Journal of Public Economics 91 (2007) 1878 – 1901
JOURNAL OF PUBLIC ECONOMICS
www.elsevier.com/locate/econbase

Public pension programmes and the retirement of married couples in Denmark

Paul Bingley [a,b,c,d,*], Gauthier Lanot [e]

[a] Danish National Centre for Social Research, Copenhagen, Denmark
[b] Department of Epidemiology, University of Southern Denmark, Denmark
[c] Economics Department, University of Warwick, UK
[d] Centre for Corporate Performance, Aarhus University Business School, Denmark
[e] School of Economic and Management Studies, Keele University, UK

Received 1 October 2006; received in revised form 7 September 2007; accepted 9 September 2007
Available online 19 September 2007

Abstract

In this paper we study the economic determinants of the joint retirement process of married couples. We propose a tractable dynamic discrete choice model for retirement decisions which allows for non-trivial saving behaviour. We estimate the model on a 1% sample of Danish couples of potential retirement age drawn from a population-based administrative register. The introduction and subsequent reforms of a publicly financed early retirement programme provide us with variation in the data to insure identification of the parameters of interest: the elasticities of participation/retirement with respect to income flows. Our estimates imply a significant asymmetry in the sensitivity of retirement behaviour of men and women with respect to variation in their own, or their spouse's, income flows.

JEL classification: J26; H24; C35
Keywords: Retirement; Pensions; Dynamic structural model

1. Introduction

A growing fraction of people of potential retirement age are in two-earner married couples. In Europe and North America over recent decades the labour force participation rate of women of all ages has risen, and now an increasing proportion of women are reaching their 50's having had a long labour market career.

Career married couples tend to retire close together in calendar time even though they are often of different ages. Given their numerical importance and correlated behaviour, couples are the

* Corresponding author. Danish National Centre for Social Research, Herluf Trolles Gade 11, 1052 Copenhagen, Denmark.
E-mail address: pbingley@sfi.dk (P. Bingley).

0047-2727/$ - see front matter
doi:10.1016/j.jpubeco.2007.09.006

natural unit for analysis of retirement decisions. Empirical studies have found support for different explanations of joint retirement: complementarity of leisure, correlated preferences, common shocks, shared finances and combinations of these factors.

In most empirical work there is a balance to be struck between complexity of the behaviour being represented and the amount of information contained in the data. The amount of useful variation depends upon the sampling frame, the set of controls and the institutional setting. The latter is of overriding importance when observational data is used. A majority of published pension and retirement studies use the US Health and Retirement Study (HRS). This is a very rich set of variables following a modest number of seniors each second year since 1992. The US institutional setting poses several modelling challenges. Employer-provided health insurance, re-entry to the labour force after retirement from a career job, joint taxation and joint benefits all require careful attention be paid to several simultaneous processes.[1]

It is the aim of this study to address some of the difficulties that have hindered couples' retirement studies to date. We set up a dynamic structural model of joint retirement decisions. This is applied to a 1% sample of career married couples drawn from 25 years of population-based longitudinal data from Denmark. The Danish context is perhaps uniquely informative because of its combination of universal health insurance coverage, individual-based programme eligibility and rare labour market re-entry. This must be contrasted with the institutions and consequent modelling difficulties that researchers using US data in particular have to face.

Furthermore, in Denmark a generous public pension programme (*efterløn*, hereafter PEW) was introduced at the end of the 1970s, with dramatic effects. Within two years, the participation rate for married men aged 61–65 fell from 70% to 50%, while the rate for those aged 54–58 remained almost unchanged. Eligibility to the programme was not exogenous, but was arguably pre-determined in the sense that it was a function of historical unemployment insurance contributions, which previously did not passport pension programme eligibility. These criteria were changed at different times during our sample period. Eligibility varied both within and between married couples, entitlement was all-or-nothing, and re-entry to the labour force after first receipt of PEW, meant permanent disqualification. An important contribution of this study is exploiting these discontinuities by way of carefully characterising income streams from potential future retirement ages to provide exogenous income variation that helps to identify the model.

The population-based longitudinal administrative data we use here covers 1977–2001 and does not suffer from sample attrition.[2] Furthermore it brings the relevant planning horizon within the observed time period for many couples. In comparison to the HRS for example, our Danish dataset can credibly sustain a good degree of model complexity.

Three broad approaches have been followed in the literature on the economic determinants of retirement. First, structural modelling of the household as a decision unit where all the elements deemed important to describe current and future retirement options are characterised in detail and embedded in a structural discrete choice model.[3]

[1] For example Van der Klaauw and Wolpin (2005).

[2] Emigration of older workers is negligible, and what remains is "natural" sample attrition due to marital separation, divorce and death.

[3] Following the seminal work of Rust and Phelan (1997), which explains the association between eligibility to health insurance and labour force participation of older individuals, Blau and Gilleskie (2004) estimate a dynamic structural model of the joint employment decision of older married couples based on the HRS. Health insurance plans are bundled with the employment contract to a certain employer and the contribution of the paper is to model how aversion to the risk of uninsured medical expenditure affects employment transitions. Again based on the HRS, Van der Klaauw and Wolpin (2005) set up an ambitious structural dynamic model with uncertainty for older married couples where savings decisions are the central channel through which pension specific characteristics determine the joint retirement age.

A second approach places more emphasis on the decision process within the household as the important determinant of association between retirement dates. Within this approach, two different routes (maintained behavioural hypotheses) have been followed: some studies assume that spouses are involved in bargaining over the outcomes of their labour force participation decisions (Nash or otherwise),[4] while others posit that decisions within the household always achieve Pareto optimality.[5]

A third approach places less emphasis on structural modelling, or even on the details of decision making within household, and relies on reduced form dynamic (or even static) models of joint retirement ages as a response to variation of earnings in-work and income out-of-work.[6,7] This has the advantage of allowing a complete statistical description of the data of interest, at the expense of less transparent economic interpretation. In particular, these models remain silent about the effects of out-of-sample policy reforms.

Our stochastic dynamic model of married couples retirement ages is such that individuals are forward looking and we allow for non-trivial (although not general) consumption/saving behaviour. The household is the unit on which to base policy analysis because it is at this level that decisions are internalized. We do not set up a model of decision making within the household, and so implicitly take a unitary approach. However, our model structure allows us in principle to distinguish between alternative explanations of the joint timing of retirement within couples: correlated preferences due to leisure complementarity or substitutability, common shocks, and shared finances. In particular, we show the effect of changes in programme generosity and non-work income on the probability of each spouse retiring at each age.

The rest of the paper is organised as follows. The next section describes the institutional specifics of the Danish social security system. This is followed by a description of the data. The fourth section presents our empirical model of couples' retirement. Section 5 presents and interprets the estimation results and provides some counterfactual policy analysis. Finally, we conclude.

2. Institutional background and data description

2.1. Income security programmes for the elderly in Denmark

There are five transfer programmes relevant to the elderly population: old age pension, social and disability pension, unemployment insurance benefit, post-employment wage (PEW) and

[4] Gustman and Steinmeier (2000) estimate a forward-looking collective model where elderly married couples bargain non-cooperatively. Uncertainty is ignored and retirement ages are decided optimally at the beginning of the lifecycle. Retirement is an absorbing state. The incomplete estimation problem of multiple Nash equilibria is addressed by assuming the first spouse to retire commits the couple to his or her preferred equilibrium. This work is extended further in Gustman and Steinmeier (2002, 2004).

[5] Michaud and Vermuelen (2004) estimate a static structural collective model where elderly couples bargain cooperatively to reach Pareto efficiency. They assume that widow(er)s have the same preferences as married individuals, and that marriage only gives utility through externalities in leisure.

[6] Blau (1997, 1998) accounts for forward-looking behaviour by including in the set of explanatory variables a small number of expected income streams corresponding to typical employment sequences. However, the dynamic programming problem is not solved. Sedillot and Walraet (2002) follow Blau for French couples. Mastrogiacomo, Alessie, and Lindeboom (2004) estimate a reduced form multinomial logit for Dutch couples.

[7] Using a similar data source to ours, An, Christensen, and Datta Gupta (2004) set up a reduced form dynamic model for the retirement of Danish couples. This is estimated as a bivariate mixed proportional hazard.

public employees pension (PEP). This sub-section outlines the relevant features of these programmes with the aim of illustrating that the PEW and PEP programmes are the most relevant early exit routes. Furthermore, amongst this variety of programmes, reform of PEW offers important variation in replacement rates which can help identify the behavioural effects of financial incentives.

The publicly financed old age pension (*folkepension*) was available at 67 during our sample period, covers everyone regardless of previous attachment to the labour market, and was worth 28 (30)% of average production worker earnings (APWE) in 1980 (1995). Provision for those not working due to disability or long-term sickness is through the social and disability pension (*fortidspension*) programme. It is available to those aged over 17 according to three levels of assessed capacity for work, based on medical and social criteria. The lowest level of entitlement corresponds to old age pension. Individuals are not allowed to work, even part-time, while in receipt of social and disability pension, as this would lead to disqualification. There is no re-assessment of health after eligibility has initially been established and consequently labour market re-entry is negligible.

Unemployment Insurance Benefit (*arbejdsløshedsforsikring*) is available to those who are not working, are actively seeking work and currently members of an unemployment insurance fund and have been insured for at least 6 of the previous 12 months. Benefit lasts for up to 30 months, is 90% of former wage up to a ceiling and was worth 71% (60%) of APWE in 1980 (1995). The Post-Employment Wage (PEW) is an early retirement programme where remuneration is the same as for Unemployment Insurance Benefit, but after 30 months, instead of support being withdrawn, it is reduced to 80% of its former level. Eligibility was restricted to those aged 60–66 who have contributed to an unemployment insurance fund.

The Public Sector Employees Pension (*tjænstemændspension*, hereafter PEP) is available to civil servants in selected occupations from age 60. Entitlement is a function of years of service and seniority (or pensionable wage).

Social and disability pension, social assistance and unemployment insurance benefit programmes experienced no major reforms relevant to the elderly population. PEW was introduced in 1979 conditional on current UI fund membership and membership for 5 out of the previous 10 years. The eligibility requirement for the number of years of membership was increased in 1980 (10/15), 1985 (15/20), 1990 (20/25) and 1995 (25/30). Before 1979 there were no pension rights associated with UI Fund membership. In this sense PEW membership is pre-determined.

Those in receipt of PEW are allowed to work for up to 200 hours each year. Exceeding this level leads to permanent disqualification from the programme. We consider this level insufficient to constitute part-time work or partial retirement. Hence labour market re-entry after initial PEW receipt is negligible.

In summary, PEW and PEP are the most attractive early exit routes from work for those who are eligible. PEW offers high replacement rates for low wage workers.

2.2. Health insurance and occupational pensions

Health insurance provision is universal, and financed through general taxation. Only prescription medicine purchases have co-payments below a ceiling, but these are still heavily subsidized. Health insurance is also independent of the employment relationship.

Employer-provided pensions are organised along occupational and industry lines into about 50 different pension funds. Membership of an occupational pension fund is mandatory following collective agreements. Consequently, the fund follows the worker rather than the employer, and individual transition between funds is uncommon. Movement of workers between employers has no consequence for pension entitlements (Seiersen, 2003). Employer lock-in through occupational pensions is not an issue in Denmark.

Most occupational funds were established in the early 1990's with contribution rates of 1–2% of salary. Rates have gradually risen to 9–10% for blue collar and 12–18% for white collar workers in 2006. Although most workers were members of occupational pension schemes by 2001 (the end of our sample period), pension funds were still far from maturity due to short contribution histories and initially low contribution rates. Even for those retiring in 2001, private occupational pension receipts represent a low proportion of retirement income. In view of the immaturity of these schemes, and the additional assumptions on fund growth expectations that would be required to model them appropriately, occupational pension receipts are not included in expected benefit calculations in the current application.[8]

These institutional features have important implications for modelling of retirement behaviour in Denmark which contrast sharply with the US setting. In the US employer-provided health insurance is an important determinant of elderly labour supply, which has consequences for movement of workers between employers and out of the labour force. This ties together work decisions with individual expected future health and attitude to risk through health insurance. Furthermore, most plans cover spouses, and couples choose to contribute to the most generous plan that their respective employers offer. This ties spouses to their employers differentially according to plan generosity.

In Denmark, health insurance is irrelevant to the retirement decision, although health itself of course might be an important determinant. Employer-to-employer movements are neither relevant to health insurance nor occupational pensions. This greatly simplifies the problem of modelling elderly labour supply, since employer-to-employer movements need not be explicitly considered and neither does health insurance plan design and choice need to be taken into account.

The effective non-work requirements of social and disability pension and PEW, which rule out part-time or partial retirement, allow us to simplify the modelling further. Moreover, the permanent benefit disqualification following breaking these non-work requirements, explains why labour market re-entry after leaving a career job is in fact negligible.

The aim of our structural model of couples' retirement is essentially to characterise the sensitivity of husband and wife retirement decisions with respect to husband and wife non-work income. This provides a common currency with which to compare evidence from different models estimated in different institutional settings on the basis of different datasets. While it is important to take distinctive local institutional features appropriately into account in modelling, and care needs to be taken when comparing elasticities between settings, it is precisely these distinctive institutional details which can inform us about the determinants of behaviour. In Denmark, during our sample period, the financial consequences of elderly labour supply decisions were relatively uncomplicated.

[8] All schemes are Defined Contribution, but many also have a Defined Benefit element in the form of guaranteed average minimum interest rates earned. Information about fund growth and guarantee mix is not available at the individual level. However, it is possible to make plausible assumptions about these parameters by occupation and year. In initial estimations of a reduced form option value model of couples' retirement on this dataset, coefficients of interest were robust to inclusion in benefits of expected fund growth equal to average realised fund performance. As private occupational pension funds mature, they will obviously grow in importance, and as more recent data becomes available it will be increasingly relevant to incorporate this into benefit calculations. Incorporating expectations of occupational pension receipts into the analysis is left for future work.

Table 1
Descriptive statistics for 1% retirement age estimation sample

	Men				Women			
	1st obs.		All obs.		1st obs.		All obs.	
Log earnings	11.7606	2.5531	11.9062	1.2380	11.3330	2.2115	11.4094	1.3915
Wealth / 1,000,000	0.9798	1.3081	1.0599	1.3391	0.9798	1.3081	1.0599	1.3391
PEP entitled	0.0000	0.0000	0.1184	0.3231	0.0000	0.0000	0.0772	0.2669
PEW entitled	0.0000	0.0000	0.3315	0.4708	0.0000	0.0000	0.2764	0.4472
Tenure	7.7581	4.0020	11.8853	5.3051	6.9187	3.2442	10.7465	4.3367
Experience	11.4935	2.2099	15.6221	4.0635	9.0709	2.2759	13.0343	3.7628
Age at obs.	56.1242	1.4526	62.8598	5.5872	53.9005	1.4912	60.5950	5.5964
Retirement age	62.3296	3.1411	62.8511	3.3482	60.3296	2.8849	60.7812	3.0755
Censoring age	62.2014	3.7283	64.5400	4.0932	59.6540	3.2642	61.5382	3.9966
# observations	1425		10392		1425		10392	

This simplifies analysis because of a rather transparent incentive structure which helps direct interpretation of results.

3. Data description

The data originate from Danish public administrative registers, which are organised thematically into general research databases by Statistics Denmark. Since 1967, each resident of Denmark has been issued with a unique social security number, which is used throughout public administration and allows individuals to be consistently identified over time. For research, this data has several virtues. It is population-based and immune from the usual sample selection and attrition issues of sample survey data, in that birth, migration or death are the only ways in or out. Data quality is know to be high in the sense that recall bias is not an issue and the purpose and mode of collection is well documented. In particular it is relevant for our purposes that there are strong financial incentives for the tax authorities, the social security administration and the individuals concerned to report correctly.[9]

In this paper we consider the population of married couples at any time during the years 1977–2001, where both are aged over 49 years, the husband is observed working until at least age 57 and the wife until at least 55. Should the couple separate or divorce, or either member die or emigrate, then both members of the couple are observed in the sample for the last time in the previous year. Men are assumed to be first at risk of retirement at age 58 and women at 56. This population data (on 153,818 couples) is used to estimate earnings functions (1,239,686 male and 1,273,664 female observations) which are used to generate unobserved potential in-work net income flows. In order to make computations manageable, the structural age of retirement model is estimated on a 1% sub-sample of these couples where at least one partner is at risk of retirement (10,392 observations on 1425 couples).

3.1. Descriptive statistics

Table 1 presents descriptive statistics for the 1% sample of married couples used in retirement age modelling. Substantive columns 1–4 show means and standard deviations for men and

[9] Our dataset is drawn from the Statistics Denmark tax-oriented income statistics register and the integrated database for labour market research. See Leth-Sørensen (1993) for further details.

columns 5–8 are for women. Columns 1–2 and 5–6 show observations when individuals are first at risk of retirement, whereas columns 3–4 and 7–8 include all observations.

Earnings originate from pay-as-you-earn employer filings with the tax authorities. Employment histories come from mandatory pension contributions (*ATP*) which have been made as a function of hours worked since 1965. Contributions match each employer to employee and these are summed on an annual basis allowing calculation of labour market experience and firm tenure since 1965.

Wealth is obtained from public administrative records collected for the purposes of wealth taxation.[10] Wealth was taxed until 1996 on a joint basis. Wealth comprises reports from financial institutions on cash holdings in banks, bonds, value of shares, own capital in domestic companies, together with the sum of liabilities. Further to these self-reported values of cooperatively owned housing, cars, boats, and un-quoted shares are included.[11] Value of owner-occupied housing is also assessed for the purposes of taxing imputed rental income. However this is a function of accommodation size and broadly defined neighbourhood which is a poor approximation of market value. An important omission is the value of pension funds, which are only taxed on receipt.

We use wealth reported by financial institutions, ignoring self-reported items, the value of owner-occupied housing and pension fund value. Only one observation of wealth is used to initialise the consumption process. This is taken from the year before the first spouse is at risk of retirement, or 1980, whichever is the later. A reform to wealth taxation in 1980 marked the beginning of the wealth database we have access to.[12]

3.2. Retirement

Retirement is defined as a transition to long-term (at least 1 year) non-work for older workers. Our definition of non-work requires that two conditions are fulfilled within a given calendar year: (1) zero mandatory pension contributions which are only made in-work; (2) primary source of income during the year is not from current employment.[13] Our register data does not allow us to time retirement within the year, and all we are able to say is that the transition from work to retirement occurred sometime during the calendar year before the first full year of non-work.

Fig. 1 presents the evolution over time of labour force participation of the population of married men and women for selected age groups. The introduction of PEW in 1979 can clearly be seen to coincide with a sharp reduction in elderly male participation. For women the reduction was somewhat more modest, falling from a lower base participation rate, but still obviously within a few years of PEW. Similar reductions are not observed at younger ages, where secular and cyclical movements predominate.

3.3. Programme eligibility

Changes in eligibility criteria to the PEW programme form the basis for identification of the effects of transfer programmes on individual income and retirement. Calculating individual eligibility requires knowledge of unemployment insurance fund membership back to 1965. We do

[10] After 1996, financial institutions continued to report individual holdings on a similar basis.

[11] See Browning and Leth-Petersen (2003) for details on wealth measurement in this database in the context of imputing individual consumption.

[12] This means that for couples observed before 1980, wealth may first be observed while they are at risk of retirement. This timing issue is taken into account in estimation.

[13] This definition is similar to that used in other retirement studies using Danish register data, for example Pedersen and Smith (1996).

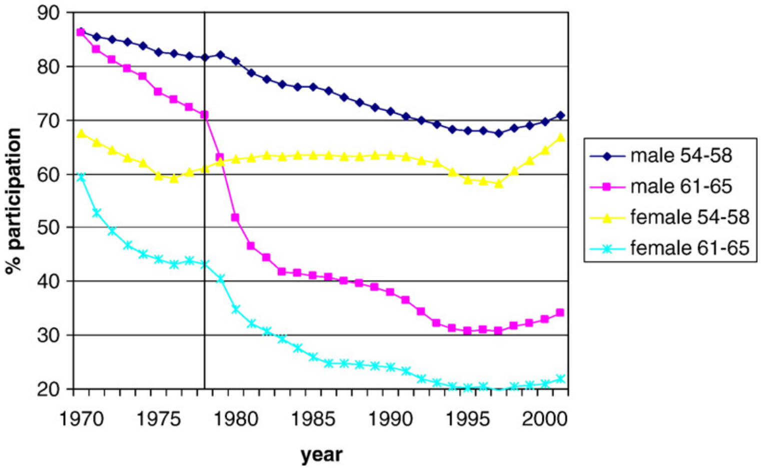

Fig. 1. Labour force participation rate for married couples.
Note: These numbers are own calculations from population-based administrative register data of all married couples resident in Denmark.

not observe membership prior to 1977. However, we do know the employment history for each person from 1965, via the history of mandatory social security contributions. Fund membership is imputed during these missing years by using employment history together with fund membership status in 1977. The assumption is that membership status in 1977 reflects membership during the period 1965–1976.[14]

Unemployment insurance fund contributions, UI Benefits and PEW entitlements are a function of hours worked. The unemployment insurance system and social security contributions measure hours of work in similar ways. Hence, conditional on the assumption regarding the equivalence of mandatory social security contributions and continued insurance fund membership, we are able to impute to each worker an age at which he or she becomes eligible to the PEW. In our sample, 66% of men and 58% of women are ever potentially eligible. Of married couples 47% are both eligible, 23% neither, 11% women only and 19% men only.

3.4. Earnings functions

The dynamic programming model requires predictions of potential future earnings in order to calculate expected future consequences of current work decisions. For retired individuals, obviously what they would have earned if they had remained in work is not observed, and counterfactual earnings predictions must be made. In generating these counterfactuals, sample selection might be an issue. If for example those who retire early have systematically lower earnings than those retiring late, all else equal, then predicting earnings for early retirees based on the earnings of later retirees will result in biases.

Table 2 presents a set of auxiliary regressions with the aim of showing the degree of sample selection involved in estimating earnings functions for our population. Each cell of the table contains a coefficient from a separate regression. Each regression explains log earnings at an age

[14] This can be justified by the fact that very few older workers (only 3% of those aged over 49) change membership status in the first 6 year period over which we observe membership 1977–1983.

Table 2
Earnings sensitivity to future retirement and attrition status

Age	Women 55				Men 57			
	Retire coefficient	S.E.	Attrite coefficient	S.E.	Retire coefficient	S.E.	Attrite coefficient	S.E.
56	−0.0598	0.0840	−0.0940	0.0899				
57	−0.1040	0.0426	−0.1054	0.0481				
58	0.0296	0.0323	0.0270	0.0298	−0.0052	0.0132	−0.0105	0.0132
59	−0.0231	0.0382	−0.0202	0.0348	−0.0159	0.0154	−0.0197	0.0151
60	0.0650	0.0390	0.0524	0.0361	−0.0737	0.0209	−0.0713	0.0199
61	0.0118	0.0637	−0.0077	0.0572	−0.0350	0.0234	−0.0370	0.0227
62	0.0222	0.0905	0.0196	0.0825	−0.0278	0.0250	−0.0299	0.0242
63	0.0830	0.1234	0.0877	0.1100	−0.0161	0.0339	−0.0357	0.0320
64	0.0497	0.1510	0.1010	0.1438	−0.0790	0.0374	−0.0450	0.0363
65	−0.2806	0.1756	−0.3313	0.1760	0.0341	0.0320	0.0410	0.0308
66	−0.1324	0.1190	−0.1371	0.1187	−0.0039	0.0381	−0.0146	0.0367
67	−0.0013	0.0773	0.0057	0.0732	−0.0200	0.0412	0.0006	0.0392
68	0.0222	0.0892	−0.0012	0.0870	0.0128	0.0491	0.0060	0.0470

Note: OLS estimates based on the population counterpart of our 1% structural retirement model estimation sample. Each coefficient is from a separate regression explaining log earnings at last age before risk of retirement for individuals retiring at different sets of two consecutive future ages. The coefficient presented is of a dummy variable for retiring at the earlier of the two ages. Other variables which are included but (not presented in the table) are year dummies and quadratics in tenure and experience.

defined in the *column* header (the last age before "retirement risk"), for a sample of individuals retiring at two subsequent future ages. For example, the first cell of the table with the *row* header 56 is the coefficient on a dummy variable for retiring at age 56 estimated on a sample of those retiring at either 56 or 57. The coefficient −0.0598 implies that earnings at age 56 are 5.98% lower than for someone who will retire at age 57. The *t*-statistic on this coefficient is a test of earnings selectivity according to future retirement age. This is insignificant.

Other explanatory variables included in the regression but not presented in the table are log earnings at age 55, year dummies and quadratics in tenure and experience. The first four columns of the table are for earnings of women aged 55 and the last four columns are for men aged 57. Columns one and five examine retirement selection and columns three and seven look at retirement *and* attrition from the sample.

Estimates are based on the entire population of (153,818) married couples, as are all of the earnings functions in the paper. For women, future retirement or attrition age is never a significant explanatory of "pre-retirement-risk" earnings. For men this is also the case for all ages except for those retiring or leaving the sample at age 60 rather than 61, who have respectively 7.4 and 7.1% lower earnings. Of course 60 is age of first eligibility to PEW, and upward bias in future earnings predictions would bias replacement rate predictions downward and in turn produce spurious incentive effects between these two ages. However, in view of the limited nature of this observed selection — only for men at one age, and the huge additional computational burden of taking into account this modest degree of earnings selectivity, we follow the vast majority of structural retirement literature and estimate these reduced form earnings functions in a first stage without selectivity correction.[15] An important virtue of this distinct first stage is that it reduces the computational burden of estimating earnings functions, which allows information from the whole population to be used and allows for a very flexible earnings specification, of effectively about 4000 coefficients.

[15] In particular, incorporating dynamic earnings selection would require calling the potential social security benefit stream calculator at each iteration of the likelihood.

Potential future real gross earnings in work are estimated from the population data. Observed log real earnings at all relevant ages are regressed on log real gross earnings measured when the individual enters the sample (this provides us with a control which is out of the period of interest and it will capture part of the variability which may be specific to the individual), quadratics in job tenure and labour market experience and dummies for 5-year birth cohorts. There are 190 separate OLS regressions for each gender. These represent all relevant combinations of present work ages (50 through 68) and future work ages (51 through 69). Each regression equation has up to 11 co-variates, plus an intercept, depending on the number of age cohorts spanned.

Descriptive statistics for the estimation data and the estimated coefficients are presented in Table 3. The substantive columns 1–4 contain means and standard deviations (not standard errors) of estimated coefficients. The regressions are run with a mean of just over 30,000 observations (and a minimum of 624), where an observation represents a unique combination of current and future ages. These produce a mean R-squared of around 0.5. Columns 5–8 contain means and standard deviations of the in-work population co-variates on which the regressions are run.

3.5. Income flows

Potential future income out of work for all possible alternative work trajectories is calculated on the basis of observed couple characteristics. These are run through a transfer entitlement routine capturing all important features of the relevant programmes and their evolution. Income taxation of couples is separate for labour income and joint for capital income, and this is taken into account in the calculations.

Average potential income flows by age of retirement are shown in Figs. 2 and 3. A sub-sample of couples is drawn from the estimation sample where a woman aged 55 is married to a man of 57. Fig. 2 illustrates individual net income in and out of work for men (respectively women) in the left

Table 3
Earnings function co-variates and estimated coefficients

	Female coefficients		Male coefficients		Female co-variates		Male co-variates	
	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.
Log earnings	0.5656	0.2086	0.6177	0.2517	11.462	1.5276	11.942	1.7891
Tenure/10	0.1871	0.2590	0.1702	0.1554	0.9965	0.4151	1.1207	0.4964
Tenure/10 sq.	−0.0513	0.1063	−0.0434	0.0601	1.1654	0.8967	1.5023	1.1904
Experience/10	−1.0024	0.9362	−1.6027	1.7330	1.2149	0.3553	1.4842	0.3713
Experience/10 sq.	0.4218	0.4368	0.6448	0.6791	1.6022	0.9046	2.3408	1.1474
Born 1920–24	−0.0142	0.0287	−0.0156	0.0539	0.1144		0.2398	
Born 1925–29	−0.0261	0.0603	−0.0333	0.0824	0.2766		0.2928	
Born 1930–34	−0.0015	0.0640	0.0037	0.0726	0.2911		0.2630	
Born 1935–39	0.0105	0.0439	0.0114	0.0415	0.2364		0.1789	
Born 1940–44	0.0114	0.0399	0.0100	0.0455	0.0798		0.0255	
Born 1945–49	0.0044	0.0241	0.0000	0.0000	0.0016			
Intercept	5.3058	2.2739	5.1441	3.2482				
Age now	57.6667	5.0662	59.0000	4.5947	56.4275	2.9636	58.6770	2.9476
Age in the future	65.3333	5.0662	66.0000	4.5947	62.3172	4.3926	64.0501	3.9553
# observations	30878	30624	33986	32120	1273664		1239686	
R-squared	0.5264	0.1824	0.4979	0.1543				

Note: OLS estimates based on the population counterpart of our 1% structural retirement model estimation sample. OLS regressions explaining future earnings as a function of present observed characteristics are run separately for men and women for all 190 combinations of relevant present (50–68) and future (51–69) potential work ages.

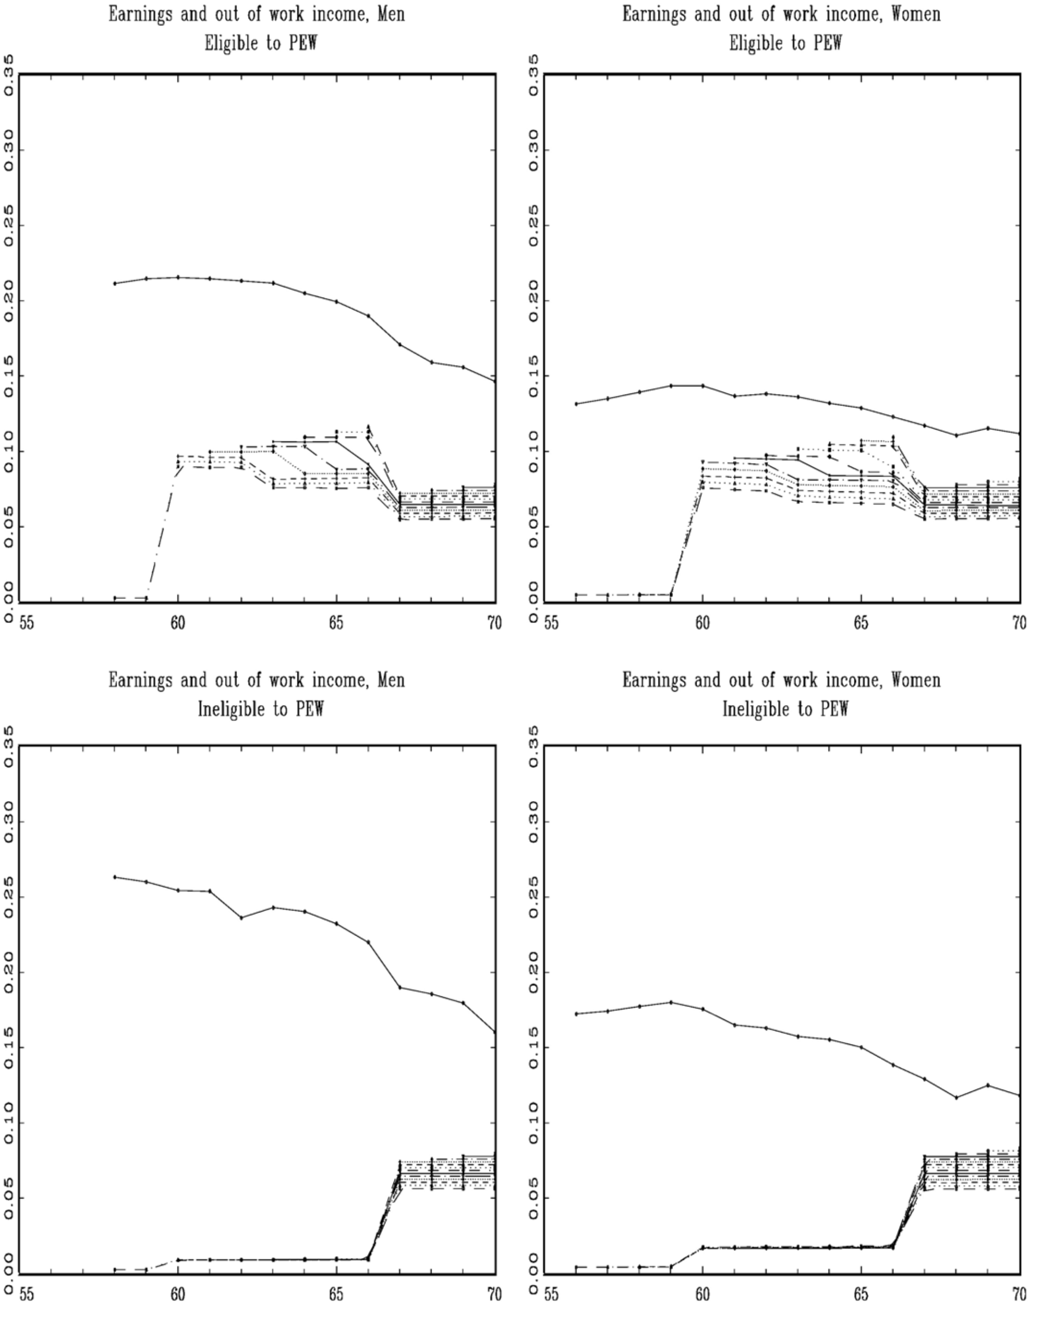

Fig. 2. Marginal income flows by retirement age.

Note: Predicted income flows based upon estimates presented in Table 3. Predictions are for women initially working aged 55 and men aged 57. Solid (uppermost) lines represent household income in work. Lower sets of lines are income flows from different ages of retirement.

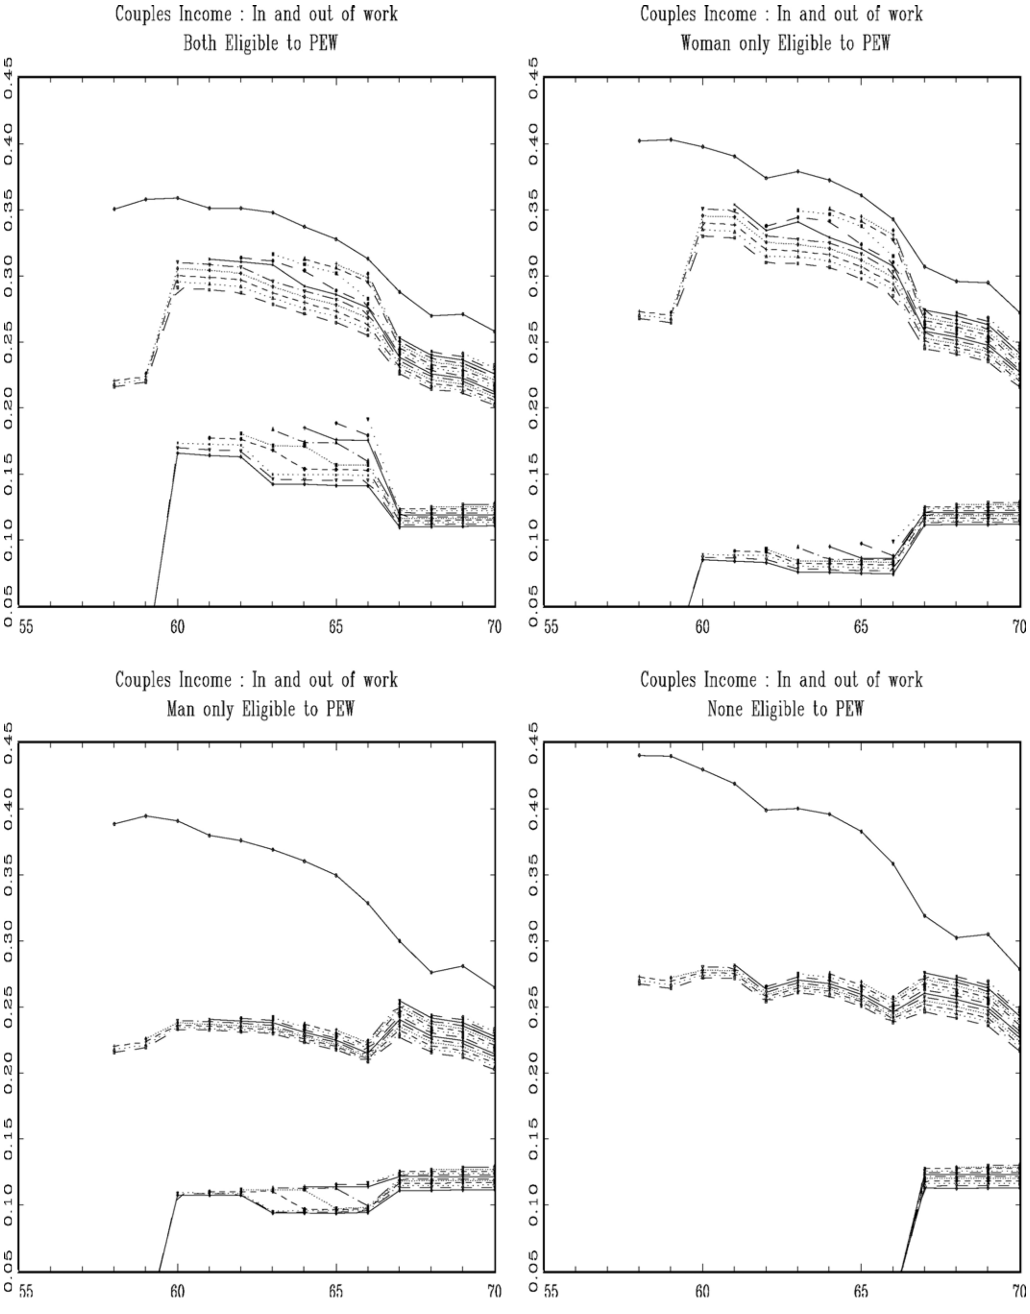

Fig. 3. Joint income flows by retirement ages.

Note: Predicted income flows based upon estimates presented in Table 3. Predictions are for couples initially both working with the woman aged 55 and man 57. Solid (uppermost) lines represent household income for both in work. Sets of lines in the middle of each pane are income when the wife retires at different ages and the husband continues working. The lower sets of lines are income when both are retired.

(right) panes and for those eligible (respectively ineligible) to PEW in the upper (lower) panes. The uppermost solid line in each pane represents mean in-work income and the lower lines are mean out-of-work income for retirement at different ages.[16] The starting (leftmost) points of each lower line represent the mean out of work income for an individual deciding to retire at this age. Along the same line we track the mean out of work income (in real terms) as (s)he ages. The upper panes illustrate well the replacement rate hump due to PEW in contrast to the lower panes where this hump is absent. Mean income in-work for individuals not eligible to PEW is somewhat higher than for ineligibles, adding to the programme-induced replacement rate difference between the two groups.

Average potential total income flows for these couples are shown in Fig. 3. The uppermost solid line represents total income with both partners in work at different ages. The middle set of lines in each pane represents household income starting from the age of female retirement, with the male continuing to work. Each line in this group is for a different age of female retirement. Income when both have retired is shown by the group of lines at the bottom of each pane. Panes on the left are for PEW eligible men and upper panes are for PEW eligible women. Hence the upper right pane corresponds to both spouses being eligible.

Comparing upper and lower panes it can be seen that retirement for eligible women has only a modest (20%) impact on total income, whereas for non-eligible women it reduces total income by about 40%. Similarly comparing left to right panes for eligible men retirement has a large (50%) negative impact on total income and for non-eligible this is a very large (70%) reduction. Mean in-work income is 10% higher for each member of the couple ineligible to PEW.

4. Empirical model

Assume that each period each household consumes a fixed proportion α of its liquid assets whatever the labour force participation state of its members.[17] At time t and in any state $(i,j,R_{\mathrm{m}},R_{\mathrm{w}})$, where i (respectively j) is equal to one if the man (woman) is active in the labour market and 0 otherwise. R_{m} (R_{w}) is equal to 0 if the man has not retired by time t and otherwise is equal to the last date when he (she) was in work. The equation of motion for household liquid assets, $A(t,i,j,R_{\mathrm{m}},R_{\mathrm{w}})$, from the previous period's state $(i',j',R'_{\mathrm{m}},R'_{\mathrm{w}})$ is:

$$A(t,i,j,R_{\mathrm{m}},R_{\mathrm{w}}) = (1+r)(1-\alpha)A(t-1,i',j',R'_{\mathrm{m}},R'_{\mathrm{w}}) + y(t,i,j,R_{\mathrm{m}},R_{\mathrm{w}}), \tag{1}$$

where r stands for the interest rate on savings. Furthermore, since retirement is absorbing the state variables must obey the following dynamics:

$$\begin{aligned}
&0 \le i \le i' \le 1,\ 0 \le R'_{\mathrm{m}} < t-1,\\
&0 \le j \le j' \le 1,\ 0 \le R'_{\mathrm{w}} < t-1,\\
&R_{\mathrm{m}} = R'_{\mathrm{m}}1[\{i'=0\} \cup \{R'_{\mathrm{m}}>0\}] + (t-1)1[\{i'=0\} \cup \{R'_{\mathrm{m}}=0\}],\\
&R_{\mathrm{w}} = R'_{\mathrm{w}}1[\{j'=0\} \cup \{R'_{\mathrm{w}}>0\}] + (t-1)1[\{j'=0\} \cup \{R'_{\mathrm{w}}=0\}].
\end{aligned} \tag{2}$$

[16] The lower lines are shifted upwards slightly (by 3%) in order to improve clarity.

[17] This line of enquiry could be pursued by allowing the consumption function to depend separately on liquid assets accumulated to date and current income flows, instead of summing them and treating them similarly as at present. We could then investigate to the extent to which liquid assets and income contribute differentially to the labour market decisions of the couple. However, we leave this for future work and thank a referee for making this suggestion.

$y(t, i, j, R_{\mathrm{m}}, R_{\mathrm{w}})$ stands for the couples joint income at time t in state $(i, j, R_{\mathrm{m}}, R_{\mathrm{w}})$, i.e. the sum of the spouses' income flows at time t (net earnings or out-of-work income):

$$y(t,i,j,R_{\mathrm{m}},R_{\mathrm{w}}) = y_{\mathrm{m}}(t,i,R_{\mathrm{m}}) + y_{\mathrm{w}}(t,j,R_{\mathrm{w}}). \tag{3}$$

At time t_0, we set the initial value of disposable income as

$$A(t_0,1,1,0,0) = A(t_0).$$

Finally we assume that by some given age all household members have retired.

At time t, when at least one member of the household is active in the labour market, we assume that household h's instantaneous utility flows are specified in the following way:

$$\begin{aligned}
U(h,t,1,1,0,0) &= \beta_{11}^{0} + \beta\frac{(\alpha A(h,t,1,1,0,0))^{1-\rho}}{1-\rho} + \varepsilon(h,t) + \eta(h,t,1,1,0,0),\\
U(h,t,1,0,0,R_{\mathrm{w}}) &= \beta_{10}^{0} + \beta\frac{(\alpha A(h,t,1,0,0,R_{\mathrm{w}}))^{1-\rho}}{1-\rho} + \varepsilon(h,t) + \eta(h,t,1,0,0,R_{\mathrm{w}}),\\
U(h,t,0,1,R_{\mathrm{m}},0) &= \beta_{01}^{0} + \beta\frac{(\alpha A(h,t,0,1,R_{\mathrm{m}},0))^{1-\rho}}{1-\rho} + \varepsilon(h,t) + \eta(h,t,0,1,R_{\mathrm{m}},0),
\end{aligned} \tag{4}$$

where the unobserved components $\eta(h,t,i,j,R_{\mathrm{m}},R_{\mathrm{w}})$ are household, time, and alternative specific, while $\varepsilon(h,t)$ is specific to the household at time t but is alternative invariant. This specification is more general than the specification used, for example, in Berkovec and Stern (1991) where $\rho=1$ and $\alpha=1$.

The parameters of interest are β, ρ which can be understood as the coefficient of relative risk aversion, and α the average propensity to consume out of liquid wealth and current income.[18]

To complete the model we specify the terminal value accruing to a household where both spouses are not working first at time t. The terminal value depends on the stock of assets accumulated so far and not on the (predicted) consumption at time t (as in the case of the utility flows). Furthermore it depends on the expected net present value to the couple of their pension income. To simplify the calculations we specify the terminal value as follows:

$$\begin{aligned}
W(h,t,0,0,R_{\mathrm{m}},R_{\mathrm{w}}) = \beta_{00}^{0} + \beta&\frac{((1+r)(1-\alpha)A(h,t-1,i',j',R_{\mathrm{m}},R_{\mathrm{w}}) + \gamma y(h,t,0,0,R_{\mathrm{m}},R_{\mathrm{w}}))^{1-\rho}}{1-\rho}\\
&+\varepsilon(h) + \eta(h,t,0,0,R_{\mathrm{m}},R_{\mathrm{w}}),
\end{aligned} \tag{5}$$

where at least one of the retirement dates R_{m}, R_{w} is equal to $t-1$ (perhaps both). The parameter γ captures the extent with which the couple discounts the value of its retirement income from retirement date until the couple's death (under the assumption of constant income after retirement). Two factors contribute to the size of γ: the couple's expectation concerning the discount rate (inflation rate and interest rate) and the couple's expectation about its survival. The higher the survival rate and the lower the inflation and interest rate the higher the value of γ. Hence a couple discounting future income flows with an interest rate of 5% and with a survival probability from one period to the next of 0.95 (i.e. with an expected life in retirement of 20 years)

[18] A large value of α means no saving and low dependence on previous history of income flows. It is important to appreciate the distinction between α and the more common notion of the average propensity to consume out of income alone.

and maximum life of 40 years would give a constant flow of one unit of income over its life time an expected net present value of around 10 units of income.[19]

Assuming that the household discounts future flows of utility at a common positive rate δ, the value at time t of a particular course of action of household h can then be evaluated recursively using the Bellman equation

$$W(h,t,i,j,R_{\mathrm{m}},R_{\mathrm{w}}) = U(h,t,i,j,R_{\mathrm{m}},R_{\mathrm{w}}) + \delta E_t\left[\max_{\substack{m\in\{0,i\}\\ n\in\{0,j\}}}\{W(t+1,m,n,R_{\mathrm{m}},R_{\mathrm{w}})\}\right], t\geq 0, \tag{6}$$

and at time t, given previous period labour force participation status (i',j'), the household decision problem becomes

$$\max_{\substack{i\in\{0,i'\}\\ j\in\{0,j'\}}} W(h,t,i,j,R_{\mathrm{m}},R_{\mathrm{w}}). \tag{7}$$

Hence at time t, given the state, both household members plan to be in employment if:

$$W(t,1,1,0,0)\geq W(t,0,1,0,0),$$

$$W(t,1,1,0,0)\geq W(t,1,0,0,0),$$

$$W(t,1,1,0,0)\geq W(t,0,0,0,0). \tag{8}$$

The other possible transitions can be characterised in a similar fashion by a similar set of three inequalities.

Because retirement is an absorbing state for any member of the household, the problem simplifies whenever one of the two spouses has retired by time t. For example assume that $R_{\mathrm{m}}=t-1$ and $R_{\mathrm{w}}=0$, then the wife will remain in work if

$$W(t,0,1,t-1,0)\geq W(t,0,0,t-1,0), \tag{9}$$

and retire otherwise. If instead it is the case that $R_{\mathrm{m}}=0$ and $R_{\mathrm{w}}=t-1$, the husband will participate if

$$W(t,1,0,0,t-1)\geq W(t,0,0,0,t-1), \tag{10}$$

Assumptions about the exact distribution of un-observables are crucial here for calculation of the quantities

$$E_t\left[\max_{\substack{i\in\{0,1\}\\ j\in\{0,1\}}}\{W(t+1,i,j,R_{\mathrm{m}},R_{\mathrm{w}})\}\right]. \tag{11}$$

[19] The expected net present value in the case of a constant survival rate p, a discount rate λ and with T as the maximum age can be calculated as

$$\frac{(1-p)}{1-p^T}\sum_{\tau=0}^{T}\left\{\sum_{k=0}^{\tau}\lambda^k\right\}p^\tau = \frac{1-\lambda^T}{1-\lambda} - \frac{1}{1-p^T}\left(\frac{1-\lambda^{T-1}}{1-\lambda} - \frac{1-(\lambda p)^T}{1-(\lambda p)}\right).$$

Given the above and given a past history, all instances of the unobservables relevant at time t are jointly normally distributed and in particular we assume that

$$\begin{pmatrix} \eta(i,t,1,0,R_\mathrm{m},R_\mathrm{w}) - \eta(i,t,0,0,R_\mathrm{m},R_\mathrm{w}) \\ \eta(i,t,1,0,R_\mathrm{m},R_\mathrm{w}) - \eta(i,t,0,1,R_\mathrm{m},R_\mathrm{w}) \\ \eta(i,t,1,0,R_\mathrm{m},R_\mathrm{w}) - \eta(i,t,1,1,R_\mathrm{m},R_\mathrm{w}) \end{pmatrix} \overset{\mathrm{iid}}{\sim} N(0,\Sigma), \tag{12}$$

with $\Sigma = \begin{bmatrix} 1 & & \\ \sigma_{12} & 1 & \\ \sigma_{13} & \sigma_{23} & 1 \end{bmatrix}$. In practice, we allow the variance–co-variance structure to also be dependent upon household age composition and decision time.

$E_t\left[\max_{i,j} W(t,i,j,R_\mathrm{m},R_\mathrm{w})\right]$ can be calculated recursively in terms of $E_t\left[\max_{i,j} W(t+1,i,j,\ R_\mathrm{m},R_\mathrm{w})\right]$.[20] At each time t, the probability that a household chooses a particular participation pattern, for example

$$\Pr_t[(1,0,0,0) = \arg\max\{W(t,1,1,0,0),W(t,1,0,0,0),W(t,0,1,0,0),W(t,0,0,0,0)\}],$$

is evaluated simply in terms of a trivariate normal joint distribution function and makes the evaluation of the likelihood straightforward.

Our specification captures some of the features that have been considered important in the retirement literature, especially patterns of complementarity/substitutability in leisure and consumption smoothing behaviour.

The complementarity/substitutability pattern that our model allows may be described along the lines presented in Train (2003). Consider first the myopic case where $\delta=0$. In this case, to simplify, we can imagine that each of the four options available (in general) to a couple at entry into our sample depends on an alternative-specific attribute. In our context this could simply be the identity of the alternative, or the alternative-specific income flow. In the case of a multinomial logit specification (where the distribution of the alternative-specific unobservables would be type I extreme value) varying the value of a single attribute affects all the choice probabilities in very particular fashion. This pattern of substitution (a consequence of the IIA property) is quite restrictive. A multinomial probit specification (where the distribution of alternative-specific unobservables is jointly normal) does not impose such restrictions and in principle allows for complex patterns of substitution between alternatives. Hence a variation in an attribute which makes the members of a household more likely to remain both in work, is allowed to have a differential effect on the probability that one spouse decides to chose one of the other patterns of labour market participation available to them.

When allowing the household to account for the future consequences of current actions ($\delta\neq 0$) the patterns of substitution become more complex since not only variation in current attributes matters but variability in future flows of income or more generally utility in all alternatives available in the future matter as well.[21] In the penultimate section, we will consider separate changes in each spouse's flows of income in order to describe the pattern of complementarity/substitution.

[20] These can be expressed analytically and are available from the authors upon request.

[21] See Maestas (2001) for an alternative model of joint retirement decision where complementarity/substitutability is studied at the individual level.

It can easily be shown that accounting for saving implies that today's labour market participation decision depends on distributed lags of previous period's income flows. It is that dependence on past income flows which in principle allows us to identify the average propensity to consume out of liquid wealth and current income.

5. Estimation, results and discussion

5.1. Estimation

Given the assumption of conditional independence formulated above, the likelihood of the sample is straight forward. For example, consider a couple such that the man retires first at some age t_{m} and the woman retires a few years later at age t_{w}. The contribution to the likelihood for such an observation is then easily written as

$$L(t_{\mathrm{m}}, t_{\mathrm{w}}) = \prod_{\tau=t_0}^{t_{\mathrm{m}}-1} \Pr_\tau[\text{both employed}] \times \Pr_{t_{\mathrm{m}}}[\text{husband alone retires}] \times \prod_{\tau=t_{\mathrm{m}}+1}^{t_{\mathrm{w}}-1} \Pr_\tau[\text{wife employed while husband retired}] \times \Pr_{t_{\mathrm{w}}}[\text{wife retires}].$$

Contributions to the likelihood for alternative retirement age patterns can be obtained in a similar fashion. Furthermore, the contribution of right-censored observations (such that one or both retirement ages are not observed within the sample period) can be easily accommodated using conventional arguments.

The simultaneous estimation of all parameters is not easy, the likelihood is very non-linear in the parameters of the model and the computation times are substantial. As a practical solution, we preformed a grid search over the values of δ[22] and maximised the likelihood given these parameter values to find estimates for the other parameters. Table 4 presents results for four sets of estimated parameters for different values of δ, $\delta=0.85$ and $\delta=0.95$. In each case we estimate the parameters of the model assuming first that the variance of the unobserved component is constant over time (the "homoscedastic" case), and then we estimated the model assuming that the variance of the unobserved components varies over time (the "heteroscedastic" case).

As an extension to the model presented in the previous section, the heteroscedastic model is such that the variances of unobserved components are allowed to vary with ages of the spouses such that:

$$\sigma_i = \exp(\theta^i_{\mathrm{m}}(\mathrm{age}_{\mathrm{m}} - 62)/10 + \theta^i_{\mathrm{w}}(\mathrm{age}_w - 62)/10). \tag{13}$$

It captures the likely changes in uncertainty about future utility flows as both spouses age. The normalisation sets the reference variance of the unobserved component to be unity when both spouses are aged 62 at the same time. This is sufficient to allow for the identification of all the parameters of the model.

[22] The rate of time preference is often not estimated but fixed (see for example Berkovec and Stern, 1991). Although our grid search is not ideal, it provides an improvement on the usual practice. We maximised the likelihood for values of δ between 0.725 and 0.975 with a step of 0.025; The results we present in Table 4 give the maximum of the likelihood function over the grid.

Table 4
Maximum likelihood estimation results

δ	Homoscedastic		Heteroscedastic		Homoscedastic		Heteroscedastic	
	0.85		0.85		0.95		0.95	
Parameters	Estimates	S.E.	Estimates	S.E.	Estimates	S.E.	Estimates	S.E.
α	0.450	0.063	0.489	0.055	0.430	0.072	0.420	0.065
ρ	0.091	0.114	0.181	0.099	0.029	0.118	0.030	0.114
γ	2.082	0.290	2.042	0.278	2.507	0.355	2.715	0.396
β	3.540	0.425	3.403	0.496	2.926	0.341	2.283	0.331
β_{10}^0 constant	−0.235	0.214	−0.161	0.159	−0.229	0.182	−0.208	0.120
β_{01}^0 constant	−0.442	0.232	−0.431	0.179	−0.437	0.208	−0.394	0.142
β_{11}^0 constant	−0.256	0.179	−0.177	0.155	−0.302	0.155	−0.308	0.116
β_{10}^0 (age man − 62)/10	0.175	0.120	0.474	0.178	0.192	0.109	0.321	0.123
β_{10}^0 (age woman − 62)/10	−0.025	0.115	0.102	0.145	0.004	0.103	0.011	0.103
β_{01}^0 (age man − 62)/10	0.333	0.112	0.410	0.157	0.325	0.102	0.287	0.117
β_{01}^0 (age woman − 62)/10	−0.327	0.114	−0.236	0.140	−0.289	0.102	−0.277	0.105
β_{11}^0 (age man − 62)/10	−0.096	0.126	0.334	0.195	−0.075	0.114	0.130	0.123
β_{11}^0 (age woman − 62)/10	−0.214	0.107	0.086	0.169	−0.188	0.096	−0.075	0.108
$\ln \sigma_1^2 = \theta_m^1(\text{age}_m - 62)/10 + \theta_w^1(\text{age}_w - 62)/10$								
θ_m^1	0	–	0.843	0.364	0	–	0.708	0.363
θ_w^1	0	–	0.675	0.361	0	–	0.607	0.371
$\ln \sigma_2^2 = \theta_m^2(\text{age}_m - 62)/10 + \theta_w^2(\text{age}_w - 62)/10$								
θ_m^2	0	–	0.832	0.332	0	–	0.723	0.337
θ_w^2	0	–	0.549	0.306	0	–	0.596	0.328
$\ln \sigma_3 = \theta_m^3(\text{age}_m - 62)/10 + \theta_w^3(\text{age}_w - 62)/10$								
θ_m^3	0	–	0.280	0.361	0	–	0.281	0.351
θ_w^3	0	–	0.571	0.351	0	–	0.533	0.337
Mean log-likelihood	−3.29825		−3.28337		−3.29722		−3.28434	
Number of observations	1425		Number of couples × year observations	10392				

Note: The homoscedastic model is as described at the beginning of Section 4 with age trends in utility. The heteroscedastic model extends this to allow the variance to depend on the ages of the couple as in Eq. (13).

Furthermore we make the alternative-specific constants β_{11}^0, β_{10}^0 and β_{01}^0 linearly dependent on the age of each spouse.[23] This captures shifts in tastes for each alternative where at least one of the spouses is still in employment relative to the alternative where both spouses retire.[24] The identification of the age dependent parameters, given the implicit normalisation of the variances, is straightforward to establish.[25]

[23] The alternative-specific constants are allowed to depend on dummies for birth cohort and month of birth. We do not present these coefficients here. They are available from the authors upon request.

[24] The model presented earlier is further modified to account for the probability of death of one or both spouses (see for example Berkovec and Stern, 1991). To that effect we use life tables based on the population of couples.

[25] The introduction of age-specific parameters plays a role in the calculations we present below. However, the age-specific effects are not responsible, on their own, for the dynamic effects we document. If we set the age-specific parameters to zero, our calculations still exhibit non-constant (and distinct) patterns as the couple ages.

δ=0.95 is the value of the rate of time preference which maximises the likelihood (over the coarse grid defined earlier in footnote 22) in the homoscedastic case, while δ=0.85 maximises the likelihood in the heteroscedastic case. Our preferred set of estimates is for δ=0.85. A likelihood ratio test (LRstat= 42.408, dof=6) rejects the restriction to the homoscedastic case. The same test performed with δ=0.95 reaches the same conclusion (LRstat= 36.708, dof=6). Overall the coefficients are precisely estimated.[26,27] Our discussion focuses on the estimates obtained when δ=0.85.

Because of the multinomial nature of the model the parameters are difficult to comment upon in isolation. The average propensity to consume out of liquid wealth and current income takes a reasonable value around 0.5.

The parameter value of ρ which maximises the likelihood is close to zero, this obviously implies almost risk neutrality of the couples we analyse here. The parameter γ, which measures the expected net present value of income flows while retired takes a value of around 2 and is fairly precisely estimated. This suggests that households may apply a large discount rate to future flows of income or that households do not expect to survive retirement for very long. If households discount future flows of income at an interest rate consistent with δ=0.85 (about 17%), our estimate of γ implies that households expect the couple to survive from one period to the next with probability 0.6 (expected average life time after retirement around 2.5 years). Alternatively if the household expects a constant couple survival with probability 0.95 (an expected life time of 14 years after retirement) the estimated value of γ implies that the household discounts future income flows at the rate of 53% (i.e. consistent with an interest rate around 88%).

5.2. Goodness-of-fit

The fit of the model can be assessed by comparing the empirical distribution of retirement ages with the distribution of retirement age predicted by the model. An example of such distribution is shown in Fig. 4 for couples which enter the sample at ages 57 for the man and 55 for the woman.

The left pane of Fig. 4 represents the proportion of retirement ages for the couples in the sample. The right pane represents the predicted probabilities. The predicted probability distribution fits the peak at 60 for both spouses. The predicted distribution is however less concentrated than the observed proportions.[28] An informal (i.e. taking the estimated parameters as given) chi squared test which compares the observed and predicted distribution over a coarser lattice of retirement ages (so that each cell contain a number of observation greater than 5) rejects the fit of the predicted distribution to the observed distribution. This is not specific to this sub-sample of households. In particular the model does not fit the tail of the data. It under-predicts the peak at ages of eligibility and over-predicts later ages. This is not entirely surprising since it is not designed simply to fit the distribution of the data exactly. In that respect a naive reduced form bivariate probit model of joint retirement ages, including interactions of age dummies as explanatory variables would, according to this criterion, fit much better.

[26] The variance–covariance of the parameters we present here is obtained using outer product of the gradients to estimate the information matrix.

[27] In a previous version we estimated a more "general" specification which involves alternative specific values for β. In this case, the parameter estimates where different but not significantly so and less precisely estimated.

[28] Predicted distribution scale is 2.5 observed distribution scale. This illustrates that the model fits the observed distribution shape quite well, but fails to predict the magnitude of the observed joint retirement mode.

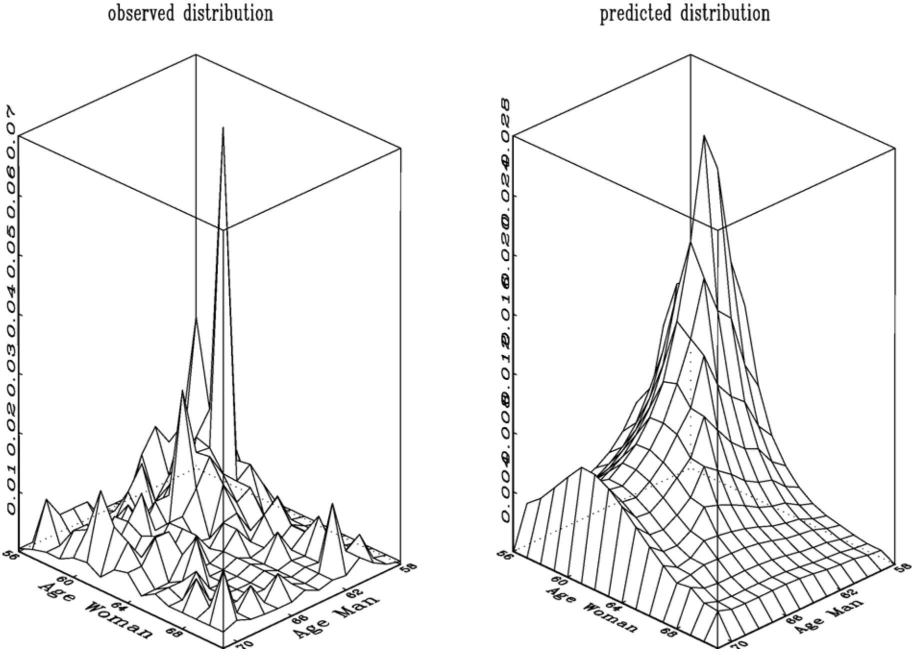

Fig. 4. Goodness-of-fit.
Note: The comparison is made for couples who are both working and are first at risk of retirement i.e. when the woman is age 55 and the man 57. Predictions are based on estimates from the heteroscedastic model with $\delta=0.85$.

5.3. Simulations

In order to illustrate the behavioural parameters that have been estimated, two simulation exercises are performed.[29] All of the simulations are based upon estimates from our preferred specification in Table 4 of the heteroscedastic model with $\delta=0.85$. Simulations are run for couples where the woman is age 55 and the man 57, i.e. both are working and first at risk of retirement. First, a 1% increase in non-work income provides participation elasticities at each age at risk. Second, a simulation of the introduction of PEW illustrates the explanatory power of the model and allows decomposition into direct and indirect programme effects within couple.

The elasticities of the probabilities of participation conditional on being in work the previous period[30] are presented for all ages in Fig. 5. These are calculated for a couple where the woman is age 55 and married to a man of 57 by simulating a 1% increase in potential net retirement income

[29] The simulations were produced for both sets of parameters with $\delta=0.85$ or $\delta=0.95$. The conclusions we reach in the text are not qualitatively changed if we use either set of parameter estimates. In the course of our analysis we estimated alternative models, for example setting all parameters involving the age of the spouses to zero or allowing parameter ρ or β to depend on the alternative, which can produce qualitatively different results. These alternative models however are always rejected on the basis of a Likelihood Ratio test.

[30] Conditional participation probabilities are defined as the probability that a member of the household still participates at age t given that they have not retired at age $t-1$. These are easily calculated from the probability of each spouse retiring at different ages. Note that this is only one way to represent the same information about the timing of retirement for each spouse. We could have instead chosen to calculate the elasticities of the marginal probabilities to retire at a given age. However, because of adding up, the elasticities of the marginal probabilities must change sign at least once as the couple ages. This complicates interpretation somewhat.

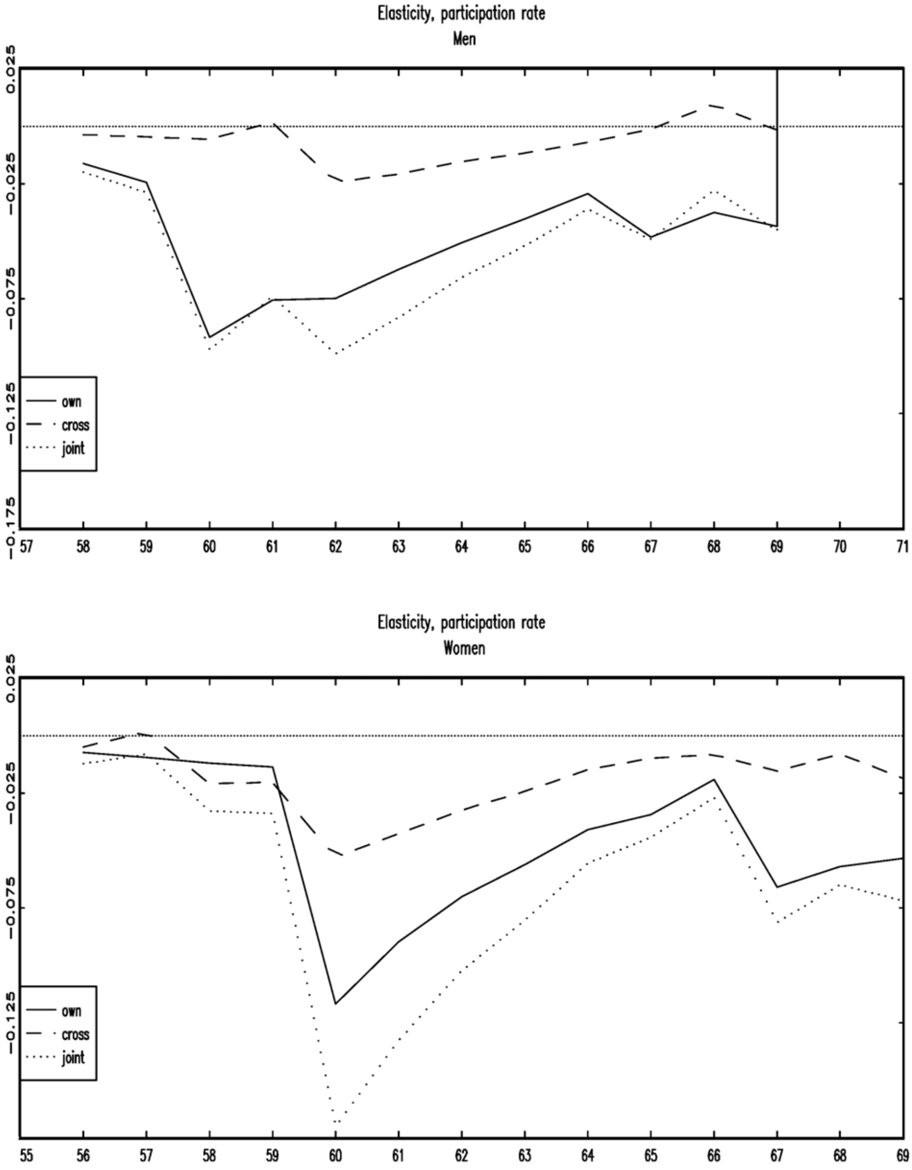

Fig. 5. Participation elasticities with respect to out-of-work income.
Note: Simulations are based on estimates from the heteroscedastic model with $\delta=0.85$ and are run on couples who are both working and both are first at risk of retirement i.e. when the woman is age 55 and the man 57.

at all future ages. The upper pane of the figure shows changes in male participation rates in response to an increase of male income only (denoted in the figure as "own"), female income only (cross), and male and female income together (joint). The *x*-axis in each pane shows the synchronised spouses' ages. The cross effect (of female income on male participation) is about one third of the own effect (of male income on male participation). As the lower pane shows, on the same vertical scale, changes in female participation rates are generally more income elastic than for males. Cross effects (of male income on female participation) are relatively larger at about half the size of the own effect (of female income on female participation). The asymmetry in the cross-income elasticities is apparent since the effect of male income on female participation is larger than the effect of female income on male participation. An important tension in couple's

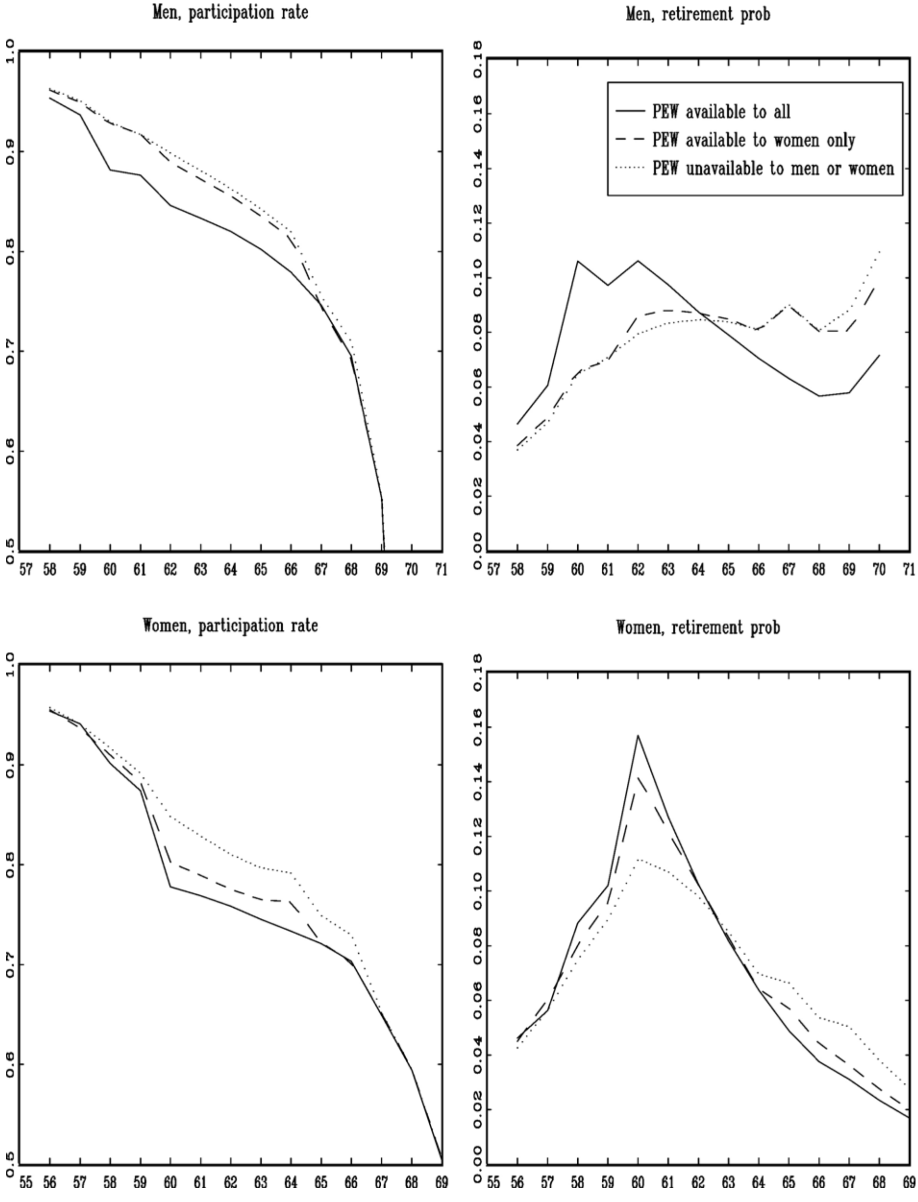

Fig. 6. Simulation of PEW introduction.
Note: Simulations are based on estimates from the heteroscedastic model with $\delta = 0.85$ and are run on couples who are both working and both are first at risk of retirement i.e. when the woman is age 55 and the man 57.

retirement studies is that the income effect at retirement works against finding complementarities in leisure. It is clear from our simulations that spousal leisures are almost always complements (since the cross elasticities are negative, i.e. an increase in one spouse's out of work income decreases the others participation probability).

The ages of 60 and 66 define the PEW eligibility window and coincide with the minimum and a (local) maximum of own income elasticities. In the case presented here where a woman is two years younger than her husband, the cross-elasticity of female income on male participation is at its minimum when the husband is aged 62, wife is 60 and her eligibility window opens.

The introduction of the PEW programme is simulated in Fig. 6. Upper (respectively, lower) panes show male (female) effects and panes to the left (right) show conditional participation (retirement) probabilities as the agent ages. Once again we consider a sub-sample of our estimation sample where the woman is 55 and the man 57. Each may or may not be potentially eligible to PEW sometime in the future, and retirement ages corresponding to this baseline are simulated and shown by the solid lines in the figure. In the first counter-factual simulation access to the PEW programme is removed for men, only women may remain eligible and predicted retirement ages are shown by the dashed lines. Finally the PEW programme is further removed for women as indicated by the dotted lines.

When the PEW retirement window opens for men or women at age 60 we observe a sizeable drop in the conditional participation probabilities (left panes). This drop persists for the duration of the window, and as expected from age 67 the conditional probabilities are virtually indistinguishable.

Total programme effects are similar for men and women with the introduction leading to higher retirement probabilities at age 60 (about 5%) as can be seen from the right panes. While this may not be a realistic policy, the intermediate case of programme availability for eligible women only, clearly illustrates direct and indirect eligibility effects. Cross-eligibility effects on participation are larger for male eligibility on female participation than for female eligibility on male participation. This should be expected from the cross-income effects shown in Fig. 5. However decomposing cross-eligibility effects shows even greater asymmetry. Male eligibility accounted for about half of the PEW-induced fall in female participation, whereas female eligibility explained only about a fifth of the induced fall in male participation. Our model predicts that without the PEW the distribution of retirement ages for men would be less concentrated around 60. The distribution of retirement age will shift to the right. For women, although the modal retirement age is still around 60, it is less pronounced.

6. Conclusions

We analyse the economic determinants of joint retirement decisions for career married couples. We set up a dynamic structural model and estimate its parameters on a 1% sample of the Danish population drawn from administrative registers. In contrast to the US, Denmark has several institutional features which simplify retirement modelling: universal health insurance coverage, individual-based programme eligibility and benefit entitlement and negligible partial retirement or labour market re-entry after retiring from a career job. Also, a generous public pension programme (PEW) was introduced within our sample period. Pre-determined eligibility criteria are used to break the simultaneity between work and potential non-work income flows and empirically identify the model.

Estimated values of preference parameters are plausible with an average propensity to consume out of liquid wealth and current income of 0.5, a rate of time preference 0.85 and a coefficient of relative risk aversion of zero. Simulations show that the model fits the retirement effects of PEW introduction reasonably well. We decompose the cross-eligibility effects of spouses, and we show that male eligibility accounts for a much larger proportion of the programme-induced fall in female participation than female eligibility accounts for male.

Female participation is more responsive to income changes than male participation. This is true for both own-income and cross-income elasticities. Spouse's leisures are found to be almost always complements. Furthermore, there is significant asymmetry in cross-income elasticities, with female participation responding much more to male income than male participation responds to female income.

Acknowledgements

The Danish Social Sciences Research Council (24-02-0064) and UK Economic and Social Research Council (RES-000-27-0200) provided funding. This paper has benefited from comments received from workshop participants in CAM Copenhagen, RES Nottingham and TAPES Uppsala, especially Martin Browning, Jim Poterba, Arthur Van Soest, Ian Walker and two anonymous referees. The usual disclaimer applies.

References

An, M., Christensen, B.J., Datta Gupta, N., 2004. Multivariate mixed proportional hazard modelling of the joint retirement of married couples. Journal of Applied Econometrics 19 (6), 687–704.

Berkovec, J., Stern, S., 1991. Job exit behavior of older men. Econometrica 59 (1), 189–210.

Blau, D., 1997. Social security and the labor supply of older married couples. Labour Economics 4, 373–418.

Blau, D., 1998. Labor force dynamics of older married couples. Journal of Labor Economics 16 (3), 595–629.

Blau, D., Gilleskie, D., 2004. Health Insurance And Retirement Of Married Couples. University of North Carolina at Chapel Hill Mimeo.

Browning, M., Leth-Petersen, S., 2003. Imputing consumption from income and wealth information. Economic Journal 113 (488), F282–F301.

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

Gustman, A., Steinmeier, T., 2002. Social Security, Pensions and Retirement Behavior Within The Family. NBER Working Paper, vol. 8772.

Gustman, A., Steinmeier, T., 2004. Personal Accounts And Family Retirement. NBER Working Paper, vol. 10305.

Leth-Sørensen, S., 1993. IDA: An Integrated Database for Labour Market Research. Statistics Denmark, Copenhagen.

Maestas, N., 2001. Labor, Love and Leisure: Complementarity And The Timing Of Retirement By Working Couples. Berkeley Economics Department Mimeo.

Mastrogiacomo, M., Alessie, R., Lindeboom, M., 2004. Retirement behaviour of Dutch elderly households. Journal of Applied Econometrics 19 (6), 777–794.

Michaud, P.-C., Vermuelen, F., 2004. A Collective Retirement Model: Identification And Estimation In The Presence Of Externalities. IZA Discussion Paper, vol. 1294.

Pedersen, P.J., Smith, N., 1996. A duration analysis of the decision to retire early. In: Wadensjö, E. (Ed.), The Nordic Labour Markets in the 1990's. North Holland, Amsterdam.

Rust, J., Phelan, C., 1997. How social security and medicare affect retirement in a world of incomplete markets. Econometrica 65 (4), 781–832.

Sedillot, B., Walraet, E., 2002. La Cessation D'activité Au Sein Des Couples: y a-t-il Interdependence Des Choix. Economie et Statistique 357–358, 80–102.

Seiersen, A., 2003. Portability of pension rights: The Danish experience. Danish Insurance Association mimeo.

Train, K, 2003. Discrete Choice Methods with Simulation. Cambridge University Press.

Van der Klaauw, W., Wolpin, K., 2005. Social Security, Pensions And The Savings And Retirement Behaviour Of Households, mimeo, University of North Carolina.