

Dynamic Complementarities: A Computational and Empirical Analysis of Couples' Retirement Decisions

María Casanova Rivas^{*†}

11th November 2007

[Job market paper.]

Abstract

Evidence from different sources shows that a significant proportion of married individuals retire within less than a year from each other, independently of the age difference between them. The existing literature on couples' retirement choices suggest that this is due to non-separabilities in spouses' preferences and/or correlations in spouses' unobserved (dis)taste for work, although there is no agreement as to the relative importance of the two motives. Disentangling the roles of these incentives to joint retirement will have important implications regarding retirement modelling.

In this paper, I present the first structural, dynamic model of older couples' saving and participation decisions which allows for the presence of non-separabilities in spouses' preferences and correlation in spouses' unobserved taste for work, while controlling in a precise way for the financial incentives to retirement of individuals in the estimation sample.

I use a subsample of individuals from the Health and Retirement Study (HRS) to estimate the laws of motion of the exogenous variables. I estimate the parameters of the wage equation accounting for the presence of a fixed initial condition and controlling for selection, which is likely to affect the sample of older individuals if those with the most negative shocks to wages choose to retire. I estimate the process determining health care costs merging data from my sample with a sample of older individuals, in order to be able to parameterise the evolution of medical care costs in old age.

Finally, I present simulation results that use these estimates as inputs, and show how the model can account for the dispersion in couples' choices that results from the large heterogeneity in observed states. I simulate the effect of a reduction of the spousal benefit in the Social Security and show how, despite this being a provision mainly intended to cover dependent wives, the policy change would not have a strong effect on women's participation, instead leading to a noticeable increase in the participation rates of men close to retirement age.

Keywords: Joint retirement, life-cycle model, household decisions.

JEL classification: J22, J14, D1, C61

^{*}University College London and Institute for Fiscal Studies

[†]Comments welcome at maria.c@ifs.org.uk. I am especially grateful to Richard Blundell and Lars Nesheim for their guidance on this project. I thank Sule Alan, James Banks, Andrew Caplin, Monica Costa Dias, Tom Crossley, Mario Fiorini, Erich French, Giovanni Gallipoli, Alan Gustman, John Bailey Jones, Nicole Maestas, Costas Meghir, Nicola Pavoni and participants at the RAND "Health and Economic Determinants of Retirement" conference in July 2007 for their help, comments, and suggestions. Finally, especial thanks are due to Michael Nolte, David Weir and all the staff at the MiCDA Data Enclave in the Michigan Center on the Demography of Aging for their assistance during my visit.

[‡]This project was supported by funding from the Institute of Fiscal Studies and Grant Number 5 P01 AG022481-04 from the National Institute On Aging. Its contents are solely the responsibility of the author and do not necessarily represent the official views of these institutions.

1 Introduction

With the first baby-boomers reaching retirement age in 2010, a massive increase in the US elderly population will be taking place during the next decade. Even under the most optimistic assumptions regarding future birthrates and immigration, a sharp rise is projected in the share of GDP devoted to Social Security and Medicare¹. Different policies have been suggested in order to alleviate the budgetary burden, some of which, such as the progressive increase of normal retirement age up to 67 years of age, are already taking place. In this context, it is crucial that we understand how savings and employment decisions respond to changes in incentives during the years around retirement age. This will allow understanding and predicting the effects of policy changes and, more importantly, measuring the effects on well-being in old age.

There is a wealth of papers studying retirement choices of individuals². However, since the majority of people reaching retirement age in the US are married, it is essential to be aware of potential interactions between spouses' behaviour. Evidence from different sources shows that a significant number of spouses retire within less than a year from each other, independently of the age difference between them³. The study of this phenomenon, known as joint retirement, has drawn increasing attention in the past few years.

The reasons that may lead spouses to retire together can be broadly classified in four categories: correlation in spouses' (dis)taste for work; correlation in observable variables such as assets, wages, pension incentives, health status, etc.; correlation in time-varying shocks; and complementarities in the spouses' preference for leisure, so that one or both of them enjoys retirement more if their partner is retired as well. The existing literature suggests that most of the coordination operates through the spouses' preferences, either because of non-separabilities in spouses' leisure or correlated tastes for work, but there is no agreement as to the relative importance of the two.

In this paper, I present the first structural, dynamic model of older couples' saving and participation decisions which allows for the presence of non-separabilities in spouses' preference for leisure and correlation in spouses' unobserved taste for work, while controlling in a precise way for the financial incentives to retirement of individuals in the estimation sample. Estimation of the model parameters will allow to ascertain the relative importance of different incentives to joint retirement.

Understanding the causes of joint retirement is crucial, due to the different implications for retirement modelling that they have. This applies, for instance, to the two potential sources of correlation in spouses' preferences, namely non-separabilities and correlated tastes for work. If most of the correlation in spouses' decisions is due to their shared taste for work or leisure, modelling their decisions separately will not cause problems with estimation. In particular, an exogenous change to one spouse's incentives to retirement will not affect their partner's behaviour, and hence considering the two retirement equations separately should yield unbiased -though likely inefficient- estimates. On the other hand, the existence of complementarities in spouses' preferences for leisure would imply that their utilities are determined simultaneously. In this context, separate estimation of the spouses' retirement equations would yield biased parameter estimates.

This point was first noticed by Burtless (1990), and since then many studies have tested for, and found evidence of complementarity in couples' preferences⁴. However, all these studies being reduced-form analyses, they do not provide evidence of the degree of the bias affecting estimates from structural models of individual behaviour which take spouses' decisions as exogenous or, more importantly, the

¹Board of Trustees. 2007 OASDI and Medicare Reports. See references.

²See, among others, Hurd (1990a), Stock and Wise (1990), Gustman and Steinmeier (1994), Blau (1994 and 2004), Rust and Phelan (1997), Coile and Gruber (2000 and 2001), French and Jones (2004b) and French (2005).

³Evidence of joint retirement of US couples is found in the New Beneficiary Survey (Hurd (1990a)), the National Longitudinal Survey of Mature Women (Gustman and Steinmeier (2000)), the Retirement History Study (Blau (1998)) and the Health and Retirement Study (Michaud (2003)). Banks, Blundell and Casanova (2007) find evidence of joint retirement of couples from the English Longitudinal Study of Ageing.

⁴See, among others, Gustman and Steinmeier (2000), Coile (2004a) and Banks, Blundell and Casanova (2007).

predictions of policy effects based on such models.

The rest of the paper is organised as follows: section 2 presents an overview of the main incentives to retirement facing individuals and couples, and how these are captured in the theoretical model. A model that attempts to uncover the structure behind couples' retirement choices must first precisely account for incentives to retirement at the individual level and then add those incentives that are specific to couples. I discuss the role of employment income, Social Security, private pensions, health insurance and health as retirement determinants affecting both individuals and couples, and how they are captured in the theoretical model. Then, I concentrate on incentives to retirement that are specific to couples, such as the existence of non separabilities in spouses' preferences or correlations in unobservables across spouses.

Section 3 describes the theoretical model, which accounts for all the factors identified in section 2 as potential determinants of retirement decisions of couples in the estimation sample. The model includes a large set of observed state variables, that allow to replicate the existing heterogeneity in elderly couples' financial and health status. It also account for the presence of unobserved state variables, whose purpose is to rationalise the dispersion in couples' choices conditional on observed states.

Section 4 reviews the procedure used to solve and estimate the stochastic, dynamic, Markov model presented in section 3, which accounts for both discrete and continuous decisions. I show how my model fits into the framework introduced by Rust (1987, 1988) for the solution of stochastic discrete processes, and then extend it in order to allow for the presence of continuous choices -hours worked and savings in the case of the model presented in section 3.

Section 5 describes the data used for estimation, and presents estimation results for the laws of motion of the exogenous variables. I use a subsample of individuals from the Health and Retirement Study (HRS). I estimate a separate wage equation for men and women, accounting for selection in both cases, since retirement decisions are likely to be affected by shocks to wages. The procedure followed to correct for selection bias, described in Wooldridge (1995), accounts for the presence of a fixed effect in the wage equation. It involves a linearisation of the fixed effect which I will later use to generate wage draws for simulated individuals and sample individuals who are never observed working. Also as part of the wage equation estimation, I follow Meghir and Pistaferri (2004) to estimate the variance of the persistent component of wages. The estimate of the variance is also used as an input in the simulations of the theoretical model.

Other important inputs to the simulations are the evolution of health and medical care costs. Individuals in the HRS cohort have not reached 75 years of age in the last available wave. This means that they cannot be used to estimate the evolution of health and health care costs for individuals beyond that age. Therefore, I merge data from the HRS cohort with a cohort of individuals born before 1923. In the estimation of the health care costs process, I account for the cohort effects that are likely to affect the two samples differently.

Finally, section 6 provides simulations of the theoretical model using the parameters estimated in section 5. I take a baseline set of couples and describe their predicted optimal behaviour. Then I show how this behaviour would change as a response to changes in the couples state variables, including assets, wages, health status, Social Security entitlement and the degree of complementarity in spouses' leisure.

Then, I present simulation results that predict the effect of a policy change involving a reduction of the spousal benefit, which allows one spouse to claim a pension based on a percentage of her partner's entitlement, from 50 to 33%. This reduction has been suggested by the Congressional Budget Office in one of their regular reports to the Senate Committees on the Budget⁵. The results show that, despite women being the main beneficiaries of the spousal benefit, its reduction would mostly affect the participation

⁵Congressional Budget Office. Report to the House and Senate Committees on the Budget. February 2007. <http://www.cbo.gov/ftpdoc.cfm?index=7821>

behaviour of men: once their incentives to retire in order to provide their wife with a pension are reduced, some men choose to delay retirement. The simulations show which men are most likely to be affected by the change.

2 Overview

The purpose of this paper is twofold. First, it focuses on the impact of health, wealth, wages, Social Security and medical expenditures in older individuals' retirement decisions. Second, it analyses the phenomenon of joint retirement, that is, the observed tendency of spouses to retire within a short time from each other, independently of the age difference between them. Both are studied within a structural model of couples' saving and participation choices which carefully accounts for the main incentives to retirement facing married individuals.

Section 3 introduces a model of couples' behaviour within a life-cycle framework. So as to be able to replicate the dispersion in household savings and spouses' participation rates and number of hours worked observed in the data, the model allows for a large degree of heterogeneity across couples, both in observed variables -such as health status, wages, wealth and accumulated lifetime earnings- and unobserved ones -such as spouses' taste for work. Moreover, the budget constraint facing couples is modelled in a precise way, so as to capture the complex institutional environment facing couples close to retirement age. Finally, in order to account for the main sources of correlation in spouses' preferences, the model allows for complementarity in spouses' leisure -one or both spouses may enjoy leisure more when it is shared by their partner- and correlation in spouses' unobserved tastes for work.

This section motivates the theoretical model, concentrating on the retirement decision: it discusses observed retirement patterns for individuals and couples, the main age-specific incentives to retirement, and how those relevant for the estimation sample are captured in the model.

2.1 Incentives to retirement from the individual perspective

Life-cycle models predict that households will accumulate savings during their working life in order to finance retirement. Given that the interest of this paper is in older couples, we would expect most of them to have accumulated a significant amount of wealth by the time they are first observed, already in their fifties. Nevertheless, 55% of the couples interviewed in the first survey wave report a net value of financial wealth -which excludes housing wealth- of less than \$10,000. Unless all these couples intend to use their primary residence to finance their retirement, it would seem that their savings are far too low to support them into old age. Financial savings, however, are not the only route to financing retirement. Income from work, Social Security pensions, private pensions and health insurance are alternative sources of funds during the retirement years. Below I discuss the incentives for retirement at particular ages provided by each of these, how they are captured into the model and how they determine the choice of estimation sample.

Labour income

Employment from work is an important source of income for a significant number of older couples. Figure 1 shows the weight of different sources of income for couples where the husband is aged 65 to 70, by income quintiles. According to the graph, income from employment represents around 25 percent of total income for households in the three highest quintiles. Some of the individuals who work past age 65 will never have retired; and yet others will have retired and then re-entered the work force -27% of married men and 35% of married women working after age 65 have had at least one spell of inactivity in

the previous four years). Both types of behaviour are allowed for in the model, where non-participation is not an absorbing state.

Work income is related to the number of hours worked, one of the continuous choices of couples in the model. Most individuals do not jump discontinuously from full-time work into full-retirement. Instead, they choose to gradually reduce the number of working hours. In order to capture the dispersion in the hours distributions for individuals not clustered at 40 hours per week, hours worked by the spouses will be treated as continuous variables.

Labour income interacts with other state variables in the model. For instance, the ability to earn income from employment can be impaired by an individual's health status. This is particularly important for individuals approaching the retirement years, as they anticipate a quicker decline in health status than ever before in their working lives careers. The role of health as a determinant of hours worked is captured in the model through a fixed cost of being in bad health which diminishes the maximum number of available hours to divide between leisure and work. The effect of health on productivity is explored in the estimation of the wage equation, described in section 5.2.

Finally, individuals who work while receiving Social Security benefits are subject to the earnings test. Accordingly, part of their benefits is withheld whenever their earnings rise beyond a threshold. This is accounted for in the model's Social Security function, described in detail in section 3.7.

Social Security

Social Security benefits represent a source of retirement income for most of the older population. In 2005, 90% of the population aged over 65 received benefits from the Social Security, and for 65% of the aged these benefits represented more than half their income⁶. Individuals can claim their full Social Security benefit entitlement by age 65. They can claim benefits as early as age 62, subject to a reduction of 6.7% per year, which is roughly actuarially fair. Liquidity-constrained individuals may have incentives to remain in work until Social Security benefits become available at age 62.

Benefit claiming can also be delayed beyond age 65. For each year of delay between 65 and 70, benefits are increased by a rate of 5.5%⁷ per year, which is less than actuarially fair. Individuals who work after claiming Social Security benefits are subject to the earnings test. This implies that their benefits are withheld at a rate of \$1 for every \$2 of earnings above a threshold if they are younger than 65, and \$1 for every \$3 above a threshold if they are aged between 65 and 70. The combination of the less than actuarially fair increase for delayed retirement and the earnings test gives incentives to claim benefits by age 65 and disincentives to work beyond that age.

Figure 2 shows retirement frequencies at different ages for married men and women. The spikes at ages 62 and 65, which can be appreciated for both sexes, are well documented in the retirement literature. Part of the peak in retirement at these ages has been attributed to Social Security incentives (Gustman and Steinmeier (1986), Rust and Phelan (1997), French (2005)).

Social Security rules are carefully captured in the theoretical model in section 3. A description of the main Social Security provisions is provided in section 3.7. So as to simplify the dynamic program, however, the decision to apply for Social Security benefits is not considered explicitly. Instead, it is assumed that individuals start claiming the first year they are observed not to work after age 62. Figures 3 and 4 in the appendix use the Social Security records of HRS respondent to compare the actual claiming age with the one assumed in the model. The two series are very close for men. For women, the assumed Social Security claiming date overpredicts the peak at age 62⁸. On the whole, however, the approximation seems to work

⁶Social Security Administration. "Fast Facts & Figures About Social Security, 2007" http://www.ssa.gov/policy/docs/chartbooks/fast_facts/2007/fast_facts07.pdf

⁷This rate is applicable to individuals born in the years 1933 and 1934. The average individual in the estimation sample was born in 1934.

⁸The discrepancy is mainly due to a significant proportion of women who start receiving benefits before the age of 62.

quite reasonably.

Private Pensions

An important source of incentives to retirement are private pensions. In particular, defined benefit (DB) pensions give strong incentives to retirement at specific ages: after a certain number of years of service in a firm, or past the early or normal retirement ages, the rate of pension accrual is greatly reduced and can even become negative. For a large proportion of DB pension holders, these incentives are likely to dominate those provided by Social Security provisions (Lumsdaine, Stock and Wise, 1994a). Benefits from defined contribution (DC) pensions, on the other hand, are typically determined only by the amount of assets accumulated in the plan at the time of retirement, and they provide no specific incentives that encourage or discourage retirement at specific ages (Lumsdaine, Stock and Wise, 1994b). Nevertheless, most DC pensions, such as 401(k) plans or IRAs, specify an earliest withdrawal age. Withdrawing benefits from the plan before this age is strongly penalised. This may lead liquidity-constrained individuals to remain in work while their money is locked up in their DC pension plan.

Figure 5 shows retirement frequencies as a function of age for men with different pension types. It is clear that DB pension holders are much more likely than DC ones to retire before the Social Security incentives kick in at age 62. Moreover, part of the exit frequencies at ages 62 and 65 for individuals with a DB pension are likely to be due to their pension plan's characteristics, rather than Social Security provisions: the most common ages in the distribution of normal retirement ages for DB pension holders are 65 and 62, followed by 55, and the rest distributed between 56 and 60. The most common early retirement ages are 62 and 55 (Károlyi, Maestas and Zissimopoulos, 2006).

The tendency of DB pension holders to retire early is confirmed by table 2: Men who have a DB pension plan are 17 percentage points less likely than DC plan holders to be employed by the time they become entitled to Social Security benefits at age 62.

Figure 6 shows retirement frequencies for women, by pension type. Even though the difference is not so noticeable as for men, DB pension holders are still more likely to retire before the age of early Social Security entitlement than DC pension holders. According to table 2, women who have a DB pension plan are 6 percentage points more likely than those who have a DC pension plan to have retired by the time they become 62.

Introducing private pension incentives into a dynamic model implies adding a sufficient number of state variables to describe pension characteristics. In the case of DB pensions, these variables would have to include the early and/or normal retirement age, a measure of job tenure and the wage. In a model of couples such as the one presented in section 3, separate state variables would have to be added for men and women, and this would render the programme computationally intractable.

Ignoring the role of DB pensions, on the other hand, would disregard an important retirement incentive. Using the sample of DB pension holders to estimate a model that does not account for DB provisions would create problems in fitting the behaviour of those who retire before age 60, upon reaching their plan's early retirement age -and in the absence of any health, health cost or wage shock. Moreover, the model would likely attribute to Social Security incentives the retirement exits of individuals whose DB-plan early or normal retirement ages are 62 and 65.

In order to maintain a computationally-tractable number of state variables, while still accounting for the main incentives to retirement of the individuals in the sample, I restrict the estimation sample to couples with no private pension or one or more DC plans. DC pension holdings are treated in the model as part of household wealth. While this can be a reasonable approximation for non-liquidity constrained

It is possible for a non-disabled woman to claim benefits at age 60 or before in exceptional circumstances. She should be a widow, and now cohabiting with a new partner (in the model, cohabiting couples are counted as married couples) or taking care of young children.

individuals, it is possible that a minority of DC pension holders may be obliged to remain in work until the earliest age at which their DC pension funds become available. Nevertheless, this is not likely to be an issue for the majority of the estimation sample, given that only couples where the husband is aged 58 or older are selected, and the typical early withdrawal age is 59 and a half.

A more important concern is the special tax treatment of DC plans. Most DC pension plans allow workers to defer income taxes on plan contributions until withdrawal. The tax-deferred nature of DC-plan assets is not accounted for in the model, which may lead to underestimation of couples' incentives to save in this type of plans.

Couples with no private pension and those where one or both of the spouses have a DC pension are considered together in the estimation sample in order to attain a reasonable sample size. It is important, though, to bear in mind that individuals who have no private pension have quite different characteristics from those with a DC plan. Table 2 shows that they tend to belong to poorer households, have worse health, less education and lower wages. The key assumption that allows to model these two groups together is that none of them face incentives from a pension plan to retire at particular ages. The model in section 3 is rich enough to account for other observable and unobservable differences between the two: differences in health, wages and household wealth are captured through the initial conditions for these variables. The effect of education is captured through the initial conditions on wages and wealth. Moreover, the model allows for heterogeneity in spouses' taste for work, which may also be related to pension type.

Health Insurance

A source of incentives to retirement often considered in the literature is the type of health insurance coverage. Gustman and Steinmeier (1994), Rust and Phelan (1997), French and Jones (2004b) and Blau and Gilleskie (2006) distinguish three types of individuals according to the type of health insurance coverage: those whose health insurance is tied to their job, and would lose their coverage if they retired -i.e. individuals with "tied" coverage-; those who can keep their health insurance even if they retire from their job before age 65 -individuals with "retiree" coverage-; and those with no work-related health insurance -individuals with no coverage⁹. They argue that individuals with tied coverage will have stronger incentives to remain in work until they become eligible for state-provided Medicare coverage at 65 than those with retiree or no coverage. Gustman and Steinmeier and Blau and Gilleskie find that the effect of health insurance on retirement behaviour is small. Rust and Phelan find that the effect is large for the subsample of individuals without a private pension. However, their model ignores the role of savings as insurance against medical shocks, and is thus likely to overestimate the importance of health insurance. Finally, French and Jones estimate a dynamic model with savings and participation decisions using the HRS data and find that individuals whose health insurance is tied to the job leave the labour force on average half a year later than workers with retiree coverage.

None of these studies models explicitly the relationship between health insurance and pension type. However, it can be seen from table 2 that there is a correlation between the two: individuals with no pension are the most likely to have no health insurance; individuals with a DB pension plan are the most likely to have retiree coverage; and individuals with DC pension plans are the most likely to have tied coverage. In their paper, French and Jones acknowledge this correlation, but do not control separately for health insurance and pension type. Instead, given that people with retiree coverage are the most likely to have a DB plan, French and Jones assign them the sharpest drops in pension accrual after age 59. In this way, they compound the effect of health insurance and pension type, and thus it is not clear what part of the later retirements of people with tied coverage is due to the type of health insurance, and what part is due to them being more likely to have a DC pension.

⁹Even though these individuals are referred to as having no coverage, some of them will have government provided or private health insurance. The key is that their coverage is not employer-provided.

In the absence of a model that explicitly accounts for pension type, I choose not to control for health insurance type either. I therefore ignore any incentives that individuals with tied coverage may have to remain in work for longer than the rest. The estimate of French and Jones that those with retiree coverage and a DB pension retire half a year earlier than those with tied coverage and a DC pension is likely to be a higher bound for individuals in my estimation sample, given that I drop all observations with a DB pension plan.

2.2 Incentives to retirement from the couple's perspective

A growing share of the retirement literature characterises retirement as a decision concerning the couple, rather than the individual¹⁰. This follows the observation that a significant share of spouses retire within less than one year of each other, independently of the age difference between them. Evidence of this phenomenon, known as joint retirement, has been found in surveys dealing with couples from several generations and countries, such as the New Beneficiary Survey (Hurd (1990a)), the National Longitudinal Survey of Mature Women (Gustman and Steinmeier (2000)), the Retirement History Study (Blau (1998)) and the Health and Retirement Study (Michaud (2003)) in the US, and the English Longitudinal Study of Ageing (Banks, Blundell and Casanova (2007)) in England.

Figure 7 shows the difference in retirement dates¹¹ for HRS couples whose members had retired by the year 2004. Each graph corresponds to a particular age difference across spouses¹²: the first graph shows the distribution of retirement date differences for couples where the husband is at least one year younger than the wife; the second graph shows couples where the husband is the same age as the wife; and so on. In all of the 6 graphs, the highest frequency corresponds to a retirement date difference of zero, that is, to spouses retiring on the same calendar year.

Further evidence of correlated retirement outcomes is provided in table 3, according to which men aged 50 to 60 who are married to working women (shown in the first and second rows) are 7 percentage points more likely to transit out of the work force between period $t - 1$ and t if their wife also stops working. For men older than 60, the corresponding difference amounts to 17 percentage points. Similarly, table 4 shows that women aged 50 to 60 and married to working men are 10 percentage points more likely to stop working between periods $t - 1$ and t if their husband also stops working. Women older than 60 are 15 percentage points more likely to transit following their husbands¹³.

There are many reasons that may lead to joint retirement decisions: financial incentives, desire of the spouses to spend time together, common shocks, caring needs of one spouse, children or grandchildren, etc. These can be broadly classified in four categories: sorting of spouses according to their tastes for leisure; correlation in observable variables such as assets, wages, pension incentives, health status, etc.; correlation in time-varying shocks; and complementarity in leisure, so that spouses enjoy retirement more if their partner is retired as well. Past research finds that most of the coordination operates through the spouses' preferences, either because of complementarities in leisure or correlated tastes for work.

It is important to distinguish between the two, because they have very different implications for couples' behaviour. Suppose that people with a strong taste for leisure tend to be married to each other, either

¹⁰See Gustman and Steinmeier(2000), Blau and Gilleskie (2004), Coile (2004a, 2004b), Michaud (2003), Michaud and Vermeulen (2004).

¹¹Retirement date difference is defined as the husband's retirement date minus the wife's retirement date. Hence positive values indicate that the husband retired at a later calendar date than the wife.

¹²Age difference is defined as age of the husband minus age of the wife.

¹³The higher degree of association in spouses' transitions when one of the partners is over 60 is mentioned in Banks, Blundell and Casanova (2007). In that paper we provide evidence of nonseparabilities in spouses' preferences for leisure, and observe that the retirement of the wife has stronger detrimental effects on the participation of older husbands. Men below age 60 who wish to follow their wife into retirement are likely to face more financial constraints than those who have reached early retirement age either for their Social Security pension or some private pension plan.

through matching in the marriage process or because they have developed similar tastes through their life together. In the absence of any leisure complementarity, and after controlling for the resulting income effect, an exogenous change in retirement provisions that forced the wife to delay her retirement should not have any effect on the husband’s behaviour. On the other hand, in the presence of leisure complementarity the husband would enjoy his retirement more when the wife is retired as well. Thus, if she was forced to delay retirement, he would have incentives to remain in work for longer, so that the two can retire together.

This paper presents the first structural dynamic model of couples’ behaviour which allows for leisure complementarities across spouses and correlation in unobserved tastes for work. Moreover, the model accounts in a precise way for financial incentives and the main observable characteristics of both spouses. Estimation of the model parameters will allow to disentangle the relative role that different incentives to retirement play in determining couples’ decisions.

3 Theoretical model

This section describes a dynamic stochastic model of labour supply and saving choices of households close to retirement age. Each household consists of two spouses (“husband” and “wife”) who make their decisions jointly. The model captures the sequential nature of the decision-making process, with households adjusting their behaviour in every period as the uncertainty regarding spouses’ wages, health status, survival and medical expenditures unfolds.

At each discrete period t , given initial assets and husband and wife’s wages, average lifetime earnings¹⁴ and health status, households make their decisions regarding consumption and spouses’ work hours in order to maximise the expected discounted value of remaining lifetime utility.

Retirement status is defined as a function of the participation decision: a spouse who chooses not to participate in a period when he is at least 62 years old is referred to as “retired”. Retirement is not an absorbing state, as retired individuals can go back to work in any future periods. Spouses’ decisions to apply for Social Security benefits are not modelled separately from the participation decision. Individuals are assumed to start receiving Social Security pension benefits the first period in which they choose not to work after age 62. Benefit claiming is an absorbing state: Social Security entitlement is determined the first time individuals claim benefits, and it is not possible for them to accrue more pension in future periods, even if they go back to work.

The model is restricted to married couples who stay married until one or both spouses die. Decisions of widowed individuals are not explicitly modelled.

3.1 Choice set

At each discrete period t , households make both discrete choices -i.e. whether to participate in the labour force- and continuous ones -i.e. hours worked by each spouse and household consumption and savings.

Some of these decisions are linked. In particular, the participation decision can be viewed as part of the hours choice -a spouse who supplies zero hours of work in a period is at the same time choosing not to participate. However, it is useful to formalise the model explicitly separating the continuous and discrete choices, assuming, without loss of generality, that households make decisions in two steps: first, they make the discrete choices, that is, whether each of the spouse will work. Then, they choose optimal household savings and hours supplied by each spouse *conditional* on the discrete alternative.

¹⁴Average lifetime earnings are one of the variables that determine pension entitlement at the time of retirement.

Both types of choices are described in detail below. For ease of exposition, I will talk about the “husband” or “wife”’s choices when referring to household decisions concerning one of the spouses’ variables, such as his or her hours of work. However, all decisions are made by the household, which acts as a sole individual who maximises a unique welfare function.

3.1.1 Discrete choices

The discrete choice variables are each spouse’s participation. As mentioned above, non-participation is not an absorbing state, and individuals can always go back to work after periods of inactivity. Therefore, the variables indicating participation status, d_t^j , can take on the values 1 or 0 in all periods:

$$\begin{cases} d_t^j = 1 & \text{if spouse } j \text{ works in period } t \\ d_t^j = 0 & \text{otherwise,} \end{cases}$$

where the superscript $j = m, f$ identifies the spouse, m being the husband or “male”, and f being the wife or “female”.

$\mathcal{D}^j$ is the set of discrete alternatives available to spouse j each period. It is defined as:

$$\mathcal{D}^j = \{w, r\}, \quad \text{for } j = m, f,$$

where w indicates that spouse j is working and r indicates that she is not working.

The set of 4 discrete alternatives available to the household each period is $\mathcal{D} = \mathcal{D}^m \times \mathcal{D}^f$. Elements of $\mathcal{D}$ are of the type $d = (d^m, d^f)$, where d^m refers to the husband’s participation status, and d^f to the wife’s. For example, $d_t = (1, 0)$ indicates that the husband works and the wife does not work in period t .

3.1.2 Continuous choices

In each period t , households optimally choose savings, s_t , and each spouse’s work hours, h_t^m and h_t^f , conditional on the discrete action d_t . $y(d_t)$ is the vector of continuous choice variables:

$$y_t = (s_t, h_t^m, h_t^f).$$

$\mathcal{C}_t$ is the choice set for the continuous controls conditional on the discrete alternative d_t and the state spaces z_t and x_t :

$$y_t \in \mathcal{C}_t(z_t, x_t; d_t) \subset R_+ \times [0, L] \times [0, L],$$

where L is the period’s leisure endowment.

The conditioning on the discrete alternative d_t indicates the restrictions on the continuous controls imposed by the discrete choice. For instance, conditional on the discrete option being $d_t = (r, r)$, according to which none of the spouses is working in period t , the following restrictions on hours apply: $h_t^m(r, r) = 0$ and $h_t^f(r, r) = 0$.

3.2 State Space

The state space in period t consists of variables that are observed both by the agent and the econometrician and variables that are observed by the agent, but not by the econometrician. The vector of observed states is the minimal set of variables needed to model the interaction of savings, Social Security and labour income of both spouses:

$$z_t = \{A_t, E_t^m, E_t^f, w_t^m, w_t^f, M_t^m, M_t^f, B_{t-1}^m, B_{t-1}^f, age_t^m, age_t^f\},$$

where A_t are assets available at the beginning of period t , E_t^j is a measure of spouse j 's lifetime accumulated earnings¹⁵, w_t^j spouse j 's hourly wage, M_t^j spouse j 's health status, B_{t-1}^j an indicator of whether spouse j has started claiming benefits before period t and age_t^j spouse j 's age in years.

The vector of unobserved state variables is

$$x_t = \{\varphi^m, \varphi^f, \varepsilon_t\}.$$

φ^m and φ^f measure the husband and wife's taste for work, respectively. These variables are known by the household when making decisions, but not observed by the econometrician. They are the source of time-constant heterogeneity across couples. ε_t is a vector of unobserved utility components associated to the discrete alternative chosen by the household:

$$\varepsilon_t = \{\varepsilon_t(d_t) | d_t \in \mathcal{D}\},$$

where $\varepsilon_t(d_t)$ affects the utility derived from alternative d at time t . The value of the vector ε_t is known by the agent when making decisions in period t .

The unobserved states generate dispersion in couples' choices conditional on observables.

3.3 Preferences

Household utility in period t is defined as the sum of each spouse's utility plus an unobserved component, $\varepsilon_t(d_t)$, associated to the discrete choice and assumed known by the household:

$$U(d_t, y_t; z_t, x_t, \theta_1) = \phi u^m(c_t, l_t^m) + (1 - \phi) u^f(c_t, l_t^f) + \varepsilon_t(d_t), \quad (3.1)$$

where ϕ represents some household sharing rule and θ_1 is the vector of preference parameters.

Within-period utility for each spouse, u^j , is assumed non-decreasing and twice differentiable in consumption, c_t , and own leisure, l_t^j . In the empirical part of the paper, the function u^j is assumed to take the following general CES form, where ρ is the coefficient of relative risk aversion, $1/(1-\gamma)$ determines the elasticity of substitution between consumption and leisure and α_1 measures the share of consumption's contribution to the spouse's period utility:

$$u^j(c_t, l_t^j; z_t, \varphi^j, \theta_1) = \frac{1}{1-\rho} \left(\alpha_1 c_t^\gamma + (1-\alpha_1) \omega_t^j (l_t^j)^\gamma \right)^{\frac{(1-\rho)}{\gamma}}$$

The weight of leisure, ω_t^j , is given by:

$$\omega_t^j = \exp(\alpha_0 + \alpha_2^j(1 - d_t^k) + \varphi^j d_t^j), \quad \text{for } j \neq k$$

The parameters α_2^m and α_2^f measure the degree of complementarity in spouses' leisure. They are not restricted to being equal, thus allowing for asymmetrical complementarity of the type found in empirical work on spouses' retirement choices (Coile (2004)). The utility derived from own leisure increases when the partner is retired if there is complementarity in spouses' leisure (i.e. $\alpha_2^k > 0$); decreases when the partner is retired if there is substitutability (i.e. $\alpha_2^k < 0$); and is unaffected by the partner's leisure if both leisures are unrelated (i.e. $\alpha_2^k = 0$).

¹⁵ E_t^j are spouse j 's average indexed monthly earnings, or AIME, the measure used by the Social Security Administration, together with age at retirement, to determine benefit entitlement.

Moreover, utility from own leisure increases (decreases) in proportion to φ^j when spouse j is working and has a (dis)taste for work.

Finally, own leisure is equal to:

$$l_t^i = L - h_t^j - \tau^j P_t^j - \delta^j I(M_t^j = \text{bad}),$$

where L is the leisure endowment, h_t^j the number of hours worked by individual j in period t , τ_j a gender-specific fixed cost of work¹⁶, I is the indicator function, and δ_j a gender-specific fixed cost of bad health. These fixed costs are measured in hours of leisure lost during the period.

3.4 Budget constraint

Household income in each period consists of asset income rA_t ; husband's labour income $w_t^m h_t^m$; wife's labour income $w_t^f h_t^f$; husband and wife's Social Security benefits ssb_t^m and ssb_t^f ; and government transfers T_t . Post-tax resources are allocated between household consumption, c_t , and savings, s_t . The budget constraint can be written as:

$$c_t + s_t = A_t + Y(rA_t, w_t^m h_t^m, w_t^f h_t^f, \tau) + ssb_t^m + ssb_t^f + T_t,$$

where post-tax income Y is a function of taxable income and the tax-rate vector τ ; r is the interest rate, assumed constant; and w_j the hourly wage¹⁷.

Next period's available assets are determined by subtracting from savings husband and wife's out-of-pocket medical costs, hc_t^m and hc_t^f , which are assumed to be realised after the saving and work decisions:

$$A_{t+1} = s_t - hc_t^m - hc_t^f$$

The following liquidity constraint is imposed:

$$s_t \geq 0$$

The borrowing constraint implies that the household net worth at the beginning of a period can be negative if the sum of the husband and wife's health costs exceed savings¹⁸.

Following Hubbard, Skinner and Zeldes (1995), government transfers are parameterised as:

$$T_t = \min \left\{ c_{\min}, \max\{0, c_{\min} - (A_t + Y_t + ssb_t^m + ssb_t^f)\} \right\}$$

Transfer payments guarantee a minimum amount of resources for the household in every period equal to $c_{\min}$. The transfer function is intended to capture the penalty on saving behaviour that means-tested programmes such as Medicaid, Supplemental Security Income (SSI) or food stamps impose on low-asset households.

¹⁶French and Jones (2004b) argue that including a fixed cost of work allows to capture the empirical regularity that hours of work are clustered around 2000 hours and 0 hours, and the distribution outside those values is roughly uniform. They include a fixed cost, rather than discretise the choice set for hours (into, say, full-time, part-time and none), because it provides a better way to capture this dispersion.

¹⁷The rest of the variables in the budget constraint are normalised by the maximum number of hours an individual can work in a year, so that they are all measured in comparable terms.

¹⁸French and Jones (2004b) argue that this is a reasonable assumption in view of the number of HRS households who report medical expense debt.

3.5 Wage process

The logarithm of wages is modelled as the sum of a function of observables and a persistent error term:

$$\ln w_{it} = X_{it}\beta + u_{it}^P \quad (3.2)$$

$$(3.3)$$

The persistent component v_{it}^P is assumed to follow a random walk:

$$\begin{aligned} v_{it}^P &= v_{it-1}^P + \zeta_{it}, \\ \zeta_{it} &\sim N(0, \sigma_\zeta^2) \end{aligned}$$

Section 5.2 describes the estimation of the wage processes for men and women and discusses the role of transitory shocks and measurement error in observed wages. It also describes the procedure used to correct for selection in the presence of the persistent component.

Involuntary unemployment is not considered, that is, any individual can find a job at the period's wage. In this context, shocks to wages can be interpreted as shocks to productivity.

3.6 Out of Pocket Medical Expenditures

Out-of-pocket medical expenditures, hc_t^j , are modelled as a polynomial on health status M_t^j and age age_t^j , plus a random term ψ_t^j . The process for health costs is estimated separately for men and women. Omitting the superscript j :

$$\begin{aligned} \ln hc_{it} &= hc(age_{it}, M_{it}) + \psi_{it}, \\ \psi_{it} &\sim N(0, \sigma_\psi^2) \end{aligned} \quad (3.4)$$

In order to keep to a minimum the number of state variables in the dynamic problem, I do not account for persistence in the health cost process beyond that due to persistence in health status. Following Hubbard, Skinner and Zeldes (1995), French and Jones (2004a, 2004b) decompose the idiosyncratic error in (3.4) into a persistent and a transitory component. Estimates of the parameters for these processes imply a serial correlation between ψ_{it} and ψ_{it+1} of around 0.3. By ignoring this, I am likely underestimating the degree of health risk, since draws for the idiosyncratic component conditional on health status are likely to have longer lasting effects than I account for.

3.7 Social Security Benefits

A vast number of studies¹⁹ have analysed how the US Social Security system shapes workers' labour supply incentives. Although there is no agreement on the magnitude of the effect, there seems to be no doubt that the Social Security plays a role in determining labour force behaviour. In particular, it appears to be closely related to the spikes in retirement rates at 62 and 65 years of age.

Because of this important role, the Social Security system must be modelled as precisely as possible, in order to account in an accurate and realistic way for the incentives it provides. Its main features -all

¹⁹See, among others, Hurd (1990), Blau (1994), Rust and Phelan (1997), Diamond and Gruber (1998), Coile and Gruber (2001), French (2003), Coile (2004b) and Michaud (2004).

of which are captured in the theoretical model- are the following²⁰:

The level of Social Security benefits is determined from a worker's lifetime earnings in several steps. First, past annual earnings are indexed to account for changes in the national average wage. Then, indexed earnings are averaged over the worker's 35 highest-earnings years. This yields the so-called averaged indexed monthly earnings (AIME). Finally, a formula is applied to AIME to obtain the worker's basic benefit, or primary insurance amount (PIA). This formula is weighed in favour of relatively low earners, so that the replacement rate falls as the level of earnings rises.

Individuals receive benefits equal to their PIA if they retire at normal retirement age, which for individuals in the sample was 65 years of age. Workers can claim benefits as early as age 62, but they have their PIA reduced by the equivalent to 6.7% per each year between retirement age and age 65. This rate is roughly actuarially fair. Workers can also delay claiming benefits beyond age 65, and have their benefits increased by 4% per year of delay beyond 65 and before 70. This is less than actuarially fair and, thus, a major disincentive to delay claiming beyond age 65.

Once a worker has claimed benefits, these will be adjusted every year for increases in CPI, hence the Social Security provides a real annuity.

Individuals who claim benefits and keep working are subject to an *earnings test*. Social Security beneficiaries below age 65 have their benefits taxed by \$1 for each \$2 of earnings above a threshold of \$7,440. For every year worth of benefits taxed away, their future benefits are increased by 6.7%. Beneficiaries aged 65 to 70 that threshold. Workers between 65 and 70 have their benefits taxed by \$1 for each \$3 if earnings over a threshold of \$10,200. For every year worth of benefits taxed away, their future benefits are increased by 4%. Again, this is far from actuarially fair, and hence a strong disincentive to work beyond age 65.

The spouses of Social Security beneficiaries are entitled to Social Security benefits in some circumstances. A widow or widower is entitled to a benefit equal to up to their partner's full amount when they reach age 60. More importantly, spouses are entitled to a benefit equal to up to one half of their partner's full amount when they reach 62. This provision -known as the dependent spouse benefit- introduces potential correlation in spouses' incentives to claim benefits. For instance, a husband may find it worthwhile to delay retirement one year in order not only to increase his own pension, but also that of his wife.

The formulae used to approximate individual Social Security benefits in the model, which take all these features of the system into account, are described in detail in the appendix.

3.8 Health transitions and mortality rates

Health affects utility, through its effect on the leisure endowment, potentially worker productivity, through the wage process, and survival probability. It is assumed to be a discrete variable which can only take on the values "good" or "bad". The evolution of health is modelled as a first-order Markov process, with transition probabilities that depend on age and sex:

$$\pi_{k,j}^i(t) = \Pr(M_{t+1}^i = j \mid M_t^i = k, age_t^i, i), \quad k, j \in \{\text{good}, \text{bad}\}, i \in \{\text{m}, \text{f}\}$$

The health processes of the two spouses are assumed independent.

Mortality rates are a function of previous health status, age and sex. In particular, the probability that an individual who is alive in period t survives to period $t + 1$ is:

$$s_{t+1}^i = s(M_t^i, age_t^i, i), \quad i \in \{\text{m}, \text{f}\}$$

²⁰I describe the benefit system corresponding to the year 1992, which is the one used in the empirical model. See the *Annual Statistical Supplement to the Social Security Bulletin*, years 1993 to 2005, for more information.

3.9 Terminal value functions and bequest function

Upon death of one spouse, the behaviour of the surviving partner is not modelled. Their remaining lifetime utility is represented by the terminal value functions B^f or B^m -depending on whether the wife or the husband survives.

$$B^i(z_t) = \theta_i \frac{(W_t^i + K)^{\alpha_1(1-\rho)}}{(1-\rho)} \quad i = m, f,$$

$$W_t^i = A_t + PDV_t(\max(ssb_t^i, ssb_t^j)) \quad i = m, f,$$

where K is a parameter determining function B^i 's curvature and W_t is the sum of assets available at the beginning of the period following the death of one spouse and the present discounted value of the remaining spouse's Social Security benefits -the surviving spouse is entitled to the highest between his own benefits and those of the deceased spouse.

If none of the spouses reaches period t alive, the household still derives utility from assets bequeathed to survivors, A_t , according to the function B^b .

$$B^b(A_t) = \theta_b \frac{(A_t + K)^{\alpha_1(1-\rho)}}{(1-\rho)}$$

4 Model Solution

The eventual objective of this work is to use the observed realisations of household choices and observed states $\{d_t, y_t, z_t\}$ to estimate the unknown parameters describing couples' behaviour. Given the complexity of the problem described in section 3, this is not an easy task. In what follows, I introduce a series of assumptions that will allow to derive a computationally tractable GMM criterion associated to the household problem, which can eventually be used to estimate the model parameters.

It follows from the description in the previous sections of the laws of motion for the state variables that the agents' beliefs about uncertain future states can be represented by a first-order Markov probability density function. That is, the joint stochastic process for $\{z_t, x_t\}$ has an associated density function that can be written as:

$$f(z_{t+1}, x_{t+1} | d_t, y_t, z_t, x_t, d_{t-1}, y_{t-1}, \dots) = f(z_{t+1}, x_{t+1} | d_t, y_t, z_t, \varepsilon_t, \theta_2, \theta_3), \quad (4.1)$$

where the vectors of parameters θ_2 and θ_3 , to be estimated from the data, will be defined below. The fact that the distribution of states in period t is only affected by the realised states and choices in period $t-1$ greatly increases the computational tractability of the problem, as it implies that it is not necessary to keep track of the whole history of the process.

There is an extensive literature dealing with the solution and estimation of this type of stochastic Markov programs. The framework was first introduced by Rust (1987, 1988) and subsequently extended in Hotz and Miller (1993) and Aguirregabiria and Mira (2002), among others. Both the theoretical literature and subsequent applications²¹ focus on discrete decision processes. However, for the couples in my sample, more information is available than their discrete choices and observable states, namely their savings and hours choices, which are also an integral part of the model presented in section 3. Hence,

²¹see Eckstein and Wolpin (1989), Rust (1994) and Miller (1997) for surveys on the topic and Rust and Phelan (1997), Hotz and Miller (1993), Keane and Wolpin (1997) and Gilleskie (1998) for applications with discrete choice sets.

below I extend the statistical model for discrete Markov decision processes so as to integrate the continuous choices.

For ease of exposition, the solution procedure described below corresponds to a model without unobserved heterogeneity in spouses' tastes for work, where the only unobserved state variables are the choice-specific disturbances ε_t . The presence of heterogeneity in spouses tastes would simply add two dimension to the state space in each period.

4.1 Optimisation problem

In order to solve the finite-horizon Markovian decision problem, the household chooses a sequence of decision rules $\Pi = \{\pi_0, \pi_1, \dots, \pi_T\}$, where $\pi_t(z_t, \varepsilon_t) = (d_t, y_t)$, to maximise expected discounted utility over the lifetime. The value function of the problem is defined as:

$$V_t(z_t, \varepsilon_t) = \sup_{\Pi} E_t \left\{ \sum_{j=t}^T \beta^{j-t} [U(d_t, y_t; z_t, \varepsilon_t, \theta_1)] \mid z_t, \varepsilon_t \right\}, \quad (4.2)$$

subject to

$$\begin{aligned} c_t + s_t &= A_t + Y(rA_t, w_t^m h_t^m, w_t^f h_t^f, \tau) + ssb_t^m + ssb_t^f + T_t, \\ A_{t+1} &= s_t - hc_t^m - hc_t^f \\ h_t^i &\geq 0, \quad \text{for } i = m, f \end{aligned}$$

The expectation in 4.2 is taken with respect to the controlled stochastic process $\{z_t, x_t\}$, with probability distribution given by equation 4.1.

Since this is a finite horizon problem, the feasible set of household choices is compact, and the utility function continuous and integrable, the value function and the optimal policy always exist by backward induction, given an appropriate terminal condition. For any period t , the value function is given by:

$$V_t(z_t, \varepsilon_t) = \max_{d_t, y_t} [U(d_t, y_t, z_t, \varepsilon_t, \theta_1) + \beta E_t V_{t+1}(z_{t+1}, \varepsilon_{t+1})] \quad (4.3)$$

And the optimal policy for period t , π_T^* , is determined from Bellman's equation by the identity:

$$\pi_t(z_t, \varepsilon_t, \Theta) = \arg \max_{d_t, y_t} [U(d_t, y_t, z_t, \varepsilon_t, \theta_1) + \beta E_t V_{t+1}(z_{t+1}, \varepsilon_{t+1})],$$

with $\Theta = \{\theta_1, \theta_2, \theta_3\}$.

The optimal policy for the model presented in section 3 could be computed by backward recursion, starting from period T and at each period computing or approximating the value function V_t at any feasible point in the state space, so that it can be recovered when integrating over the random state variables from period $t - 1$. There are, however, several difficulties with this approach, associated to the presence of the vector ε_t . On the one hand, this vector adds up to 4 dimensions to the state space in every period, which must be added to the already high-dimensional vector z_t . Even taking a rough grid approximation of the (continuous) vector ε_t , this considerably increases the number of points at which the problem needs to be solved in every period. Moreover, the vector ε_t enters nonlinearly the unknown function $E_t V_{t+1}$. Hence it must be integrated out in order to obtain the conditional choice probabilities associated to each discrete option and the continuous policies. This would add 4 more dimension to the 4-dimensional integrals over both spouses' wage and medical cost shocks.

In order to simplify the problem I make two assumption regarding the role of unobservable state variables that were first introduced by Rust (1988) and have been widely used in the literature²², namely:

ASSUMPTION 1 (Additivity): The within-period utility function U has the additively separable representation given by equation 3.1:

$$U(d_t, y_t, z_t, \varepsilon_t, \theta_1) = \phi u^m(d_t, y_t, z_t, \theta_1) + (1 - \phi) u^f(d_t, y_t, z_t, \theta_1) + \varepsilon_t(d_t) = u(d_t, y_t, z_t, \theta_1) + \varepsilon_t(d_t)$$

where $\varepsilon_t(d_t)$ is the d^{th} component of the $J \times 1$ vector ε_t . The support of $\varepsilon_t(d_t)$ is the real line for all d_t .

ASSUMPTION 2 (Conditional Independence): The conditional probability density function of the state variables factors as

$$f(z_{t+1}, \varepsilon_{t+1} | d_t, y_t, z_t, \varepsilon_t, \theta_2, \theta_3) = q(\varepsilon_{t+1} | z_{t+1}, \theta_2) g(z_{t+1} | z_t, d_t, y_t, \theta_3).$$

These assumptions render the household problem tractable by simplifying it in two important respects: (i) first, the additivity assumption facilitates the integration over ε_t when computing $E_t V_{t+1}$. Below I make a further assumption on the specific functional form for the distribution of ε_t which, together with the additivity assumption, yields closed form solutions for the integral. (ii) Second, the assumptions imply that V_t can be obtained from a version of the Bellman equation that involves the reduced state space $\{z_t\}$, rather than the much larger $\{z_t, \varepsilon_t\}$ ²³.

According to the additivity assumption, the Bellman equation can be written as:

$$V_t(z_t, \varepsilon_t) = \max_{d_t, y_t} [u(d_t, y_t, z_t, \theta_1) + \varepsilon_t(d_t) + \beta E_t V_{t+1}(z_{t+1})],$$

where the dependence of $E_t V_{t+1}$ on z_{t+1} only is intended to indicate that ε_t has been integrated out.

As explained in section 3, continuous choices are conditional on the discrete option chosen by the couple. This allows to rewrite the Bellman equation as:

$$V_t(z_t, \varepsilon_t) = \max_{d_t} \left[\max_{y_t} \{ [u(j, y_t, z_t, \theta_1) + \varepsilon_t(j) + \beta E_t V(z_{t+1})] \mid d_t = j \} \right] = \max_{d_t} \left\{ \left[\max_{y_t} \{ [u(j, y_t, z_t, \theta_1) + \beta E_t V(z_{t+1})] \mid d_t = j \} \right] + \varepsilon_t(d_t) \right\}, \quad (4.4)$$

where the equality follows because ε_t is independent of the choice of y_t . Equation (4.4) shows that the solution of period t 's problem can be divided, without loss of generality, in two stages. Proceeding backwards, there is first an inner maximisation with respect to the continuous choices *conditional* on the discrete choice $d_t = j$, for $j = 1, \dots, J$. Then, the option that yields the highest value is chosen. These two steps are described in detail below.

4.2 Inner maximisation

The first stage in the solution of period t 's problem involves the *choice-specific value function*, defined as:

$$r(z_t, j, \theta_1) \equiv \max_{y_t \in C(d_t)} \{ [u(j, y_t, z_t, \theta_1) + \beta E_t V_{t+1}(z_{t+1})] \mid d_t = j \} \quad (4.5)$$

²²See Rust (1987, 1988), Hotz and Miller (1993) and Aguirregabiria and Mira (2002).

²³For a proof see Rust (1988).

This function has to be computed for each possible value of d_t , subject to the contemporaneous budget constraints, the liquidity constraint and the constraints that hours worked by each spouse must be greater or equal to 0.

The associated first-order conditions for $d_t = j$ are:

$$\phi u_c^m(j, y_t, z_t, \theta_1) + (1 - \phi) u_c^f(j, y_t, z_t, \theta_1) - \beta(1 + r) E_t \frac{\partial V_{t+1}(z_{t+1})}{\partial A_{t+1}} + \eta_t^A = 0 \quad (4.6)$$

$$P_t^m \left[(\phi u_c^m(j, y_t, z_t, \theta_1) + (1 - \phi) u_c^m(j, y_t, z_t, \theta_1)) \frac{\partial Y_t}{\partial h_m} + \phi u_{h_m}^m(j, y_t, z_t, \theta_1) + (1 - \phi) u_{h_m}^f(j, y_t, z_t, \theta_1) + \beta E_t \left(\frac{\partial V_{t+1}}{\partial h^m} \right) + \eta_t^{h_m} \right] = 0 \quad (4.7)$$

$$P_t^f \left[(\phi u_c^m(j, y_t, z_t, \theta_1) + (1 - \phi) u_c^m(j, y_t, z_t, \theta_1)) \frac{\partial Y_t}{\partial h_m} + \phi u_{h_f}^m(j, y_t, z_t, \theta_1) + (1 - \phi) u_{h_f}^f(j, y_t, z_t, \theta_1) + \beta E_t \left(\frac{\partial V_{t+1}}{\partial h^f} \right) + \eta_t^{h_f} \right] = 0 \quad (4.8)$$

$$-\eta_t^A A_{t+1} = 0$$

$$-\eta_t^{h_m} h_t^m = 0$$

$$-\eta_t^{h_f} h_t^f = 0$$

$$\eta_t^A, \eta_t^{h_m}, \eta_t^{h_f} \geq 0$$

where Y_t represents post-tax income, η_t^A is the multiplier for the liquidity constraint and $\eta_t^{h_m}$ and $\eta_t^{h_f}$ for the male and female hours constraints, respectively.

The terms involving $u_{h_m}^f(j, y_t, z_t, \theta_1)$ and $u_{h_f}^m(j, y_t, z_t, \theta_1)$ are equal to zero in the absence of complementarity in spouses' leisure, that is, when α_2^f and α_2^m , respectively, are equal to zero.

Whenever a spouse has reached the maximum value of accumulated lifetime earnings, further hours of work will not increase her pension entitlement. Similarly, a spouse's Social Security entitlement is computed the first time she retires, and she cannot accumulate further benefits even if she goes back to work. In any of these two cases, work in period t will not have any effect on period $t + 1$'s value function, so the term $\frac{\partial V_{t+1}}{\partial h^j}$ will be zero for spouse j .

The unobserved utility components ε_t do not feature in the first-order conditions for a given $d_t = j$. This implies that any differences in choice-vectors y_t across households with equal realisations of the observed state vector z_t must be due to differences in spouses' tastes for work.

4.3 Outer maximisation

Let us define the vector $r(z_t, d_t, \theta_1) \equiv \{r(z_t, j, \theta_1) \mid d_t = j \in \mathcal{D}_t(z_t)\}$, where $r(z_t, j, \theta_1)$, representing the indirect utility function associated to option $d_t = j$, has been defined in equation 4.5.

Let us also define the surplus function corresponding to the density $q(\varepsilon_{t+1} | z_{t+1}, \theta_2)$ as

$$G(r(z_t, d_t, \theta_1)|z_t, \theta_2) \equiv \int_{\varepsilon} \left[\max_{d_t \in \mathcal{D}_t(z_t)} [r(z_t, d_t, \theta_1) + \varepsilon_t(d_t)] \right] q(\varepsilon_{t+1}|z_{t+1}, \theta_2) d\varepsilon_{t+1}.$$

If the conditional independence assumption holds, the conditional probability of choice $d_t = j$ is given by:

$$P(j|z_t, \theta_2) = G_j(r(z_t, j, \theta_1)|z_t, \theta_2),$$

where G_j denotes the partial derivative of G with respect to $r(z_t, j, \theta_1)$. And the function $E_t V_{t+1}$ can be obtained from:

$$E_t V_{t+1} = \int_x G(r(x, d_t, \theta_1)|z_t, \theta_2) g(z_{t+1}|x_t, d_t, y_t, \theta_3) dx$$

Choosing a specific functional form for $q(\varepsilon_t|z_t, \theta_2)$ yields a more explicit formula for the choice probabilities. In what follows, I assume that $q(\varepsilon_t|z_t, \theta_2)$ is a multivariate extreme value distribution with parameters $\theta_2 = (\mu, \sigma)$:

$$q(\varepsilon_{t+1}|z_{t+1}, \theta_2) = \prod_{j \in \mathcal{A}_t(z_t)} \exp\left\{-\frac{(\varepsilon_t(j) + \mu)}{\sigma}\right\} \exp\left\{-\exp\left\{-\frac{(\varepsilon_t(j) + \mu)}{\sigma}\right\}\right\}$$

Then the social surplus function G is given by:

$$G(r(z_t, d_t, \theta_1)|z_t, \theta_2) \equiv \sigma\gamma + \sigma \ln \left[\sum_{j \in \mathcal{D}_t(z_t)} \exp\left\{\frac{r(z_t, j, \theta_1)}{\sigma}\right\} \right],$$

where $\gamma = 0.577216$ is Euler's constant.

As it is well known in the random utility literature, the value of G in (4.3) is increasing in the number of options. This is a cause of concern in this particular application, where the agents do not have the same number of choices in every period. In order to avoid agents spuriously deriving value from maximising the number of options available in future periods, I normalise the location parameter and substitute μ by $\tilde{\mu}$, given by:

$$\tilde{\mu} = \mu + \sigma \log(1/J_t)$$

where J_t is the number of discrete choices available in period t . The social surplus function becomes:

$$G(r(z_t, d_t, \theta_1)|z_t, \theta_2) \equiv \sigma\gamma + \sigma \ln \left(\frac{1}{J_t} \left[\sum_{j \in \mathcal{D}_t(z_t)} \exp\left\{\frac{r(z_t, j, \theta_1)}{\sigma}\right\} \right] \right),$$

which is no longer increasing in J_t .

The conditional choice probabilities are given by the multinomial logit formula:

$$P(j|z_t, \theta_2) = \frac{\exp\{r(z_t, j, \theta_1)/\sigma\}}{\sum_{j \in \mathcal{D}_t(z_t)} \exp\{r(z_t, j, \theta_1)/\sigma\}},$$

and EV_{t+1} is given by:

$$EV_{t+1}(z_{t+1}, \varepsilon_{t+1}) = \sigma\gamma + \sigma \int \ln \left(\frac{1}{J_t} \left[\sum_{j \in \mathcal{D}_t(z_t)} \exp \{r(x, j, \theta_1)/\sigma\} \right] \right) g(x \mid z_t, d_t, y_t, \theta_3) dx.$$

5 Data and First Stage Results

5.1 Data

I use data from the Health and Retirement Study (HRS) to estimate the model. The HRS is a panel data set, representative of non-institutionalised individuals born between 1931 and 1941 and their spouses. There are currently 7 available waves, covering every two years from 1992 to 2004.

The HRS provides extensive information on economic status -including comprehensive measures of wealth, income from work, private pensions, Social Security and other government transfers-; health -including subjective assessments, ADLs, IDLs and evaluations of cognitive condition-; retirement -including expectations for individuals not yet retired- and demographics.

It is possible to merge the HRS data with administrative data from the Social Security Administration (SSA), which contain information on quarters of work in Social Security-covered employment, wage and self-employment income of individuals through their work histories. I use the SSA HRS data to estimate the regressions used to input initial values of Averaged Indexed Monthly Earnings (AIME) to sample individuals.

The initial HRS sample, after eliminating all couples who divorce, separate or change partner during the sample period, consists of 4,752 observations. Out of these, I drop 364 where at least one of the spouses reports receiving Social Security disability insurance (SSDI) before age 62. In order to maintain a homogeneous sample, I drop couples who have wealth over 1,125,000 1992 dollars at any wave, the 1% with the lowest wealth at every wave and those where any spouse has a wage below 3 1992 dollars per hour or above 50 1992 dollars per hour.

Of the 4,050 remaining couples, I use observations for all men who have no pension or a DC pension -independently of their wife's pension type- in the estimation of the men's wage process, and all women who have no pension or a DC pension in the estimation of the women's wage process. By not selecting the sample according to the partner's pension type I make the implicit assumption that this does not influence the wage. It is possible, however, that it does have an effect of participation. Indeed, an indicator of the partner's pension type will be included among the exclusion restrictions in the estimation of the selection process (see section 3.5).

The sample used to estimate the health cost process consists of couples where both partners have either no pension or a DC pension. Unlike in the case of wages, the partner's pension type does play an important role in the determination of out-of-pocket medical expenditures. This is because spouses -especially women- often have health coverage through their partner's insurance. Thus, since individuals with a DB pension are the most likely to have retiree health coverage, individuals whose partner has a DB pension are likely to benefit from their partner's retiree coverage. I minimise the amount of individuals with retiree coverage, either on their own or through their spouse, by dropping all DB pension holders from the sample.

Finally, for the estimation of the preference parameters I restrict the sample according to the constraints imposed by the dynamic programme. The ages of the spouses are state variables in the dynamic model, since a couple where both partners are 60 years of age solves a different optimisation problem from a couple where the husband is aged 60 and the wife 59. Hence, the model presented in section 3 has to be solved separately for each possible age gap between husband and wife. Therefore, I restrict the sample to

couples where the husband is not younger and at most 3 years older than the wife. Added to the restriction that both spouses must have either no private pension or a DC pension, the final sample consists of 789 couples and 3,780 couple-year observations.

5.2 Estimation of the wage equation

Observed wages are assumed to evolve according to the following specification:

$$\ln w_{it} = X_{it}^1 \beta^1 + u_{it}, \quad (5.1)$$

where u_{it} can be decomposed into a persistent component v_{it}^P , a transitory component v_{it}^T and i.i.d. measurement error m_{it} :

$$u_{it}^1 = v_{it}^P + v_{it}^T + m_{it}$$

The transitory component follows an MA process of order q :

$$v_{it}^T = \varepsilon_{it}^T + \sum_{j=1}^J \theta_j^T \varepsilon_{it-j}^T,$$

and the persistent component is a random walk:

$$v_{it}^P = v_{it-1}^P + \zeta_{it}$$

The permanent and transitory components are assumed to be uncorrelated at all leads and lags.

Repeatedly substituting backwards in time for the persistent component yields the following expression for observed wages at any point during the sample period:

$$\ln w_{it} = X_{it}^1 \beta^1 + v_{i0}^P + \sum_{j=0}^t \zeta_{ij} + v_{it}^T + m_{it}, \quad (5.2)$$

where v_{i0}^P is an unobserved individual initial condition that will be related to unobserved individual characteristics such as ability.

So as to solve and estimate the dynamic model presented in section 3, one must be able to assign initial wages to all simulated individuals, including those who eventually choose not to work. This poses two different problems. First, for those individuals who are never observed working during the sample period, there is no wage information available that can be used to infer the value of their initial conditions. Second, in the estimation of the vector of coefficients β^1 it is important to control for selection. Selectivity bias has been found to be important in the estimation of female wage equations (Heckman 1979, 1980). For the sample of workers in the retirement years, it can conceivably affect male wage equations as well, as men decide whether to retire in a particular period based on a set of variables which presumably includes the period's shocks to the permanent and transitory components, plus the initial condition.

In order to control for selection in the presence of the individual unobserved initial condition I follow the procedure described in Wooldridge (1995 and 2002). This allows to control for selection in the presence of an individual fixed effect in the wage equation and any type of serial correlation in the idiosyncratic errors of both the wage equation and the selection process. The procedure involves a linearisation of the fixed effect as a function of observables, which will eventually allow to generate initial conditions for all sample individuals, regardless of their participation status.

The variance of the shock to the persistent component is estimated following Meghir and Pistaferri (2004). As it is shown in their paper, the variances of the measurement error and the transitory component are not separately identified. In the theoretical model I make the assumption that individual decisions are affected by persistent shocks to wages only. Hence, the transitory component does not feature into the programme, and its variance does not need to be estimated.

I begin by describing the procedure to estimate β_1 and how it is applied to the data. The wage process is that of equation 5.2, which simplifies to

$$\ln w_{it} = X_{it}^1 \beta^1 + v_{i0}^P + u_{it}^1, \quad (5.3)$$

with $u_{it}^1 = \sum_{j=0}^t \zeta_{ij} + v_{it}^T + m_{it}$. X_{it}^1 is a subset of the exogenous variables available at time t , denoted by X_{it} . Wages are only observed when a binary selection indicator -in this case the participation indicator P_{it} - is equal to 1. The selection process is determined by the following probit equation:

$$\begin{aligned} P_{it} &= I\{X_i \beta^2 + \nu_{it}^2 > 0\}, \\ \nu_{it}^2 | X_i &\sim N(0, \sigma_{\nu_{it}^2}^2), \end{aligned} \quad (5.4)$$

where $X_i = \{X_{it}\}_{t=1}^T$. Equation 5.4 is the reduced form of many usual selection processes, including dynamic selection mechanisms and selection equations featuring a fixed effect, such as the following:

$$P_{it} = I\{X_i \beta^2 + f_i^2 + u_{it}^2 > 0\},$$

where f_i^2 is a time-constant fixed effect, and f_i^2 and u_{it}^2 are jointly normally distributed.

The following assumptions on u_{it}^1 , v_{i0}^P and ν_{it}^2 are needed in order to correct for selection in the presence of arbitrary serial dependence between u_{it}^1 and ν_{it}^2 :

ASSUMPTION W.1: The selection equation is given by 5.4.

ASSUMPTION W.2: $E(u_{it}^1 | X_i, \nu_{it}^2) = E(u_{it}^1 | \nu_{it}^2) = \mathbf{L}(u_{it}^1 | \nu_{it}^2) = \rho \nu_{it}^2$ for $t = 1, \dots, T$

ASSUMPTION W.3: $E(v_{i0}^P | X_i, \nu_{it}^2) = \mathbf{L}(v_{i0}^P | 1, X_i, \nu_{it}^2)$,

where $\mathbf{L}$ denotes the linear projection operator.

The conditional mean independence assumption W.2 always holds if (u_{it}^1, ν_{it}^2) are independent of X_i , a standard assumption in the selection context. A sufficient condition for the last equality in W.2 to hold is joint normality of (u_{it}^1, ν_{it}^2) . Notice that W.2 imposes no restriction on the temporal dependence of $\{u_{it}^1\}$, or on how u_{it}^1 relates to ν_{it}^2 , $r \neq t$.

Assumption W.3 is analogous to Chamberlain's approach to unobserved component probit models. It is clearly more restrictive than W.2, but it still allows for arbitrary serial dependence in (u_{it}^1, ν_{it}^2) . A sufficient -though not necessary- condition for assumption W.3 to hold is bivariate normality of (v_{i0}^P, ν_{it}^2) conditional on X_i , with constant variance matrix and expectation linear in X_i .

Assumption W.3 allows to write

$$E(v_{i0}^P | X_i, \nu_{it}^2) = X_i \Pi + \phi_t \nu_{it}^2. \quad (5.5)$$

As mentioned above, this linearisation of the fixed effect as a function of observable variables provides a method to assign initial conditions to all individuals in the sample regardless of whether their wage is

observed and, eventually, to simulated individuals. (5.5) implies

$$E(\ln w_{it} \mid X_i, \nu_{it}^2) = X_{it}^1 \beta^1 + X_i \Pi + \gamma_t \nu_{it}^2, \quad (5.6)$$

where $\gamma_t = \rho + \phi_t$. Conditioning on $P_{it} = 1$ yields

$$E(\ln w_{it} \mid X_i, P_{it} = 1) = X_{it}^1 \beta^1 + X_i \Pi + \gamma_t \lambda(X_i \beta^2). \quad (5.7)$$

A crucial implication of assumptions W.1 to W.3 is that the Π coefficients are constant across time²⁴. This ensures that β^1 is separately identified from Π in (5.7) for time-variant variables. For time-invariant variables, β^1 is not separable from Π .

Under assumptions W.1 to W.3, consistent estimates of β^1 can be obtained by first estimating equation 5.4 by pooled probit across i and t , generating the estimated inverse Mills ratios $\hat{\lambda}_{it}$, and then running a pooled OLS regression of $\ln w_{it}$ on X_{it}^1 , X_i , $\hat{\lambda}_{it}$ for the selected sample.

The first step in applying this procedure to the data is estimating the reduced-form selection process in (5.4). The vector of regressors X_i contains observations of the X_{it} variables in all time periods. Thus those observations that are missing in at least one period must be dropped, restricting the sample to the subset of observations that form a balanced panel. In order to keep the whole sample and conserve on degrees of freedom, I adopt a Mundlak (1978) approach and replace X_i with $(X_{it}, \bar{X}_i)$, where $\bar{X}_i = \sum_{t=1}^T X_{it}$. Equation 5.4 becomes:

$$\begin{aligned} P_{it} &= I\{\mu_0 + \bar{X}_i \mu_1 + X_{it} \mu_2 + \nu_{it}^2 > 0\}, \\ \nu_{it}^2 \mid X_i &\sim N(0, \sigma_{\nu_{it}^2}^2), \end{aligned} \quad (5.8)$$

which can be estimated by pooled probit. Results for the sample of married men aged 51 to 75 are provided in the first column of table 5. The exclusion restrictions used are an indicator of whether the wife is older than 62, which is the early Social Security retirement age, and should hence make her more likely to have retired; and an indicator of whether the wife has a DB pension, which increases her likelihood to retire early. The choice of exclusion restrictions is based on the evidence presented in section 2.2, which shows that men are more likely to be retired if their wife is retired as well. The assumption is made that the indicators of whether the wife has a DB pension and whether she is past age 62 can only affect participation directly, and not through the fixed effect in the selection equation. Wealth is also used as an exclusion restriction, as the assumption is made that it does not affect the wage directly. However, unlike the indicators of the wife's age and pension, wealth also enters the linearisation of the fixed effect.

Table 5 shows that both having a wife who is older than 62 or who has a DB pension decrease men's probability of participation, as does wealth. The probability of selection decreases for men when they are in bad health and as their age increases, with significant drops when reaching the Social Security early and normal retirement ages of 62 and 65. Finally, more educated men are more likely to be observed working.

Results from the estimation of the selection processes for married women aged 51 to 75 are provided in the second column of table 5. The same exclusion restrictions as for men are used.

Women whose husband is above the Social Security early retirement age of 62 or whose husband has a DB pension are less likely to be working. Wealth, however, is not found to have a significant effect on participation. Finally, women's probability of participation decreases with age, with a significant drop at

²⁴See Wooldridge (1995).

age 62, and when they are in bad health. As it was the case for men, more educated women are more likely to be observed working.

The estimates for the selection process are used to generate estimated inverse Mills ratios for each gender. These can now be used to estimate equation 5.7. As with the selection process, I make the simplifying assumption that the dependence of the fixed effect on the observed variables can be captured through a time average. Hence the estimation equation becomes:

$$E(\ln w_{it} \mid X_i, P_{it} = 1) = X_{it}^1 \beta^1 + \bar{X}_i \pi + \gamma \lambda (\hat{\mu}_0 + \bar{X}_i \hat{\mu}_1 + X_{it} \hat{\mu}_2). \quad (5.9)$$

An important point regarding equation 5.9 is that β^1 is separately identified from π for time-variant variables only. For time invariant regressors, such as the education dummies, β^1 and π are not separable. In what follows I assume that the effect of the education variables operates through the permanent component of wages, and hence they are part of the initial condition.

Estimates for men and women are reported in table 6. Notice that the time averages of wealth is included among the regressors, since wealth enters the reduced-form expression for the fixed effect. In order to allow maximum flexibility, the inverse Mill's ratios are interacted with the dummies indicating whether the individual has no pension and whether she has a DC pension.

The first noticeable point is that the coefficients on the inverse Mill's ratios are never significant. The significance of these coefficients can be interpreted as a test for selection²⁵. Hence, we fail to reject the absence of selection bias for both men and women who have either a DC pension or no private pension.

There is no significant effect of bad health on wages for either men or women. For men, wages increase with age at a decreasing rate. For women, there is evidence of a similar trend with age, but it is not significant.

Regarding the fixed effect, it is positively correlated with education for men with no pension, and even more strongly for men with a DC pension. The fixed effect of women is higher for those who have gone to college. Being a high-school graduate increases the fixed effect of women who have a DC pension, and has no effect for women with no pension. The time average of wealth is a strong indicator of higher fixed effects for both men and women, while the time average of the dummy measuring whether the individual is in bad health is an indicator of lower fixed effects.

I now follow the procedure described in Meghir and Pistaferri (2004) to estimate the variance of the shock to the persistent component of wages. Using the robust estimates of β^1 , I generate estimates of the error term u_{it} in equation 5.1:

$$\hat{u}_{it} = \ln w_{it} - X_{it}^1 \hat{\beta}^1$$

The unexplained component of the rate of growth of wages is defined as:

$$\begin{aligned} g_{it} &= \Delta_2 u_{it} = \Delta_2 v_{it}^P + \Delta_2 v_{it}^T + \Delta_2 m_{it} = \\ &= \zeta_{it} + \zeta_{it-1} + \Delta_2 v_{it}^T + \Delta_2 m_{it}, \end{aligned}$$

where Δ_2 are two-year differences ($\Delta x_t = x_t - x_{t-2}$), used here because this is the time gap between available survey waves.

An estimate of g_{it} is obtained as $\hat{g}_{it} = \hat{u}_{it} - \hat{u}_{it-2}$. I first use the estimated unconditional autocovariances of g to identify the order the MA process and find that $q = 1$ (see appendix for details). Exploiting the

²⁵Wooldridge (2005) points out, however, that this test makes stronger assumptions under H_0 than are needed for the fixed effect estimator to be consistent.

structure of the MA process, it follows that the moment condition that identifies the variance of the permanent shock is:

$$E \left[g_{it} \left(\sum_{j=-1}^{j=1} g_{it+2j} \right) \right] = E(\zeta_{it}^2) + E(\zeta_{it-1}^2) = 2E(\zeta_i^2), \quad (5.10)$$

where the last equality follows assuming that the variance of the permanent component is constant over time. An estimate of $E(\zeta_i^2)$ is obtained replacing 5.10 with its sample equivalent. The standard error is computed using non-parametric bootstrap. Both are shown in table 1 below.

Table 1: Estimated variance of the permanent shock.

	Men	Women
$\hat{\sigma}_{\zeta}^2$	0.0378** (0.0041)	0.0173** (0.0044)

5.3 Health Expenditures

Health costs are assumed to follow the process in equation (3.4). The function $hc(\cdot)$ is approximated as a polynomial on individual age and health status.

Estimates of the polynomial coefficients are used in the dynamic program to predict future health costs of sample individuals. In order to appropriately capture the evolution of health cost risk as individuals age, the estimation sample must include observations of health costs for individuals of all ages. However, individuals in the HRS cohort are rarely observed beyond age 75. Core²⁶ HRS individuals were aged 51 to 61 when interviewed for the first time in 1992, and 63 to 73 in 2004, the last year of the panel. The only way for an individual beyond age 75 to be present in the HRS cohort is if he or she is married to a younger, age-eligible spouse. Hence, in order to obtain estimates that are also representative of older individuals, I combine data from the HRS cohort with data from the AHEAD (Aging and Health Dynamics) cohort, which samples individuals born before 1923 and their spouses.

As explained above, I have selected from the HRS cohort only those couples where both spouses have either no private pension or a DC pension. I have no information on private pension type for AHEAD individuals. So as to obtain a comparable sample, I exclude AHEAD individuals who report to be receiving any private pension income. This eliminates from the sample those AHEAD couples where one of the spouses is receiving a DC pension. However, given the relatively recent expansion of this type of pensions, the number of such couples is unlikely to be large. On the other hand, the sample would include any AHEAD couple who were expecting to receive a DB pension in the future, but were not doing so during the sample years. Again, since AHEAD individuals are observed until they are aged 81 or over, this is unlikely to be a concern for a significant number of observations.

There are several respects in which individuals from the HRS and AHEAD cohorts may differ. An important one is pension coverage, which is more widespread in the younger cohort, especially for women. This is partly related to another important difference across cohorts: the accumulated work experience of younger women is noticeably higher than that of their older counterparts.

In the estimation of equation (.), the presence of cohort effects is accounted for by interacting all terms of the hc polynomial with cohort dummies. Time dummies are also included in the regression, in order to account for differences in the way out-of-pocket health costs are measured across waves.

²⁶I denote as *core* individuals those who belong to a cohort based on their year of birth, and not just because their are married to an age-eligible spouse.

Finally, I explore the possibility of changes in health costs after age 65, when all individuals become entitled to the public health insurance programme Medicare, by adding to the regressions a dummy indicating whether an individual is 65 years of age or older, and its interactions with other variables of interest.

Results of the separate regressions for men and women are shown in table 7.

5.4 Remaining Calibrations of Exogenous Parameters

Gender-specific health transition probabilities, conditional on health status on the previous period, are calibrated to those observed in the data.

I take unconditional survival probabilities from the life table used by the US Social Security Administration²⁷ for the cohort born between 1930 and 1939 -the HRS sample includes individuals born from 1931 to 1941 and their spouses-. Survival probabilities conditional on health status are obtained applying Bayes' rule, separately for men and women.

$$prob(survival_t | M_{t-1} = good) = \frac{prob(M_{t-1} = good | survival_t)}{prob(M_{t-1} = good)} \times prob(survival_t),$$

where all probabilities except for the unconditional survival probability are calibrated from the data.

The means-tested consumption floor provided by transfers is set to \$633 per household, per month. This is the (means-tested) amount of Supplemental Social Security Income that a couple aged 65 or older and on income support would have received in 1992.

6 Results from Simulations

This section presents results obtained using the parameters estimated in section 5 to describe the evolution for the exogenous variables and assuming values for the preference parameters taken from studies of individual behaviour -mainly French (2005) and French and Jones (2004). These parameters will not necessarily be those which minimise the distance between empirical moments and those simulated using this model. They are used here to solve the numerical model in order to illustrate its dynamics and the type of behavioural patterns it is able to capture, rather than to make quantitative statements relating to agents' behaviour or welfare.

6.1 Effects of observed state variables in retirement behaviour

In order to illustrate the effect of changes in state variables on couples' behaviour, and thus how the heterogeneity in states allows to model the dispersion in observed behaviours, I start by defining a baseline set of couples.

Baseline couples have some of their states fixed at the following values: the husband is 64 years old and the wife is 62; the husband has a wage equal to the mean for married men aged 60 to 65 (i.e. \$11.83 per hour); the wife has a wage equal to the mean for married women aged 60 to 65 (i.e. \$8.25); both spouses are in good health; they have not started claiming Social Security benefits; and they have no assets at the beginning of the period, which implies that they rely heavily on their Social Security pension to finance retirement. Moreover, there are no complementarities in spouses' leisure, so that the utility each spouse derives of their own leisure is independent of the partner's work status. This is equivalent to setting the parameters α_2^m and α_2^f equal to zero.

²⁷"Life Tables for the United States Social Security Area 1900-2100". Social Security Administration. Office of the Chief Actuary. August 2005.

Graph 8 plots optimal household participation behaviour for baseline couples as a function of husband and wife's accumulated earnings. The x axis in figure measures deciles of the maximum earnings a woman may have accumulated by age 62, with the left-most corner corresponding to no accumulated earnings -and thus no Social Security pension on her own behalf if the woman retires- and the right-most corner corresponding to maximum accumulated earnings -and maximum pension when the woman retires. The y-axis measures deciles of the maximum earnings a man may have accumulated by age 64.

The different-coloured areas in the graph represent the optimal participation choices of the couples in the current period. The light-brown area on the bottom left corner of the graph corresponds to points where accumulated earnings of the spouses are relatively low. In this set of points, the optimal discrete choice is for both partners to work. As we move towards the right on the x axis, accumulated earnings of the wife increase. The blue area covers those points where the discrete option yielding the highest value is for the husband to work and for the wife to retire. Notice that the minimum amount of accumulated earnings that makes it optimal for the wife to retire is a function of the husband's accumulated earnings. This is because as the husband's accumulated earnings increase, the household anticipates a higher pension for the husband in future periods, and thus more means to finance future consumption. Hence the wife's income in the current period -either from work or, if she retires, from her Social Security pension-, which would contribute to finance current consumption but also to increase savings, is less needed, and she can afford to retire with a lower replacement rate.

The red area on the upper-right corner of the graph, which corresponds to combinations of high accumulated earnings for both spouses, covers points where it is optimal for both spouses to retire.

The most interesting feature of figure 8, however, is the red area on the upper-left corner of the graph. This corresponds to points where the husband has more than 70% of the maximum accumulated earnings, while the wife has at most 10% of the maximum accumulated earnings for a woman her age. For combinations of accumulated earnings in this region, it is optimal for both spouses to retire. The rationale behind this area of the graph is the retirement spouse benefit provision. This Social Security provision allows one spouse -in this case the wife- to claim a pension equal to 50% of that of the husband as long as he is retired. Hence in this part of the graph, where the woman has almost no accumulated earnings, and therefore no right to a pension on her own behalf if she retires, it is optimal for the husband to retire together with the wife in order to provide her with a pension. It is interesting to see that, where the wife's accumulated earnings are slightly higher, it is no longer optimal for the husband to retire with her. It is only when her pension entitlement is sufficiently low that both of them retire together.

Figure shows the two-dimensional value function for the baseline couples as a function of husband and wife's accumulated earnings. As expected, the lowest value corresponds to couples where both spouses have no accumulated earnings, while the highest value corresponds to couples where both spouses have the maximum accumulated earnings for their age.

The next step is to change some of the baseline states so as to show how they affect optimal choices. Figure 10 shows optimal participation behaviour for couples with all states equal to those at baseline except initial assets, which in this case are positive and equal to \$40,000. Since these couples are richer than those at baseline, they do not rely only of their Social Security pension to finance retirement, and hence can afford to stop working with lower amounts of accumulated earnings. Accordingly, the graph shows that the light-brown area, which corresponds to combinations of accumulated earnings that make it optimal for both spouses to work, is smaller than before, while the red areas, where it is optimal for both spouses to retire, are bigger. In this graph, a red area appears in the bottom-right corner, covering a set of points where it is optimal for both spouses to retire together, with the man benefiting from the dependent spouse provision.

We can see from figure 11 that the value function for couples with positive assets dominates that of

couples at baseline -who have no assets.

Figure 12 shows optimal participation choices for couples that differ from those at baseline in that the husband is in bad health. Being in bad health diminishes the number of hours available for the husband to divide between work and leisure. Thus working as much as before would considerably reduce the amount of leisure that he enjoys. In this situation, couples will find it optimal for the husband to reduce the number of hours worked or to retire. Indeed, figure 12 shows that the set of points where it is optimal for both spouses to retire has increased considerably with respect to the baseline. Having to reduce the amount of hours worked has decreased the employment income husbands can provide. Hence it is not worthwhile for them to work for values of accumulated earnings that did not lead to retirement at baseline.

Figure shows that couples where the husband is in bad health attain lower values than those at baseline for every possible combination of husband and wife's accumulated earnings.

Finally, figure shows optimal participation for spouses whose wages have been exchanged with respect to baseline: now the wife has a wage of \$11.83 per hour and the husband has a wage equal to \$8.25 per hour. In this figure, a black area appears in the upper-left corner which corresponds to points where it is optimal only for the man to retire. Notice as well that, even though wages have been exchanged between spouses, the resulting figure is not a mirror image of the baseline one. In particular, there are still points in this case where it is optimal for the wife to retire while the husband is working, despite her wage being higher. This is due to the fixed cost of work, which I take from previous studies. This cost has been estimated to be much higher for women than for men. The consequence is that, for the same wage, women will choose to work less hours than men, and will be more likely to retire than men. Therefore, they will provide less employment income to the household than a man would do with the same wage. This is the reason why the value function for these couples is lower everywhere than that for couples at baseline, as can be seen in figure .

6.2 The role of non-separabilities

So far I have considered spouses whose valuation of own leisure is independent of their partners' participation status. All the incentives to retire jointly for these couples were financial, such as those provided by the dependent spouse benefit provision. The next graph shows how the existence of non-separabilities can increase spouses' tendency to retire together. Figure 16 shows the participation decisions of couples where both spouses enjoy their leisure more if their partner is retired. In particular, all states are equal to baseline except for the complementarity parameters α_2^m and α_2^f , which have been set equal to 0.15. The result has been an increase of the areas where both spouses find it optimal to retire together. It is interesting to note how the existence of non-separabilities exacerbates the effect of the financial incentives to joint retirement provided by the dependent spouse provision. The number of points where it is optimal for the man to retire with the wife so that she can benefit from a pension based on his entitlement has increased with respect to the baseline. Moreover, there are now some combinations of accumulated earnings that make it optimal for the wife to retire with the husband, so that he can benefit from a pension as a dependent spouse.

The type of non-separabilities assumed in this experiment rises each spouse's utility whenever their partner is retired. For spouses who choose to retire in the current period, their current utility is higher than it would have been at baseline. And for all couples, independently of their participation status in the current period, the value of expected future utility is higher than it would have been at baseline, since they anticipate retiring before they die, and benefiting from the utility derived from future shared leisure. This implies that the value function for these couples is higher than that at baseline for every combination of accumulated earnings, as can be seen in figure 17.

6.3 Policy Analysis Application

One of the advantages of structural models like this one is related to the analysis of policy reforms. These types of models provide a complete description of agents' behaviour. Therefore, once the change in individual incentives caused by a policy reform is identified, they can be used to predict how it will affect agents' behaviour -even that of agents not directly targeted by the policy-. Structural models allow to draw both qualitative and quantitative conclusions about the likely effects of a policy change that would normally not be identified from a reduced-form analysis.

In the rest of this section, I illustrate how this model can be used to predict the effects of a policy reform involving a reduction of the dependent spouse benefit. The spousal benefit gives spouses the right to choose between a pension based on their own entitlement and one based on 50 per cent of their partner's entitlement. It was initially intended to cover nonworking, dependent wives, at a time when households where only the husband worked were the norm. Lately, doubts have been raised about the wisdom of maintaining the program, given the increasing participation rate of women and the likely disincentives to work past early retirement age that it involves for some spouses. A report from the Congressional Budget Office published in February 2007 argued for a reduction of the dependent spouse benefit from 50 percent to 33 percent, under the rationale that the reform would strengthen the connection between taxes paid and benefits received: "[The dependent spouse benefit rules] weaken the link between the Social Security taxes that are paid and the benefits that are received. Relative to Social Security taxes paid, a one-earner couple currently receives substantially higher benefits than either a single worker who has the same earnings history or a two-earner married couple."²⁸

In the next graph I simulate a reduction of the spouse benefit from 50% to 33%. The effect with respect to baseline can be seen in figure 18. The proportion of spouses retiring together when the husband has high accumulated earnings and the wife has no or very low accumulated earnings has decreased. On the other hand, the proportion of spouses choosing to retire together when both of them have high accumulated earnings remains stable, since their incentives for joint retirement were unrelated to the spousal benefit.

Figure 19 shows that the value function for spouses with maximum earnings has not changed with respect to baseline. This is because they already have the maximum value of accumulated earnings, and independently of their participation behaviour in the future they will never benefit from the spouse benefit. The value function for couples where the husband has maximum accumulated earnings and the wife has no accumulated earnings has decreased with respect to baseline, as the wife is now receiving a Social Security pension based on 33 percent of her husband's entitlement, instead of 50 percent. Finally, the value function for couples where the husband has no accumulated earnings is also lower than at baseline. Notice that none of these couples was benefiting from the spousal benefit at baseline. However, depending on both spouses' participation choices, they may benefit from the spousal benefit in the future. Hence their expected future utility is lower than it would have been at baseline.

This experiment shows how the model presented in section 3 may be used to predict the effect of a policy change. According to the results presented in 18 we can predict that, even though the spousal benefit is a policy designed to cover women with low accumulated earnings, decreasing it from 50 to 33% would have no impact on baseline women's participation behaviour. The behaviour of their husbands, instead, would be affected: since the reform diminishes their incentives to retire in order to provide a pension for their wife, we observe less of them retiring after the reform, and hence an increase in baseline men's participation.

²⁸Congressional Budget Office. Report to the House and Senate Committees on the Budget. February 2007. <http://www.cbo.gov/ftpdoc.cfm?index=7821>

7 Conclusions

In this paper, I present a stochastic dynamic model of older couples' choices concerning consumption and savings, participation and number of hours worked.

I present results of the estimation of the laws of motion of the exogenous variables. The wage equation is estimated separately for men and women, controlling for selection in both cases, since retirement decisions are likely to be correlated with individuals' wage shocks. The procedure followed to correct for selectivity bias, described in Wooldridge (1995), accounts for the presence of a persistent component in the wage equation. Then, I follow Meghir and Pistaferri (2004) to estimate the variance of the shock to the persistent error component of wages. The estimate of the variance, together with the estimates from the wage equation and the linearisation of the persistent component are used to input wages to simulated individuals.

The process for health costs is estimated merging data from the HRS cohort with data from older individuals, in order to be able to predict expected health costs in very old age. I control for the presence of cohort effects that would affect differently the medical care costs of individuals from the two samples.

A growing body of literature in recent years (see Heckman, Lochner and Cossa (2003), Adda, Dustman, Meghir, and Robin (2005) and Meghir (2006)) has made a case for the need of structural models in the analysis of policy reforms. I show, using results from simulations, how this model can capture the direct and indirect effects that would result from a reduction of the dependent spouse benefit in the Social Security. The spousal benefit is a policy designed to cover women with very short working histories. The simulations suggest, however, that the participation behaviour of women would not be as affected by the change in policy as that of men, who become much less likely to retire once their incentives to do so in order to provide their wife with a pension are reduced.

As part of my current research, I am now working on the estimation of the model's preference parameters. The objective of this paper is to show the use of the model once the preference parameters estimates have been obtained. When this has been achieved, it will be possible to make meaningful quantitative statements about the effects of a policy change like the one discussed above. The model will allow not only to analyse the process through which a change in incentives affects agents' behaviour, but also to quantify resulting changes in welfare, which are of the utmost importance when analysing the effects of a policy on older individuals.

References

- [1] Abowd, J. M. and D. Card. 1989. "On the Covariance Structure of Earnings and Hours Changes". *Econometrica*. Vol 57. Issue 2, pp 411-445.
- [2] Abramowitz, M. and Stegun, I. A. (eds.) 1964. *Handbook of Mathematical Functions*. National Bureau of Standards, Washington DC.
- [3] Adda, J., C. Dustmann, Costas Meghir and Jean-Marc Robin. 2005. "Career Progression and Formal versus on the Job Training". Mimeo, University College London (UCL)
- [4] Aguirregabiria, V. and P. Mira. 2002. "Swapping the Nested Fixed Point Algorithm: A Class of Estimators for Discrete Markov Decision Models". *Econometrica*. Vol. 70, pp. 1519-1543.
- [5] Banks, J., R. Blundell and M. Casanova Rivas. 2007. "The dynamics of retirement behavior in couples: Reduced-form evidence from England and the US". Mimeo, University College London (UCL).
- [6] Blau, D. 1994. "Labor Force Dynamics of Older Men". *Econometrica*. Vol 62. Issue 1, pp 117-156.
- [7] Blau, D. 1998. "Labor Force Dynamics of Older Married Couples". *Journal of Labor Economics*. Vol 16. Issue 3, pp 595-629.
- [8] Blau, D. 2004. "Retirement and Consumption in a Life Cycle Model". Manuscript. University of North Carolina at Chapel Hill.
- [9] Blau, D. and D. Gilleskie. 2006. "Health Insurance and Retirement of Married Couples". *Journal of Applied Econometrics*. Vol 21. Issue 7, pp 935-953.
- [10] Blundell, R. and T. MaCurdy. 1999. "Labor Supply: A Review of Alternative Approaches". In *Handbook of Labor Economics*. Vol. 3A, Chapter 23, pp 1560-1695. Ed. by O. Ashenfelter and R. Layard.
- [11] De Boor, C. 1979. "Efficient Computer Manipulation of Tensor Products". *ACM Transactions on Mathematical Software*. Vol. 5, no 2, pp. 173-182
- [12] De Boor, C. 2001. "A practical Guide to Splines". Revisited Edition. Springer-Verlag New York.
- [13] Burtless, G. 1990b. "Comment on M. Hurd's 'The Joint Retirement Decisions of Husbands and Wives'". In *Issues in the Economics of Aging*", ed. David A. Wise, Chicago: University of Chicago Press.
- [14] Cagetti, M. 2003. "Wealth Accumulation over the Life-Cycle and Precautionary Savings". *Journal of Business and Economic Statistics*. Vol 21, pp. 339-353.
- [15] Coile, C. 2004a. "Retirement Incentives and Couples' Retirement Decisions". *Topics in Economic Analysis & Policy*. Vol. 4, No. 1, Article 17.
- [16] Coile, C. 2004b. "Health Shocks and Couples' Labor Supply Decisions". NBER Working Paper 10810.
- [17] Coile, C. and J. Gruber. 2000. "Social Security and Retirement". NBER working paper 7830.
- [18] Coile, C. and J. Gruber. 2001. "Social Security Incentives for Retirement" in *Themes in the Economics of Aging*. Ed. by D. Wise. The University of Chicago Press.

- [19] Diamond, P. and J. Gruber. 1999. "Social Security and Retirement in the United States". In *Social Security and Retirement Around the World*. Ed. by J. Gruber and D. Wise. Chicago: University of Chicago Press.
- [20] Eckstein, Z. and K. Wolpin. 1989. "The Specification and Estimation of Dynamic Stochastic Discrete Choice Models". *Journal of Human Resources*. Vol 24, pp. 562-598
- [21] French, E. 2005. "The Effects of Health, Wealth, and Wages on Labour Supply and Retirement Behaviour". *Review of Economic Studies*. Vol. 72, issue 2, pp. 395-427.
- [French and Jones, 2004a] French, E. and J. B. Jones. 2004a. "On the Distribution and Dynamics of Health Care Costs". *Journal of Applied Econometrics*. Vol 19, issue 4, pp. 705-721.
- [22] French, E. and J. B. Jones. 2004b. "The effects of Health Insurance and Self-Insurance on Retirement Behavior". Federal Reserve Bank of Chicago Working Paper.
- [23] French, E., J. B. Jones and M. De Nardi. 2006. "Saving Motives After Retirement". Working paper.
- [24] Gilleskie, D. 1998. "A Dynamic Stochastic Model of Medical Care use and Work Absence". *Econometrica*. Vol 59, pp. 189-210.
- [25] Gourinchas, P.-O., and J. Parker. 2002. "Consumption over the life cycle". *Econometrica*. Vol. 70, no. 1, pp. 47-89.
- [26] Gustman, A. and T. Steinmeier. 1986. "A Structural Retirement Model". *Econometrica*. Vol. 54, No. 3, pp. 555-584.
- [27] Gustman, A. and T. Steinmeier. 1994. "Employer-Provided Health Insurance and Retirement Behavior". *Industrial and Labor Relations Review*. Vol. 48, No. 1 , pp. 124-140
- [28] Gustman, A. and T. Steinmeier. 2000. "Retirement in Dual-Career Families: A Structural Model". *Journal of Labor Economics*. Vol. 18, issue 3, pp. 503-545.
- [29] Hausman, J. and P. Ruud. 1984. "Family Labor Supply with Taxes". *American Economic Review*. Vol. 74. No. 2.
- [30] Heckman, J., L. Lochner, and R. Cossa (2003). "Learning-By-Doing versus On-the-Job Training: Using Variation Induced by the EITC to Distinguish Between Models of Skill Formation", in *Designing Inclusion: Tools to Raise Low-end Pay and Employment in Private Enterprise*. Ed by E. Phelps. Cambridge: Cambridge University Press
- [31] Hotz, J. and R. A. Miller. 1993. "Conditional Choice Probabilities and Estimation of Dynamic Models". *Review of Economic Studies*. Vol 60, pp. 497-529
- [32] Hubbard, G., J. Skinner and S. Zeldes. 1995. "Precautinary Savings and Social Insurance". *Journal of Political Economy*. Vol. 103, issue 2, pp. 360-399.
- [33] Hurd, M. 1990a. "Research on the Elderly: Economic Status, Retirement, and Consumption and Saving". *Journal of Economic Literature*. Vol 28. Issue 2 pp 565-637.
- [34] Hurd, M. 1990b. "The Joint Retirement Decisions of Husbands and Wives". In *Issues in the Economics of Aging*, ed. David A. Wise, Chicago: University of Chicago Press.
- [35] Hurd, M. and S. Rohwedder. 2003. "The Retirement-Consumption Puzzle: Anticipated and Actual Declines in Spending at Retirement". NBER Working Paper No 9586.

- [36] Internal Revenue Service. Department of the Treasury. 1992. "Estimated Tax for Individuals". IRS.gov
- [37] Judd, K. 1998. "Numerical Methods in Economics". The MIT Press, Cambridge, MA.
- [38] Karoly, L., N. Maestas and J. Zissimopoulos. 2006. "Retirement Transitions of the Self-Employed in the United States and England". Manuscript, RAND.
- [39] Keane, M. and K. Wolpin. 1994. "The Solution and Estimation of Discrete Choice Dynamic Programming Models by Simulation and Estimation: Monte Carlo Evidence". *The Review of Economics and Statistics*. Vol 76, No 4, pp. 648-672
- [40] Keane, M. and K. Wolpin. 1997. "The Career Decisions of Young Men". *Journal of Political Economy*. Vol 105, pp. 473-522.
- [41] Van der Klaauw, W. and K. Wolpin. 2006. "Social Security, Pensions and the Savings and Retirement Behavior of Households". Manuscript, University of North Carolina at Chapel Hill.
- [42] Low, H. 1999. "Self-Insurance and Unemployment Benefit in a Life-Cycle Model of Labour Supply and Savings". IFS Working Paper W99/24.
- [43] Lumsdaine, R., R. Stock and D. Wise. 1994a. "Pension Plan Provisions and Retirement: Men & Women, Medicare, and Models," in *Studies in the Economics of Aging*. Ed. by D. Wise. University of Chicago Press.
- [44] Lumsdaine, R., R. Stock and D. Wise. 1994b. "Retirement Incentives: The Interaction Between Employer-Provided Pensions, Social Security, and Retiree Health Benefits". NBER working paper number 4613.
- [45] MaCurdy, T. 1982. "The Use of Time Series Processes to Model the Error Structure of Earnings in a Longitudinal Data Analysis". *Journal of Econometrics*. Vol 18, pp. 82-114
- [46] Meghir, C. 2006. "Dynamic Models for Policy Evaluation". Mimeo, University College, London (UCL) and IFS.
- [47] Meghir, C. and L. Pistaferri. 2004. "Income Variance Dynamics and Heterogeneity". *Econometrica*. Vol 75, Issue 1, pp. 1-32
- [48] Michaud, P.-C. 2004. "Labor Force participation dynamics and Social Security Claiming Decisions of Elderly Couples". CentER Tilburg University and IZA .
- [49] Michaud, P.-C. and F. Vermeulen. 2004. "A Collective Retirement Model: Identification and Estimation in the Presence of Externalities". IZA Working Paper.
- [50] Miller, R. 1997. "Estimating Models of Dynamic Optimization with Microeconomic Data". In *Handbook of Applied Econometrics: Microeconometrics*. Vol 2, pp. 246-299. Ed. by H. Pesaran and P. Smidh. Oxford, U.K.: Blackwell.
- [51] Mitchell, O., J. Olson and T. Steinmeier. 1996. "Construction of the Earnings and Benefits File (EBF) for Use With the Health and Retirement Study" HRS/AHEAD Documentation Report DR-001. Institute for Social Research, University of Michigan.
- [52] Rust, J. 1987. "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher". *Econometrica*. Vol 55, No 5, pp. 999-1033

- [53] Rust, J. 1988. "Maximum Likelihood Estimation of Discrete Control Processes". *SIAM Journal of Control and Optimization*. Vol 26, No 5, pp. 1006-1024
- [54] Rust, J. 1994. "Structural Estimation of Markov Decision Processes", in *Handbook of Econometrics*. Vol 4. Ed. by R. E. Engle and D. McFadden. Amsterdam: North Holland.
- [55] Rust, J. and C Phelan. 1997. "How Social Security and Medicare Affect Retirement Behavior in a World of Incomplete Markets". *Econometrica*. Vol. 65, issue 4, pp 781-831.
- [56] Scholz, John Karl, Ananth Seshadri and Surachai Khitatrakun. 2004. "Are Americans Saving Optimally for Retirement?". NBER Working Papers 10260. National Bureau of Economic Research, Inc.
- [57] Stock, James H. and David Wise. 1990. "Pensions, the Option Value of Work and Retirement". *Econometrica*. Vol. 58, issue 5, pp. 1151-1180.
- [58] U.S. Social Security Administration. 2005. "Annual Statistical Supplement to the Social Security Bulletin". <http://www.ssa.gov/policy/docs/statcomps/supplement/2005/supplement05.pdf>
- [59] U.S. Social Security Administration. "Social Security Online". <http://www.ssa.gov>
- [60] Wooldridge, J. M. 1995. "Selection Correction for Panel Data Models Under Conditional Mean Independence Assumptions". *Journal of Econometrics*. Vol. 68, pp. 115-132.
- [61] Wooldridge, J. M. 2002. "Econometric Analysis of Cross Section and Panel Data". The MIT Press, Cambridge, MA.

Appendix A: Mathematical Appendix

Computation of the integral with respect to out-of-pocket medical expenditures.

In order to solve period t 's problem, we need an approximation to the expected value of V_{t+1} . This expected value is taken with respect to health status in $t + 1$, survival into $t + 1$, health costs in period t ²⁹ and, when the husband is younger than 75 and/or the wife younger than 70, wages in period $t + 1$. In this section I describe the steps involved in the computation of the expected value with respect to health costs.

$$E_{hc^m, hc^f} \widehat{V}_{t+1}(z_{t+1}, x_{t+1} \mid z_t, x_t) \quad (7.1)$$

Recall that the probability that health costs are positive, and the logarithm of health costs have been modelled as follows:

$$p(hc_{it} > 0) = p_{it}^1 = X_{it}\beta_1 + \psi_{it}^1$$

$$\begin{aligned} \ln hc_{it} &= X_{it}\beta_2 + \psi_{it}^2, \\ \psi_{it}^2 &\sim N(0, \sigma_{\psi^2}^2) \end{aligned}$$

Hence the positive health costs are lognormally distributed:

$$hc_{it} \mid X_{it} \sim \log N(X_{it}\widehat{\beta}^2, \widehat{\sigma}_{\psi^2}^2)$$

Omitting the conditioning on the state variables that are not relevant to the solution of this integral, (A1) can be re-written as follows:

$$\begin{aligned} E_{hc^m, hc^f} \widehat{V}_{t+1}(z_{t+1}, x_{t+1} \mid z_t, x_t) &= \\ &+ (1 - p_{it}^{1,m}) \times (1 - p_{it}^{1,f}) \times \widehat{V}_{t+1}(hc_{it}^m = 0, hc_{it}^f = 0) + \\ &+ (1 - p_{it}^{1,m}) \times p_{it}^{1,f} \times \int_0^\infty \widehat{V}_{t+1}(hc_{it}^m = 0, hc_{it}^f) f(hc_{it}^f \mid X_{it}) dhc^f + \\ &+ p_{it}^{1,m} \times (1 - p_{it}^{1,f}) \times \int_0^\infty \widehat{V}_{t+1}(hc_{it}^m, hc_{it}^f = 0) f(hc_{it}^m \mid X_{it}) dhc^m + \\ &+ p_{it}^{1,m} \times p_{it}^{1,f} \times \int_0^\infty \int_0^\infty \widehat{V}_{t+1}(hc_{it}^m, hc_{it}^f) f(hc_{it}^m \mid X_{it}) f(hc_{it}^f \mid X_{it}) dhc^m dhc^f. \end{aligned}$$

Below I describe in detail the computation of the integral with respect to the husband's health costs. The value of the integral with respect to the wife's health costs is computed in a symmetric way, while the double integral is solved using a two-dimensional Gauss-Hermite integration rule.

Define K as:

$$K \equiv \int_0^{+\infty} \widehat{V}_{t+1}(hc_{it}^m) f(hc_{it}^m \mid X_{it}) dhc^m$$

Since hc_{it}^m is lognormally distributed,

²⁹It is assumed that the realisation of the medical cost draw happens at the end of the period, after households have made their consumption and work decisions.

$$K = \int_0^{+\infty} \widehat{V}_{t+1}(hc_t^m) \frac{1}{hc_{it}^m \widehat{\sigma}_{\psi^2}^m (2\pi)^{1/2}} \exp \left\{ - \left(\frac{\ln hc_t^m - X_{it} \widehat{\beta}_2^m}{2^{1/2} \widehat{\sigma}_{\psi^2}^m} \right)^2 \right\} dhc^m$$

Using the following change of variable,

$$z_{it} = \frac{\ln hc_{it}^m - X_{it} \widehat{\beta}_2^m}{2^{1/2} \widehat{\sigma}_{\psi^2}^m},$$

yields

$$K \equiv \int_{-\infty}^{+\infty} \widehat{V}_{t+1} \left(\exp \{ 2^{\frac{1}{2}} \widehat{\sigma}_{\psi^2}^m z_{it} + X_{it} \widehat{\beta}_m^2 \} \right) \frac{1}{\pi^{1/2}} \exp \{ -z_{it}^2 \} dz_1$$

The value of K is approximated using Gauss-Hermite quadrature.

$$K \approx \frac{1}{\pi^{1/2}} \sum_{j=1}^P \widehat{V}_{t+1} \left(\exp \{ 2^{1/2} \widehat{\sigma}_{\psi^2}^m \xi_j + X_{it} \widehat{\beta}_m^2 \} \right) \omega_j,$$

where $\{\xi_j, \omega_j\}_{j=1}^P$ are the abscissae and weights of a one-dimensional Gauss-Hermite integration rule with P points, which can be found in standard references (e.g. Abramowitz and Stegun, 1964).

Appendix B: Social Security function

The steps followed by the SSA in computing individual benefits are the following: First, annual individual earnings are indexed to reflect the general earnings level in the indexing year. Then, average indexed monthly earnings (AIME) are computed as an average over the highest 35 earning years in the individual's career. Next, the primary insurance amount (PIA) is obtained applying a three-piece linear progressive schedule to AIME. Finally, individual benefit figures are obtained by adjusting the PIA according to retirement age: for individuals who claim benefits at the normal retirement age (NRA), the monthly benefit amount equal the PIA. For those who claim before NRA and after the early retirement age (ERA), the PIA is reduced by an actuarial factor of 5/9 of 1% per month, meaning that a worker claiming at the earliest possible age of 62, would receive a monthly benefit equal to 80% of his or her PIA. If claiming is delayed beyond NRA, benefits are increased by 3% per year of delay.

Individuals whose spouse is claiming benefits receive the higher amount between their own entitlement and 50% of their spouse's. This includes individuals who have not worked enough periods to qualify for entitlement in their own right. Surviving spouses can also claim benefits based on their partner's PIA. Provisions are in place regarding minimum and maximum benefit levels. Benefits are adjusted for increases in CPI after age 62.

Individual benefits

Benefits depend on indexed lifetime earnings. For each year of work, there is a maximum amount of earnings, from which payroll tax is deducted, which will contribute to the pension. E_t is the measure of lifetime earnings used in the model:

$$E_t \equiv \begin{cases} \sum_{j=0}^t \omega_j e_j^* & \text{if } t \leq R \\ \sum_{j=0}^R \omega_j e_j^* & \text{if } t > R \end{cases},$$

where $t = 0$ is the first year of earnings, R is the first year of receipt of Social Security benefits (subject to the restriction $R \geq 62$), ω is the weight used by the Social Security administration to index yearly earning, and e_t^* is defined as the minimum between yearly earnings $e_t = w_t \times h_t$ and maximum taxable earnings for that year, $e_t^{\max}$:

$$e_t^* = \min \{e_t, e_t^{\max}\}.$$

In order to avoid the need to keep track of every individual's whole earnings history, Average Indexed Monthly Earnings ($AIME_t$) are approximated as a function of E_t as follows:

$$AIME_t = \frac{E_t}{12 \times \max\{(t - 25), 35\}}$$

Full retirement entitlement, also known as Primary Insurance Amount (PIA_t) is obtained from $AIME$ according to the Social Security formula, re-scaled by the weight κ :

$$\begin{aligned} PIA_t = & \kappa_t [0.90 \times \min\{AIME_t, b_0\} + 0.32 \times \min\{\max\{AIME_t - b_0, 0\}, b_1 - b_0\} + \\ & + 0.15 \times \max\{AIME_t - b_1, 0\}], \end{aligned}$$

where κ_t is calibrated to give an individual retiring at each possible age with the maximum possible accumulated earnings exactly the same pension she would have been awarded under the 1992 Social Security rules³⁰. The bendpoints for the year 1992 are $b_0 = \$387$ and $b_1 = \$2,333$.

³⁰The Social Security benefit approximation just described yields a very accurate fit for individuals with the highest

Benefit entitlement is determined as a function of the PIA in the period in which an individual claims benefits, PIA_R .

$$ssb_t =$$

. Individuals who apply for benefits at age 65, receive their full PIA , that is benefits equal to PIA_R . Individuals who claim benefits before, have their b

Benefit entitlement is lower than full entitlement for individuals who claim benefits before age 65 (by 6.7% per year before 65). It is higher than full entitlement for individuals who delay claiming beyond age 65 (by 4% per year between 65 and 70).

Earnings test

Individuals who keep working after claiming benefits have their benefits taxed away for every dollar of earnings above a threshold, as described in section ?. Benefits lost throughout the earnings test during ages 62 to 65 yield an increase in future benefits equivalent to 6.7% per a whole year lost of benefits. Benefits lost throughout the earnings test during ages 66 to 70 yield an increase in future benefits equivalent to 4% per a whole year lost of benefits. This is accounted for in the benefit approximation when computing ssb_t .

Benefit claiming is an absorbing state. Moreover, benefits are a function of the PIA in the period in which an individual claims benefits, or PIA_R , which cannot be increased further, even if the individual keeps working after claiming.

Labour income

Spouses and Widowed individuals

In periods when both spouses are claiming benefits, the spouse with lowest PIA receives benefits ssb_t equal to the highest amount between her individual entitlement and entitlement based on 50% of her partner's PIA .

Individuals who become widowed can claim benefits based on their individual entitlement or that of their deceased partner.

Appendix C: Taxes

This section describes the tax function applied to couples' income in the model. Households pay federal and payroll taxes on income. Due to the great cross-sectional variation in state taxes, those are not accounted for here. I used the rates applying to married couples filing jointly. Also, I use the standard deduction, and hence do not allow households to defer medical expenses as an itemised deduction. The tax rates and exempt amounts used below are those corresponding to the year 1992.

Payroll Tax

The payroll tax is a proportional tax imposed on employees, which is used to finance the Social Security's OASDI programme and Medicare's hospital insurance programme. The social security tax rate for employees is 6.2% of earnings up to an upper limit of \$55,500. The Medicare tax rate for employees is 1.45% of earnings, and it is uncapped.

Defining individual annual earned income as

$$e_t^i \equiv w_t^i \times h_t^i, \quad \text{for } i = m, f,$$

each spouses' payroll tax contribution is given by:

$$\tau^P e_t^i = (0.062) \times \min\{\$55,500, e_t^i\} + 0.0145 \times e_t^i, \quad \text{for } i = m, f$$

possible pensions. Consequently, the highest weight used in the model is equal to 1.0517, and the lowest to 0.9915.

Federal Income Tax

The income tax is a progressive tax on labour and nonlabour income. The standard deduction for a married couple filling jointly was \$6,000 in 1992. Additionally, each spouse was entitled to a further deduction of \$700 if aged 65 or over.

Defining household income subject to federal income tax as

$$I_t \equiv (1 - \tau^P)e_t^m + (1 - \tau^P)e_t^f + rA_t,$$

generates the following level of post-income tax for a couple where both spouses are below age 65:

Table ??: Federal Income Tax Structure

Taxable income (I) (in dollars)	Post-tax Income (in dollars)	Marginal rate
0 – 6,000	Y	0.00
6,000 – 41,800	$6,000 + 0.85(Y-6,000)$	0.15
41,800 – 92,500	$36,430 + 0.72(Y-41,800)$	0.28
92,500 and over	$72,934 + 0.69(Y-92,500)$	0.31

Denoting the federal income tax structure by the vector τ^I , households' post-tax income is given by:

$$Y(rA_t, e_t^m, e_t^f, \tau^P, \tau^I) = (1 - \tau^I)I_t$$

Table 2: Descriptive statistics by pension type.

	Men			Women		
	None	DB only	DC only	None	DB only	DC only
Percentage population ¹	39.59	27.24	14.34	52.23	18.39	12.08
Employment						
% working at age 61	76.89	60.70	77.68	44.27	56.83	62.58
Average log wage in 1992 dollars ²	2.29	2.56	2.56	1.98	2.39	2.28
	(0.84)	(0.62)	(0.66)	(0.71)	(0.59)	(0.47)
Health insurance and costs before age 65						
% with retiree health coverage ³	10.67	32.25	27.22	7.13	31.75	19.43
% with tied health coverage ³	5.83	11.87	20.56	5.12	16.74	17.88
% with no empl. health insurance ³	51.23	22.39	27.78	66.92	33.91	44.59
Median health costs per period	293.34	301.17	304.03	417.23	416.52	365.96
Average health costs per period	961.11	876.20	777.21	1,202.55	1,065.71	919.15
	(3,883)	(2,706)	(1,816)	(4,502)	(3,415)	(2,404)
Health Status						
% in bad health	25.67	19.48	17.74	23.59	16.00	12.96
Total wealth in 1992 dollars ⁴						
Median	113,000	127,500	127,000	110,000	148,000	142,750
25 th percentile	41,000	65,550	63,648	44,000	75,500	72,925
Financial wealth in 1992 dollars ⁵						
Median	50,000	60,450	71,000	48,200	82,000	82,575
25 th percentile	8,200	19,000	19,200	9,800	24,400	25,000
Demographics						
Average age	62.81	62.73	61.68	58.36	58.51	57.68
% College education	35.30	36.04	41.30	27.54	46.06	38.59
% High-School graduates	33.32	40.97	36.02	42.84	40.46	47.09
N (couple-year observations)	10,031	6,903	3,634	13,234	4,661	3,060

¹ Percentages do not sum to 100 because of individuals for whom it is not possible to derive pension type.

² For participating individuals only.

³ Percentage measured with respect of individuals who report type of health insurance. Due to concerns about the measurement of health insurance type in the first two waves, reported figures correspond to wave 3.

⁴ Wealth measure includes housing but excludes private pension holdings.

⁵ Includes value of checking and saving accounts, stocks, mutual funds, investment trusts, CD's, Government bonds, Treasury bills and all other savings minus the value of debts such as credit card balances, life insurance policy loans or loans from relatives. It does not include housing wealth or private pension holdings.

Table 3: Participation status in period t of husbands who were working in $t - 1$.

	Husbands aged 50 to 60	
	% of husbands working in period t	N
wife working in $t - 1$, working in t	96.22	6,618
wife working in $t - 1$, not working in t	89.35	432
wife not working in $t - 1$, working in t	95.05	283
wife not working in $t - 1$, not working in $t - 1$	94.58	3,063
	Husbands older than 60	
	% of husbands working in period t	N
wife working in $t - 1$, working in t	89.92	5,139
wife working in $t - 1$, not working in t	72.68	527
wife not working in $t - 1$, working in t	85.78	204
wife not working in $t - 1$, not working in $t - 1$	84.69	4,410

Table 4: Participation status in period t of wives who were working in $t - 1$.

	Wives aged 50 to 60	
	% of wives working in period t	N
husband working in $t - 1$, working in t	91.96	4,874
husband working in $t - 1$, not working in t	81.49	443
husband not working in $t - 1$, working in t	85.41	185
husband not working in $t - 1$, not working in $t - 1$	87.04	1,096
	Wives older than 60	
	% of wives working in period t	N
husband working in $t - 1$, working in t	81.56	1,128
husband working in $t - 1$, not working in t	66.96	230
husband not working in $t - 1$, working in t	89.36	94
husband not working in $t - 1$, not working in $t - 1$	70.53	811

Table 5: Estimates of selection equations for married men and women, by pension type.

	Men	Women
age_{it}	0.012 (0.070)	0.174* (0.075)
$age_{it}^2/100$	-0.106* (0.055)	-0.189** (0.064)
$d62_{it}$	-0.396** (0.053)	-0.166** (0.049)
$d65_{it}$	-0.130* (0.053)	-0.080 (0.065)
bad_h_{it}	-0.239** (0.050)	-0.142** (0.051)
$wealth_{it}/1000$	-0.003* (0.002)	-0.002 (0.001)
$d62_{it}^{sp}$	-0.169** (0.037)	-0.092** (0.033)
$penDB_{it}^{sp}$	-0.132** (0.040)	-0.133** (0.030)
$edu1_{it}$	0.188** (0.039)	0.401** (0.038)
$edu2_{it}$	0.098** (0.037)	0.259** (0.034)
$\overline{age}_i$	0.138 (0.086)	-0.068 (0.099)
$\overline{age}_{it}^2$	-0.051 (0.066)	0.065 (0.090)
$\overline{d62}_{it}^{sp}$	-0.193 (0.154)	-0.360* (0.159)
$\overline{d65}_{it}^{sp}$	-0.071 (0.175)	0.236 (0.208)
$\overline{bad_h}_{it}$	-0.822** (0.070)	-0.983** (0.069)
$\overline{wealth}_{it}/1000$	-0.001 (0.002)	-0.005** (0.002)
constant	-3.897 (2.757)	-0.869 (2.949)
N	9,970	11,301

ⁱ Robust standard errors in parentheses.

ⁱⁱ * indicates the coefficient is significant at 5%. ** indicates significance at 1%.

ⁱⁱⁱ Both regressions include, year dummies and a measure of the unemployment rate at period t for men aged 55 and older in order to control for economy-wide effects.

^{iv} Dummies of the form $dage_j$ are equal to 1 if the individual is older than age_j . Dummies of the form $dage_j^{sp}$ are equal to 1 if the spouse is older than age_j . The dummy bad_h is equal to 1 if the individual is in fair or poor health. The dummy $work5y$ is equal to 1 if the individual has kept at least one job for 5 or more years. The dummies $redu1$ and $redu2$ indicate whether the individual has at least some college or is a high school graduate, respectively.

Table 6: Wage process.

	Dependent variable: $\ln w_{it}$	
	Men	Women
age_{it}	0.082* (0.039)	0.082 (0.044)
$age_{it}^2/100$	-0.077* (0.035)	-0.078 (0.040)
bad_h_{it}	0.002 (0.031)	-0.016 (0.034)
$edu1_i$	0.217** (0.023)	0.303** (0.038)
$edu1_{it} \times DC$	0.197** (0.030)	0.134** (0.035)
$edu2_i$	0.072** (0.022)	0.054 (0.029)
$edu2_{it} \times DC$	0.139** (0.032)	0.142** (0.034)
$\overline{age}_{it}$	-0.035 (0.044)	-0.075* (0.038)
$\overline{age}_{it}^2/100$	0.022 (0.035)	0.060 (0.031)
$\overline{bad_h}_{it}$	-0.227** (0.052)	-0.258** (0.088)
$\overline{wealth}_{it}$	0.010** (0.001)	0.006** (0.001)
$\hat{\lambda}_{it}^0$	-0.021 (0.086)	0.120 (0.114)
$\hat{\lambda}_{it}^{DC}$	0.016 (0.092)	0.181 (0.116)
constant	1.511 (1.431)	2.193 (1.385)
N	5,955	4,560

ⁱ Standard errors obtained from 2,500 bootstrap replications.

ⁱⁱ * indicates the coefficient is significant at 5%. ** indicates significance at 1%.

ⁱⁱⁱ The regression include a constant, year dummies, and a measure of the unemployment rate at period t for men aged 55 and older in order to control for economy-wide effects.

^{iv} Dummies of the form $dage_j$ are equal to 1 if the individual is older than age_j . The dummy bad_h is equal to 1 if the individual is in fair or poor health. The dummy $work5y$ is equal to 1 if the individual has kept at least one job for 5 or more years. The dummies $redu1$ and $redu2$ indicate whether the individual has at least some college or is a high school graduate, respectively.

Table 7: Estimates of parameters from out-of-pocket health expenditures.

	Men	Women
age	0.444** (0.137)	0.013 (0.145)
age ² /100	-0.353** (0.115)	-0.000 (0.001)
age × d65	0.071* (0.029)	-0.006 (0.029)
h_bad	1.316** (0.285)	-7.512* (3.414)
h_bad × age	-0.018** (0.007)	0.246* (0.100)
h_bad × age ² /100	-1.509 (0.785)	-0.182* (0.072)
h_bad × d65	-0.016 (0.121)	-0.387** (0.147)
cohort	-4.778* (1.968)	-1.754 (1.321)
cohort × h_bad	0.149 (0.167)	0.331* (0.163)
cohort × age	0.074** (0.021)	0.030 (0.021)
cohort × d65	-0.745 (1.368)	-0.412 (0.356)
d65	-4.495* (1.868)	0.519 (1.906)
wave=2	-0.720** (0.064)	-0.599** (0.073)
wave=3	-0.411** (0.064)	-0.290** (0.069)
wave=4	-0.425** (0.065)	-0.359** (0.069)
wave=5	-0.396** (0.067)	-0.271** (0.070)
wave=6	-0.255** (0.074)	-0.125 (0.075)
wave=7	-0.172* (0.077)	-0.078 (0.082)
constant	-7.820 (4.093)	5.551 (4.335)
σ_ψ^2	2.100** 0.032	1.955** (0.031)
N	8,442	7,777

Figure 1: Sources of income for couples where the husband is aged 65 to 70, by income quintiles.




Figure 2: Retirement frequencies for married men and women, as a function of age.




Figure 3: Comparison of actual and assumed Social Security claiming date. Men.




Figure 4: Comparison of actual and assumed Social Security claiming date. Women.




Figure 5: Retirement frequencies by pension type. Men.




Figure 6: Retirement frequencies by pension type. Women.




Figure 7: Differences in spouses' retirement dates as a function of age difference between them.




Figure 8: Optimal participation choices of couples as a function of male and female accumulated earnings. Baseline.




Figure 9: Household value function as a function of male and female accumulated earnings. Baseline.




Figure 10: Optimal participation choices of couples as a function of male and female accumulated earnings. Assets = \$40,000.




Figure 11: Household value function for couples as a function of male and female accumulated earnings. Assets = \$40,000.




Figure 12: Optimal participation choices of couples as a function of male and female accumulated earnings. Husband in bad health.




Figure 13: Household value function for couples as a function of male and female accumulated earnings. Husband in bad health.




Figure 14: Optimal participation choices of couples as a function of male and female accumulated earnings. Wages exchanged with respect to baseline.




Figure 15: Household value function for couples as a function of male and female accumulated earnings. Wages exchanged with respect to baseline.




Figure 16: Optimal participation choices of couples as a function of male and female accumulated earnings. Non-separabilities in spouses' preferences.




Figure 17: Household value function for couples as a function of male and female accumulated earnings. Non separabilities in spouses' preferences.




Figure 18: Optimal participation choices of couples as a function of male and female accumulated earnings. Policy change.




Figure 19: Household value function for couples as a function of male and female accumulated earnings. Policy change.

