individuals *exit* the SSD program. Individuals can exit SSD by three means: dying; reaching regular retirement age and transitioning to SSR benefits; or returning to the labour force. Table 5 shows the distribution of these three alternatives for the cooperative and fully non-cooperative models (results from the benchmark model are similar to those from the complementarities model and are omitted for brevity.) The first column reports figures taken from SSA data, reported by Autor and Duggan (2006). The last two columns present corresponding results from the models. As can be easily seen, both the models greatly underpredict the number of deaths that occur among recipients.[24] However, the non-cooperative model generates better predictions regarding the likelihood of returning to work. In the complementarities model, a spouse on SSDI will return to the labour force if he recovers because he receives a relatively

[24]Note that mortality in the models depends on disability status, but, under current calibrations underpredicts the amount of severe sickness leading to terminal illness among the disabled population.

small utility gain from leisure while the wife is working; it is in the best interest of both spouses to work at the same time and accumulate assets. In the non-cooperative model, a spouse on SSDI has much less incentive to leave the roles and return to state NC because doing so reduces his leisure and increases his effective obligation to his spouse under the marriage contract.

Table 5: Transitions into and out of retirement: Data and models

Transitions out of SSDI			
	SSA data (Autor & Duggan [2006])	Non-coop	Compl-in-leis
Exit by death	42.0%	23.8%	23.0%
Exit by retirement	44.0%	61.1%	52.4%
Exit by return to work	12.0%	15.1%	24.6%

4.1.3 Two-step retirement and bridge jobs

Using the 1992-2002 waves of the HRS and focusing on the 1931 to 1941 cohort, Cahill et al. (2005) estimate the rates at which individuals transition into full retirement via non-career or "bridge" jobs. The authors argue that two-step retirement represents the most important recent shift in the behaviour of older workers, with an estimated 62-66% of men and 57-61% of women in the 1931-1941 cohort exiting the labour market via a bridge job. As Cahill et al. (2005) stress, however, the reasons for bridge job behaviour remain somewhat uncertain. Table 6 shows the retirement transitions generated in each of the three models among simulated workers who have held career jobs (that is, who have been in lfs state C) in the past. Again, the benchmark model seriously understates the degree of two-step retirement. Since bridge work is relatively poorly compensated, individuals have an incentive to remain in their career jobs as long as possible. Some bridge work is still observed, because individuals may be forced out of career jobs due to health problems that raise the marginal value of leisure. However, this actually produces a second counterfactual prediction of the benchmark model. In the data, bridge job behaviour is positively correlated with good health outcomes; yet poor health outcomes are the main driver of bridge behaviour in the benchmark model. In the simulation, bridge workers are 50% less likely to be in δ state 1 than regular workers.

The complementaries model understates the bridge behaviour of men and women. For women, the reason is obvious: the couple's incentive is for the wife to stay in her career job until age 62 then claim her Social Security benefits and retire. Husbands technically do transition out of career jobs and into non-career work before entering retirement. 68.5% of men use bridge work transitions under this simple definition. However, under the marriage contract, many of these (in fact, 42%) work zero hours while waiting to claim their benefit. Since these men would not be recognized as bridge workers in the data, removing them reduces the amount share of men who work bridge jobs to only 51.5%, as reported in the table.

Finally, the non-cooperative model generates the most realistic distribution of bridge work transitions. The reason is that both spouses have an incentive to avoid main earner obligations since this raises their required work load under the marriage contract and induces their spouse to reduce their own work further. Between the ages of 48 and 60, this dynamic leads to a fair amount of quitting. As well, in the non-cooperative model, bridge job behaviour is correlated with good health outcomes. Individuals in poor health who wish to quit do so mainly by applying for benefits under SSD. Healthier individuals prefer to avoid

Table 6: Share of individuals using two-step retirement

	HRS data (Cahill et al [2005])	Benchmark	Non-coop	Compl-in-leis
Men	66.0%	49.5%	69.7%	51.5%
Women	61.0%	33.7%	52.5%	41.7%

the hassle cost of applying and simply quit into lower-paying bridge work.

4.2 Policy implications

We conclude this section by examining the model's basic implications for policy. Table 7 reports results (ex-ante lifetime utilities in the first period of life) from the complementarities and non-cooperative models examined above, plus results from three standard Social Security policy experiments: (1) a 50% non-redistributive reduction in the size of the program; (2) elimination of the disability insurance program; (3) elimination of the spousal benefit. In each case, we reduce payroll taxes by the amount that finances the corresponding reduction in total Social Security payouts. The best outcome by gender in each model is denoted in italics, and the worst outcome in bold.

From table 7 we see that in partial equilibrium the Social Security system is welfare-improving for all married individuals in the complementarities model. By contrast, husbands in the non-cooperative model prefer a world in which the spousal benefit is eliminated. The explanation is that, in the complementarities model, the existence of a spousal benefit reduces the need for the wife to work through her 60s. She and the husband are able to effectively retire at age 62 and live off her regular benefit and their combined assets until it becomes optimal for him to claim benefits. The spousal benefit imposes costs on the the husband under the marriage contract—forcing him to trade off current for future consumption to benefit the wife—but the increase in wealth (in partial equilibrium) and facilitation of jointly enjoyed leisure is enough to offset these costs, at least given the calibrated λ. In the non-cooperative world, by contrast, the existence of the spousal benefit reduces the wife's attachment to the labour force, raising the likelihood of her applying for benefits under SSD when semi-healthy (in δ state 2) and encouraging her to quit work permanently as soon as she reaches age 62. The obligations of primary earner status are therefore left to the husband. For similar reasons, the wife supports SSD more strongly in the non-cooperative than the complementarities model.

Table 7: Lifetime utility under four policy experiments

Regime	Non-coop		Compl-in-leis	
	Wives	Husbands	Wives	Husbands
Value functions				
benchmark	*-2.030*	-2.394	*-1.955*	*-2.378*
50%	**-2.095**	**-2.436**	**-2.009**	**-2.422**
SSD elim	-2.038	-2.397	-1.956	-2.383
Spouse ben elim	-2.048	*-2.390*	-1.967	-2.379

Table entries are ex-ante expected life time utilities under each model and policy experiment. Bolded entries denote the worst payoffs across the policy regimes; italicized entries represent the best payoffs across policy regimes.

5 Conclusion

In this paper we have considered the retirement behaviour of married households under different assumptions about intrahousehold interaction and cooperation. In a non-cooperative model, spouses deviate from household-level optimization rules in deciding when to unilaterally retire or transition out of career work, because they cannot internalize the entire benefit from their own productivity (i.e. their career status). This mechanism produces early retirement and Social Security benefit claiming patterns that appear at odds with standard household intertemporal theory, but are consistent with empirical observation. A competing model for explaining intrahousehold retirement patterns is one in which couples benefit from complementarities in leisure when neither spouse is working. This model can also explain observed labour force departure hazards as well as the non-cooperative model, but it does a poor job of explaining the early spike in Social Security benefit claiming at age 62, the use of Social Security disability benefits, or transitions into retirement via bridge jobs. The non-cooperative model rationalizes these phenomena as results of individual desire to reduce the externalities associated to main earner status under the marriage contract. First, individuals who enter the disability roles rarely exit back into the labour force since doing so requires the individual to give up leisure for her partner's benefit, which is much more likely to occur in a cooperative complementarities model. Second, spouses in the non-cooperative model enter into lower-paying bridge jobs in their fifties in part to reduce main-earner burden after children leave the house.

The different models also generate different implications for policy. In the cooperative or 'complementarities' model, the current policy environment increases welfare relative to any scenario in which benefits are reduced or a piece of the Social Security program eliminated. The progressivity of Social Security and its availability to individuals in poor health states provides insurance against permanent or transitory bad luck. Intracouple obligations reduce the distortionary effects of early retirement incentives. In the non-cooperative retirement model, by contrast, these distortionary effects are larger, effectively enabling destructive strategic behaviour. As a result, the spousal benefit under current SSA rules is actually welfare-reducing for husbands, even in partial equilibrium. Note, however, that an across-the-board benefit cut is still welfare-reducing even in the non-cooperative world. This is because the gains from retirement do not depend very strongly on the amount of benefit received since the retiree gains leisure from his retirement. Premature retirement, however, impacts negatively on the partner of the retiring spouse. Reductions in Social Security benefits further increase the welfare loss due to a spouse's retirement, without inducing an a sufficient reduction in the *amount* of premature retirement to compensate.

Some issues remain open. A major limitation of our 'non-cooperative' approach is that, in its current form, it does not account for the possibility of endogenous marriage. Specifically, the threat of divorce could reduce the potential for non-cooperative behaviour if consumption and non-pecuniary gains from marriage are not too large. Moreover, the model has little to say about the retirement behaviour of singles. Single men and women in all the models, unsurprisingly, behave much like married couples in the standard cooperative model, choosing optimal retirement dates that occur later than those observed in the data. However, if marital status were determined endogenously, rather than assigned exogenously as in the current model, then single men and women would be drawn from parts of the wealth and health distributions that are most likely to retire early and we might derive more realistic predictions about retirement behaviour across different types of household. This possibility is left for future work.

References

Aguiar, M., Hurst, E., 2007. Measuring Trends in Leisure: The Allocation of Time over Five Decades. The Quarterly Journal of Economics 122 (3), 969–1006.

Au, D. W. H., Crossley, T. F., Schellhorn, M., 2005. The effect of health changes and long-term health on the work activity of older canadians. Health Economics 14 (10), 999–1018.

Autor, D., Duggan, M., 2003. The rise in the disability rolls and the decline in unemployment. Quarterly Journal of Economics 118 (1), 157–205.

Autor, D., Duggan, M., 2006. The growth in social security disability rolls: A fiscal crisis unfolding. NBER Working Paper 12436.

Baker, M., Solon, G., 2003. Earnings dynamics and inequality among canadian men: Evidence from longitudinal tax records. Journal of Labor Economics 21 (2), 289–321.

Batavia, A. I., Beaulaurier, R. L., 2001. The financial vulnerability of people with disabilities: Assessing poverty risk. Journal of Sociology and Social Welfare 38 (1), 139–162.

Blundell, R., MacCurdy, T., 1999. Labor supply: A review of alternative approaches. In: Ashenfelter, O., Card, D. (Eds.), Handbook of Labor Economics, Vol 3. North Holland Amsterdam.

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

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

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

Burkhauser, R., Butler, J., Gumus, G., 2004. Dynamic programming model estimates of social security disability insurance application timing. Journal of Applied Economics 19, 671–685.

Cahill, K., Giandrea, M., Quinn, J., 2005. Are traditional retirements a thing of the past. New evidence on retirement patterns and bridge jobs. US. Department of Labor Bureau of Labor Statistics Working Paper 384.

Campolieti, M., 2002. Disability and the labor force participation of older men in canada. Labour Economics 9 (3), 405–432.

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.

Coyle, C. C., 2004. Health shocks and couples' labor supply decisions, nBER Working Paper.

Coyle, C. C., Gruber, J., 2000. Social security incentives for retirement, nBER Working Paper 7651.

Domeij, D., Floden, M., 2005. Labour-supply elasticity and borrowing constraints: Why estimates are biased. Review of Economic Studies 9, 242–262.

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

Gustman, A. A., Steinmeier, T. L., 2004. Personal accounts and family retirement, nBER Working Paper No. W10305.

Gustman, A. A., Steinmeier, T. L., 2005. The social security early entitlement age in a structural model of retirement and wealth. Journal of Public Economics 89 (2), 441–463.

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.

Jiménez-Martín, S., Labeaga, J., Martínez Granado, M., 1999. Health status and retirement decisions for older European couples. IRISS Working Paper Series.

Kreider, B., 1999. Social security disability insurance: Applications, awards and lifetime income flows. Journal of Labour Economics 17 (4), 784–827.

Levine, D., 1983. A remark on serial correlation in maximum likelihood. Journal of Econometrics 23 (3), 337–342.

Low, H., Meghir, C., Pistaferri, L., 2007. Wage and employment risk over the lifecycle, working Paper: Institute for Fiscal Studies.

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

Maestas, N., 2001. Labor, love and leisure: complementarity and the timing of retirement by working couples. Berkeley: Xerox.

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

Mazzocco, M., 2004. Saving, risk sharing and preferences for risk. American Economic Review 94 (4), 1169–1182.

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

Mroz, T., 1987. Examining the sensitivity of estimates of women's labour supply to alternative econometric assumptions. Econometrica 10 (1), 160–191.

Rivas, M. C., 2007. Dynamic complementarities: A computational and empirical analysis of couples' retirement decisions, working paper, University College London.

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

Taber, K., Hurley, R., 2006. Functional Neuroanatomy of Sleep and Sleep Deprivation. Journal of Neuropsychiatry and Clinical Neurosciences 18 (1), 1.

Wolff, E. N., 2007. Recent trends in household wealth in the united states: Rising debt and the middle class squeeze, working Paper No. 502: The Levy Economics Institute at Bard College.

Wooldridge, J., 1995. Selection corrections for panel data models under conditional mean independence assumptions. Journal of Econometrics 68, 115–132.