TRANSITIONS OUT OF AND BACK TO EMPLOYMENT AMONG OLDER MEN AND WOMEN IN THE UK

David Haardt

Institute for Social and Economic Research
University of Essex

ISER Working Paper
2006-20

Acknowledgement:

E-mail: damhaa@essex.ac.uk. This will be the first paper of my PhD thesis. It would not have been possible without the outstanding supervision which I have been receiving from Stephen P. Jenkins. Thanks are also due to my second supervisor John F. Ermisch and to David M. Blau, Marco Francesconi, James W. Hardin, Cheti Nicoletti, and Steve Pudney for helpful suggestions. I am grateful to the data depositors of the BHPS (ISER, University of Essex) and of the FES (Office for National Statistics), and to the UK Data Archive, University of Essex, for providing access to the data. Finally, I would like to thank the Austrian Academy of Sciences, the Economic & Social Research Council, the University of Essex, and the Provincial Government of Upper Austria for funding. The author alone is responsible for errors and opinions.

Readers wishing to cite this document are asked to use the following form of words:

> Haardt, David (May 2006) 'Transitions Out Of and Back To Employment Among Older Men and Women in the UK', ISER Working Paper 2006-20. Colchester: University of Essex.

The on-line version of this working paper can be found at http://www.iser.essex.ac.uk/pubs/workpaps/

The Institute for Social and Economic Research (ISER) specialises in the production and analysis of longitudinal data. ISER incorporates

- MISOC (the ESRC Research Centre on Micro-social Change), an international centre for research into the lifecourse, and
- ULSC (the ESRC UK Longitudinal Studies Centre), a national resource centre to promote longitudinal surveys and longitudinal research.

The support of both the Economic and Social Research Council (ESRC) and the University of Essex is gratefully acknowledged. The work reported in this paper is part of the scientific programme of the Institute for Social and Economic Research.

Institute for Social and Economic Research, University of Essex, Wivenhoe Park,
Colchester. Essex CO4 3SQ UK
Telephone: +44 (0) 1206 872957 Fax: +44 (0) 1206 873151 E-mail: iser@essex.ac.uk
Website: http://www.iser.essex.ac.uk

ABSTRACT

This paper analyses the labour market transitions of older men and women using data from the British Household Panel Survey (BHPS). I find large peaks in exit rates out of employment at ages 60 (women) and 65 (both sexes) which occur in the exact birthday month. This suggests that pension schemes have strong incentive effects. Discrete-time hazard regression analysis shows that benefits and health status are the two most important determinants of retirement, with effects that are larger than found in previous studies for British and US men. When modelling unobserved heterogeneity I find that the share of 'movers' between work and non-work is twice as high among women as among men.

NON-TECHNICAL SUMMARY

In the United Kingdom, as in most industrialised countries, life expectancy has risen substantially while the number of workers per pensioner and work income as a share of older people's total income have fallen considerably. These developments manifest themselves on the level of the individual pensioner in the form of inadequate retirement income and healthcare entitlements and have led to a great interest in the understanding of older people's labour market participation and, in particular, their retirement behaviour.

When does somebody retire and why? Which retirees might return to work and why? How are these transitions affected by how much one could earn in the labour market and by how much income one could have without working? I address these questions, analysing the labour market transitions of older men and women using data from the British Household Panel Survey (BHPS) for 1990-2004. The individuals considered are aged 40 to 70 years.

First, I look at the risk for a worker of exiting employment within the next twelve months and at the risk for a non-worker of resuming employment within the next twelve months, considering differences in risks by sex as well as by other characteristics. I find that the probability of employed men to exit employment within the next year is highest at age 65, and that of employed women at age 60, i.e., in both cases at state pension age. In fact, these peaks mostly occur in the exact birthday month: a man who is employed directly before his 65th birthday has a 45.38% risk of stopping to work within one month. When looking at a woman, the exit risks are 11.78% in the month of her 60th and 19.49% in the month of her 65th birthday. This suggests that pension schemes have strong incentive effects. The return to work probability for men is approximately twice as high as that for women until about age 57 when returns become rare for both sexes. I also find that men and women who have ever contributed to, or received money from, an occupational pension scheme are much more likely to return to work than those who have not.

When I considered differences in these risks among men and women using multivariate regression models, there were several clear findings which are to be interpreted 'other things being equal'. First, there are no significant differences between single people and those with a partner in the probability of exiting or re-entering employment. Second, people who gained a lot of employment experience prior to their 40th birthday have a lower risk of exiting employment and higher chances of re-entering employment between ages 40 and 70, but the effects for women are not that large. Third, women generally make fewer exits or returns if regional unemployment is high. Fourth, people in poor or very poor health have a much higher risk of exiting employment and a much lower risk of returning to work, with effects that are larger than found in previous studies on UK and US men. Fifth, for women, occupational pension membership is associated with a higher labour market participation while for men the effects are opposite or not significant. Sixth, women who have ever had a child have a lower risk of exiting employment, which is contrary to what one may have expected. Seventh, as in earlier literature, earning capacity does not seem to be an important determinant of labour market exit or entry, but the benefit entitlement has important effects on the return probability: the higher the benefit entitlement, the lower the probability to return to the labour market. These effects are again larger than found in earlier studies. Summing up, health status and benefit entitlement are the two most important determinants of retirement.

Finally, I also find that among women the share of 'movers', i.e., of individuals with a higher risk of making transitions out of and back to employment given the same values of the observed explanatory variables, is twice as high as among men.

1 Introduction

In most industrialised countries, the average life expectancy has risen substantially while the average retirement age has often fallen. This has increased the number of pensioners each worker has to finance through public pension systems. The United Kingdom is no exception: while today there are three working people for each person aged 65+, this value is projected to fall to two by 2030 (OECD 2000: 213–214). This has led to gaps in government budgets which, in turn, have led to important changes in the provision of pensions and healthcare: the share of public pension income has fallen in many countries; the financing need for healthcare, in particular for long-term care, will rise sharply.

In the United Kingdom, the level of net social transfers as a share of older people's total income has remained relatively constant over the last decades (although on a lower level than in most countries in Continental Europe) while the share of work income has fallen and the share of capital income risen (OECD 2000: 44). The GDP share of publicly financed long-term care is projected to rise by 15–20% from 2000 to 2020 (OECD 2000: 64).

A particular feature of the pension provision in the United Kingdom is the important role which occupational pension (OP) schemes play: capital income (which includes income from OP schemes) contributed less than 30% of pension income in the mid-70s, but more than 40% in the mid-90s (OECD 2000: 44).[1] In my estimation sample from the British Household Panel Survey (BHPS), 65% of men and 46% of women had ever been members of an OP scheme.

These developments manifest themselves on the level of the individual pensioner in the form of people without adequate retirement income and healthcare entitlement, whether public or private, and have led to a great interest in the understanding of older people's labour market participation and, in particular, their retirement behaviour. When does somebody retire and why? Who might return to work and why? How are these transitions affected by how much one could earn in the labour market and by how much income one could have without working? I address these questions.

There are many descriptive papers and policy studies of these issues. There are also two types of more analytical approaches: structural models of the retirement decision and reduced-form hazard regression models which allow for more freedom in the functional form of agents' underlying economic decision-making.

Lazear (1986) is a valuable survey of structural retirement models. There are (a) purely theoretical microeconomic models, (b) econometric implementations of life-cycle models, and (c) dynamic programming models whose parameters can (or could) be fitted

[1] For a more detailed discussion of the UK pension system and a description of its evolution, see Emmerson (2003).

to real-world data. Based on assumptions about individuals' utility functions, these models allow analyses of the effects of exogenous factors on retirement behaviour. For example, Kingston (2000) shows in a theoretical model of the optimal choice of the date of retirement that people with higher assets, lower wages, higher disutility of effort, and lower life expectancy will retire earlier than others.

Hazard regression models (also called duration, survival, or time-to-event models) do not impose assumptions about utility functions, etc. They are used to model discrete choices along a time axis. In the case of retirement, this choice is typically between work and non-work, but can also include a bigger variety of labour market states. While structural retirement models are usually limited to 'the exit' from the labour market (considered to be a one-off event without any option to return to the labour market), many hazard regression models also model the return to the labour market. This facilitates a more precise analysis of how different factors affect retirement behaviour, separating exit behaviour from return behaviour. In the case of the UK, taking the return into account is particularly important since previous research has highlighted how rarely older people return to work (cf. Smeaton and McKay 2003 and Humphrey *et al.* 2003) which makes it important to understand when it does occur.

Research on retirement timing using time-to-event analysis has not been very popular, partly because detailed longitudinal data are not available for many countries. Examples include Blau (1994) about the US (RHS) and Mastrobuoni (2000) about Germany (GSOEP) and Italy (SHIW). Meghir and Whitehouse (1997), using the Retirement Survey (RS), is the key article about the UK.[2] All emphasise the importance of pension benefits in determining when people retire, especially through their (decreasing) effect on the return probability.

This paper attempts to build upon Meghir and Whitehouse's research, extending our knowledge of older men and women's labour market transitions in the UK.

First, while Meghir and Whitehouse, like most previous research, only analyse men, I compare men and women. This will enable some interesting comparisons across sexes of the effects of the explanatory variables. While men can start to receive their Basic State Pension starting from their 65^{th} birthday, women can do so starting from their 60^{th} birthday. It is interesting to see whether this difference in the state pension age has any effects on older men and women's labour market behaviour, and if so, which. A related point is that I analyse the exact timing of transitions in more detail than previously done.

Second, while Meghir and Whitehouse do not include members of occupational pension schemes in their regressions (thus excluding 60% of their sample), arguing that

[2]Since I will refer to Blau (1994) and Meghir and Whitehouse (1997) often, I will not mention the publication year of these two articles every time.

occupational pension schemes were too heterogenous to be modelled without detailed data about their incentive structures, I try to keep my sample restriction criteria as minimal as possible to be able to analyse this important feature of the UK pension system. I compare members and non-members rather than examining the impacts of detailed differences in institutional features. Even such a simple approach produces some interesting results.

Third, I employ a wider set of explanatory variables. Since it is well known that health is an important determinant of retirement, I include more detailed health variables than previous studies. I examine the effects of work experience gained in younger years, which Meghir and Whitehouse do not. I also include more detailed measures of marital status and unemployment. A discussion of how I improve upon previous specifications is provided in Section 3.

Fourth and related to the previous point, I use a more up-to-date data set (Meghir and Whitehouse analyse the time period from 1968 to 1989) which is valuable in and of itself since the labour market behaviour of older Britons has changed a lot over the last decades (cf. Campbell 1999).

The data which I am using come from the British Household Panel Survey (BHPS), a longitudinal general-purpose household survey. The first wave was carried out in 1991 and there are currently 13 waves of data available. The maturity of the panel enables us to follow people over a relatively long period in their life. The BHPS has the advantage of being a rich, current data set with high data quality.[3]

Section 2 of this paper describes the modelling framework used; Section 3 the data. Section 4 presents non-parametric empirical transition hazard rates as well as the results of the hazard regression models, including predicted hazard rates and predicted median transition ages. Section 5 draws some conclusions.

2 Modelling Framework

2.1 Definition of Labour Market States

I use a discrete-time survival analysis model with two hazards. I study not only the exit from employment (E) to non-employment (N), but also the return from N to E which is important for reasons which I have just mentioned. Choosing survival analysis (more

[3]There is now also the English Longitudinal Study of Ageing (ELSA), the first two waves of which have been carried out in 2002 and 2004. There are plans for a one-off life history interview in wave 3 (*i.e.* in 2006), subject to funding. ELSA is a specialist survey geared towards many of the aspects of this paper, but I do not use it since it is still in an early stage in terms of the maturity of the panel dimension.

specifically, a mixed proportional hazards model), a reduced-form approach, enables me to study a host of potential factors—occupational pension membership, the potential incomes in and out of work, demographic and other personal characteristics, work-related variables, etc.—simultaneously without having to impose a lot of structural assumptions.

I treat everybody who reports positive working hours as employed and those who report zero hours as non-employed. The corresponding transition types are then called EN (if you change from employment to non-employment from one month to the next) and NE (*vice versa*). Men and women are analysed separately.

Using a more aggregative non-employment state (N) rather than distinguishing between separate labour market states such as retired, unemployed, and out of the labour force is in line with earlier literature since the definitions of 'retirement' and 'unemployed' are ambiguous. As age increases, an increasing proportion of those with zero hours (*i.e.* of those in my state N) will also call themselves 'retired'.

Treating all those with zero work hours as non-employed follows Blau. By contrast, Peracchi and Welch (1994), who match a number of year-to-year transitions from the US Current Population Survey (CPS) to analyse the labour market transitions of older men and women in a simpler methodological framework, regard the unemployed, *i.e.* those actively looking for a job, as full-time or part-time employed depending on their previous job. I justify my approach with the fact that the focus of my paper is really on genuine labour market activity rather than on search or the receipt of unemployment benefits. Furthermore, finding good measures for whether somebody is 'actively searching' is rather difficult.

There is a literature on the question whether unemployment and out of the labour force are different from each other. Both Flinn and Heckman (1983) and Jones and Riddell (1999) find that they are, with large differences between the states in the return probabilities to employment. Nonetheless, in my case, the focus is really on employment as opposed to everything else (since employment is associated with labour market earnings, while unemployment and out of the labour force are not), so I decided to merge all non-employment states in accordance with what Blau and Meghir and Whitehouse did. A multi-state analysis would of course be possible but the parsimonious focus on just two labour market states will make it easier to identify the principal mechanisms.

Figure 1 plots the proportions of men and women working during the BHPS panel period (1990–2004) using the definitions just defined. Comparing this to Blau (1994: 122) reveals striking differences between the UK today and the US in the 1960s and 1970s: he reports that less than 10% of men do not work at age 55, and only approximately 45% at age 65, while in Britain the corresponding figures are 28% and 73%, respectively. In fact the labour market participation among US men at age 55 is roughly comparable to

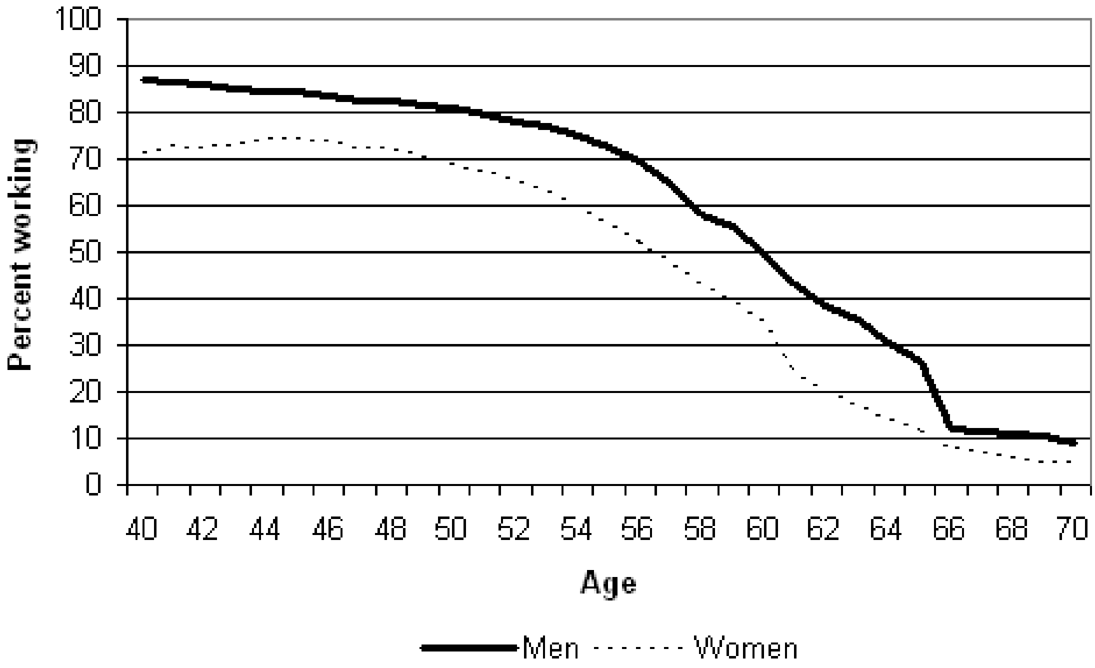

Figure 1: Shares of men and women working

that among UK men at age 40, for the labour market withdrawal process starts much earlier in the UK than in the US, and the non-employment share is higher at any given age. For women (whom Blau does not analyse) my results are even more dramatic (44% and 88%, respectively). In terms of labour market participation, the UK today is indeed very different from the 1960s and the 1970s in the US.

2.2 The Hazard Regression Model

I assume that the transitions out of and back to employment follow a proportional hazards model. Since my data are interval-censored, I use a complementary log-log (or cloglog) model for the hazard regressions.

A proportional hazards model implies that the duration profile of the hazard is the same for everybody, with the explanatory variables shifting this profile upwards or downwards:

$$\chi(t, \mathbf{x}) = \chi_0(t) \exp(\beta' \mathbf{x}), \tag{1}$$

where χ is the underlying continuous-time hazard rate. The proportional hazards property is intuitively appealing and also has analytical advantages.

I have monthly observations about each person's labour market status although people may make their transitions on a daily basis. My data are therefore interval-censored with monthly observations where 'interval-censored' means that although the actual transition process is continuous (or discrete with smaller time units than observed in the data), the data are grouped into intervals. Prentice and Gloeckler (1978) show that the cloglog model is the interval-censored discrete-time equivalent of a continuous-time model with the proportional hazards assumption, with the same coefficient vector β as in the continuous-time case.

A result of the above-mentioned analytical advantages of a proportional hazards model is that each coefficient can be transformed to a hazard ratio which is easier to interpret than the underlying coefficient:

$$HR_{x_k} = \frac{\chi(x_k = a)}{\chi(x_k = a - 1)} = \exp(\beta_{x_k}). \tag{2}$$

This is the relative risk associated with a one-unit change in the value of the explanatory variable x_k, all other elements of $\mathbf{x}$ held constant. The usual null hypothesis of $\beta_{x_k} = 0$ is of course equivalent to the null hypothesis of $HR_{x_k} = 1$ since $\exp(0) = 1$. In the discussion of my regression results, I will talk about hazard ratios most of the time.

The survivor function is

$$S(i, j) = \prod_{k=i}^{j} (1 - h_k), \tag{3}$$

i.e. the fraction of people remaining in the origin state at time j out of all those who were in the origin state at time i (*i.e.* $i < j$). I am using h for the interval-specific (*i.e.* discrete) hazard rate.

If there were no unobserved heterogeneity, one could simply estimate two separate discrete-choice regressions, one for the exit from employment (each observation then corresponding to a month spent in employment), and one for the return to employment (each observation representing a month spent in non-employment).

It is however questionable whether all individuals with the same vector of observed explanatory variables face the same expected hazard of making a transition in the labour market. In particular, it seems to be reasonable to assume that there are some individuals who, due to unobservable factors such as 'labour market attachment', are more or less likely than others to experience certain transitions. Ignoring unobserved heterogeneity can lead to various biases (cf. Jenkins 2004: 79–87). In my case, unobserved heterogeneity means that I have to take multiple spells into account (one person may make either transition several times) as well as allow for correlated error terms across the two transitions, estimating them jointly.

To model unobserved heterogeneity, I am using the non-parametric approach proposed by Heckman and Singer (1984). One could also assume a bivariate normal heterogeneity distribution; however, Abbring and van den Berg (2003) show that in duration models the heterogeneity distribution usually converges to a Gamma distribution. Also Blau, and Meghir and Whitehouse, use the approach of Heckman and Singer.

Heckman and Singer present a way of fitting a non-parametric distribution of unobserved heterogeneity where the position and probability of each mass point is determined from the data themselves, conditional on the number of mass points chosen by the researcher (typically, one starts off with two mass points and can then try to increase their number, although convergence is usually only achieved with a small number of mass points). Each mass point can be interpreted as an estimated fixed effect for a group of people who share a certain unobserved *ceteris paribus* propensity to make the corresponding transition, and the probability of each mass point as the estimated share of the sample with this specific propensity.

With two mass points, there are $2 \times 2 = 4$ types of people in this model. However, the model never converged when allowing for four groups, or when constraining only one of the four mass point probabilities to zero. I therefore constrained two mass point probabilities to zero, keeping the remaining two.

My model therefore reduces to the following. There are two types of people, group A with a corresponding probability π, and group B with the probability $1 - \pi$. Formally, the transition model for group A can be written as

$$h_{i,EN} = 1 - \exp(-\exp(\beta' \mathbf{x_i})); \tag{4}$$

$$h_{i,NE} = 1 - \exp(-\exp(\gamma' \mathbf{z_i} + \theta_{NE})); \tag{5}$$

and for group B as

$$h_{i,EN} = 1 - \exp(-\exp(\beta' \mathbf{x_i} + \theta_{EN})); \tag{6}$$

$$h_{i,NE} = 1 - \exp(-\exp(\gamma' \mathbf{z_i})), \tag{7}$$

where $\mathbf{x_i}$ and $\mathbf{z_i}$ are the vectors of explanatory variables (*i.e.* demographic and other personal characteristics as well as variables relating to work, including incentives) and β and γ the corresponding parameter vectors. The θ terms are the mass points of the estimated unobserved heterogeneity distribution; if there were no unobserved heterogeneity, they would disappear.[4]

[4] An alternative equivalent notation would be $\text{cloglog}(h_{i,EN}) = \beta' \mathbf{x_i}$ and $\text{cloglog}(h_{i,NE}) = \gamma' \mathbf{z_i} + \theta_{NE}$ for the transition model for group A and $\text{cloglog}(h_{i,EN}) = \beta' \mathbf{x_i} + \theta_{EN}$ and $\text{cloglog}(h_{i,NE}) = \gamma' \mathbf{z_i}$ for group B.

The signs of the two mass points tell us about the two groups. If θ_{EN} and θ_{NE} are both positive, group A exhibits higher labour market attachment (lower exit risk, higher return probability) and group B lower labour market attachment (higher exit risk, lower return probability). If θ_{EN} and θ_{NE} are both negative, we will have the same two groups, but *vice versa*. If θ_{EN} is positive and θ_{NE} negative, group A are stayers (lower exit risk, lower return probability) and group B movers (higher exit risk, higher return probability). If θ_{EN} is negative and θ_{NE} positive, we will again have the same two groups *vice versa*. Other combinations of these four groups are not possible due to the choice of the equations in which the two mass points appear.[5]

A further issue arising when estimating a transition model with unobserved heterogeneity is the potential non-randomness of the labour market status of a person's first spell. In my case, whether the first observation of a person in the sample is in employment or in non-employment may be non-random, because belonging to either one of the two groups of people just mentioned will also influence the probability or being employed or non-employed when first observed. Ignoring this may bias the estimates. Typically, this problem is tackled by including an additional equation in the model which models the probability of being employed when first observed, allowing for correlation in unobservables across equations. The above-mentioned model did however not converge when adding such an initial conditions equation. This may be due to the fact that the variable summarising somebody's work history between ages 15 and 40 is already a powerful predictor of the labour market status during a person's first month in the sample. Experiments which I carried out with a univariate probit model showed that the employment experience variable can indeed explain 25% of men's variance of the labour market status when first observed and 29% of women's.

3 Data, Sample Selection Criteria, and Variables

3.1 The Data

The main data source for this paper is the British Household Panel Survey (BHPS), a longitudinal survey of households containing a wealth of socio-demographic information as well as information on many economic variables. The individuals of a representative sample of 5,500 British households were interviewed in 1991 for the first time and have since been followed, with data from 13 waves (annual interviews) currently available (Taylor 2005). Apart from being a rich general-purpose household survey, the key advantages

[5]The general version of the model with all four groups, *i.e.* without constraining any mass point probabilities, has four rather than just two θ parameters, one for each equation in the above scheme.

of the BHPS over other surveys are its data quality and its up-to-dateness. While Meghir and Whitehouse, using the Retirement Survey (RS), had 1989 as the last year in their data set, I am able to include data spanning up to 2004 (since 6.4% of wave-13 interviews were carried out in the early months of 2004).

The BHPS provides us with a large number of individuals and spells: the final sample which I use for my regressions contains 8,361 people (thereof 3,828 men and 4,533 women) and 14,412 spells (thereof 6,543 men's and 7,869 women's), or slightly less than 700,000 person-months. Meghir and Whitehouse (1997: 330) analyse 641 men or 2,479 spells; Blau (1994: 121) 7,157 men or 16,385 spells.[6]

Since information on earnings and benefits (and also on wealth, which I am not using in this paper) is not so detailed in the BHPS, I use an auxiliary data source, the Family Expenditure Survey (FES). The FES is a series of repeat cross-sections, with data being available from 1961. It contains detailed information about income from different sources as well as about the receipt of different types of benefits. Since 1968, the FES has been carried out each year without any gaps. The last FES was carried out in 2000–01 after which it merged with the National Food Survey (NFS) to become the Expenditure and Food Survey (EFS), first carried out in 2001.

I follow Meghir and Whitehouse who run earnings and benefits regressions using the FES and then use these coefficients to predict lifecycle profiles of earnings and benefits for RS respondents. Arellano and Meghir (1992) provide a detailed treatment of this methodology. As always when using predicted variables as explanatory variables, the hazard regression models using the such 'augmented' RS have to take generated regressor bias into account; otherwise, the standard errors would be underestimated. Pagan (1984) offers a comprehensive treatment of this issue. In my case the prediction of earnings and benefits is carried out for the respondents of the BHPS, and therefore with all the above-mentioned advantages. I am using the corrected standard error estimator of Murphy and Topel (1985).

3.2 Sample Selection Criteria

The number of sample selection criteria is kept as low as possible, and is the same for the presentation of empirical hazard rates as well as for my regressions.

First, I restrict the age range under consideration since I am only interested in older

[6]The differences in the number of spells per person arise from differences in the definition of what a 'spell' is. In my data, a spell can be either employment or non-employment, in Blau it can be full-time employment, part-time employment, or non-employment. Meghir and Whitehouse do not provide detailed information about how they define a spell but looking at Disney, Meghir, and Whitehouse (1994) suggests that their large number of spells is because they regard job changes as separate spells, even though they do not analyse these job-to-job transitions.

people. I start the observation window at age 40, disregarding all previous months. This is similar to Meghir and Whitehouse although they also disregarded the job or unemployment spell which was underway at a person's 40th birthday. Blau considered the spell underway at a person's 55th birthday as well as all subsequent spells. His higher age bound, 55 rather than 40, can be explained by the different labour market behaviour of older men in the US *versus* the UK, as already discussed. I censor spells at age 70 since labour market transitions become rare beyond that age. Still, this allows me to take a look at five years beyond men's state pension age. Meghir and Whitehouse censor at age 65; Blau at age 73.

Second, there is a small number of spells (less than 5%) which start before people are born or when people are still very young. Not correcting this could bias the duration variable (some women who never worked report being out of the labour force since birth), and it would bias the variable which I use to summarise a person's work history up to age 40. These erroneously and/or illogically recorded spells are a genuine artefact of the retrospective employment histories in the BHPS and not related to the data organisation procedures I used. I edit all such cases to start with the person's 15th birthday. My employment experience variable therefore summarises a person's work history from age 15 to age 40. All of this also affects a few men who already worked before their 15th birthday.

Third, I drop people whose birth year and/or month is missing. Thanks to the high data quality of the BHPS, this criterion drops very few observations (less than 0.1% of spells). The rationale for dropping people whose birth month is missing, rather than generating a (pseudo-)random number with which I have also experimented, is that I want to analyse whether certain transitions are unusually often observed directly after certain birthdays (so-called 'birthday effects'), pointing towards the importance of incentives within the pension system. Generating missing birth months randomly could blur the analysis of birthday effects.

Fourth, I only analyse transitions which occur in or after September 1990 (first wave minus one year). In other words, I do not analyse transitions which are recorded in the retrospective work histories. This needs some explanation. In addition to the question about last year's employment history which is asked in each BHPS wave, BHPS respondents were also asked to provide a full retrospective work history (as well as other life histories) in waves 2 and 3. As already mentioned, I do use this information to generate a work-history variable which summarises one's labour market attachment between ages 15 and 40, but I do not use the retrospective months for the analysis of transitions.[7] There

[7]As far as the detailed procedure is concerned, I first use two Stata programmes written by Bardasi (2002) to integrate the retrospective information from waves 2 and 3 into the genuine annual panel information from waves 1 to 13. After generating the work history variable, I then delete all months

is a substantial literature on survey recall bias in economics, sociology, and statistics (cf. for instance Elias 1997 and Paull 2002 using the BHPS; Morgenstern and Barrett 1974 for a more general treatment) which suggests that especially for the lifespan period under consideration in the present paper, and particularly among older people, there are non-negligible problems with the precision of retrospective data. Furthermore, and even more importantly, retrospective data in the BHPS are limited to only a couple of variables. I would for instance not be able to use health status as an explanatory variable if I used retrospective transitions data.

Fifth and last, I have to drop people for whom I do not observe any work history data before age 40. This is because I cannot generate a meaningful employment experience variable for them. Experiments with a dummy variable for missing work history were unsuccessful. This criterion drops a bit less than one third of the remaining person-months. Most of the people dropped here are those who were not present in waves 2 and 3 since this is where the retrospective life histories were included in the interview, or who did not answer these questions fully. I do not impute this variable since it is crucial for my analysis. I have carried out some tabulations showing that the average characteristics of people dropped are similar to those of people kept which should give us some confidence.

Summing up, I look at all transitions taking place after August 1990 made by people aged 40 to 70 for whom I observe a retrospective work history.

Self-employed people are kept in the sample since self-employment is an important mode of employment among UK males and growing in its importance among older Britons. It is analysed just like being an employee. I do not use a dummy variable for self-employment for endogeneity reasons.

3.3 Variables

3.3.1 Dependent Variable

I am interested in the length of employment and non-employment spells. I operationalise this by looking at month-to-month changes in the labour market status. I therefore look at EN and NE transitions as mentioned earlier on.

Similarly to Blau, I fill labour market status gaps of a length of up to two years, assuming midway changes (if the gap length is odd, I create a pseudo-random variable to decide with which information to fill the middle month). Three examples follow. If somebody is observed working in August 1995 and not working in November 1997, the gap between these months is not filled since the gap length exceeds 24 months. If somebody is observed working in August 1995 and not working in November 1995, 'working' is

earlier than September 1990.

imputed for September and 'not working' for October (even gap length). If somebody is observed working in August 1995 and not working in October 1995, a pseudo-random number is generated to decide whether to impute 'working' or 'not working' for September (middle month of a gap with odd length). The first labour market status observed is used to fill all previous months of a person for which the labour market status is missing; the last one for all subsequent missing months, again up to a gap length of 24 months. An example: If somebody enters the sample in August 1994, has twelve months for which the labour market status is missing, and is then observed working in August 1995, the months from August 1994 to July 1995 are imputed with 'working'.

3.3.2 Explanatory Variables

Although I am estimating reduced-form models of older people's transitions out of and back to employment, job search models which have attracted a lot of attention in the unemployment literature since the 1970s can be useful in motivating the explanatory variables used. Job search models typically explain unemployment durations and reservation wages using individuals who maximise their utility over an infinite horizon, conditional on the wage distribution. Neumann (1999) gives an excellent survey of empirical implementations of job search models using survival analysis.

Much of the job search literature has focused on duration dependence and unemployment insurance, so it is natural to have elapsed spell duration and the benefit entitlement among my explanatory variables. Also age should be included since we are interested in people who are approaching the end of their working phase. Some measure of the unemployment rate will enable us to evaluate the arrival rate of new jobs. We need to find something like the wage distribution used in job search models to measure employment prospects. Similarly, we should also find some pre-determined measure of individuals' past work experience since this will have an impact on human capital as well as measuring attachment to the labour force. Demographic variables related to household/family circumstances could also be interesting, although we have to be cautious with respect to potential endogeneities. Lastly, we should include information about each person's health for obvious reasons, and information about OP membership for reasons which I discussed in the introduction to this paper.

In what follows, I will now go on to present the specific explanatory variables used in my regression analysis.

- Elapsed spell duration (in months)
- Age (in months, minus 479)

- Age 60 dummy = 1 if month of 60^{th} birthday or the month thereafter, 0 otherwise
- Age 65 dummy = 1 if month of 65^{th} birthday or the month thereafter, 0 otherwise
- Age (in months, minus 479) × Duration interaction variable
- Marital status
- Health status
- Occupational pension dummy = 1 if ever contributed to or received money from an occupational pension, 0 otherwise
- Child dummy = 1 if ever had a child, 0 otherwise
- Employment experience share = months working between ages 15 and 40 divided by total number of months in that person's work history in that age range (using retrospective as well as panel elements)
- Regional unemployment rate
- Predicted (log) income in employment ('earning power')
- Predicted (log) income out of employment ('benefit entitlement')

The specification of **elapsed spell duration**, included in levels as well as squared and divided by 1000, is in line with the literature on older people's labour market transitions. I also carried out experiments with dummy variables for certain durations (such as one dummy variable for a one-month duration and one for a duration of more than one year, as in Meghir and Whitehouse) but they were not statistically significant at any reasonable level.

Similarly, I use **age** as well as its squared value divided by 1000. Again, this is in line with the literature. In addition, to allow for birthday effects as discussed earlier on, I include two age dummy variables for the two exit peaks at ages 60 and 65. Unfortunately, the maximum likelihood model with unobserved heterogeneity does not converge if one includes separate dummy variables for the two birthday months 60 and 65 as well as the $2 \times 2 = 4$ surrounding months. Therefore, and since some regressions which I have run without unobserved heterogeneity have shown that for women, also the coefficients of the dummy variables for the months directly after these two birthday months are statistically significant, I include two variables which are 1 if it is the month of the corresponding birthday itself or the month directly thereafter.

The coefficients of age and age squared, as well as of all the other explanatory variables, are hardly affected by whether one includes such dummy variables or not, or by which ones one includes (just two dummy variables for the 60^{th} and 65^{th} birthdays, six dummy variables when also covering the surrounding months, etc.). Without such dummy variables, these spikes caused by birthday effects would simply be completely masked.

Finally, I also include an **age × duration interaction variable** (divided by 1000). The coefficient size is small, but it turns out to be statistically significant. The coefficients of the other age and duration variables are therefore not affected very much by including this interaction term. This variable allows for some interaction between age and duration which may be reasonable in terms of economic theory.

Age and duration can be identified separately because of the multiple-spell structure of the data. If we analysed a single event which can only occur once, such as death, age and duration would simply be linear transformations of one another. In my analysis however, there are two types of events (exit and return), both of which can occur repeatedly. If somebody stops employment and makes an exit to non-employment, the duration variable will be reset to one while the age variable will of course continue to 'tick'. Therefore, it is the multiple-spell structure of the data which enables separate identification of age and duration effects.

I use four **marital status** dummy variables: divorced; married; separated; widowed. The base category is never married. Missings in this variable are treated as described for labour market status.

The **health status** variables are based upon the following BHPS question: *Please think back over the last 12 months about how your health has been. Compared to people of your own age, would you say that your health has on the whole been...* The answer *Excellent* forms my base category, and I have four dummy variables for the remaining answers *Good*, *Fair*, *Poor*, and *Very Poor*. Missings in this variable are again treated in the same way as in the labour market status variable.

The **regional unemployment rate** uses information about the claimant count rate in twelve regions of the UK obtained from National Statistics: East, East Midlands, London, North East, Northern Ireland, Northwest and Merseyside, Scotland, Southeast, Southwest, Wales, West Midlands, and Yorkshire and Humberside. I use the not seasonally adjusted across-sex time series.[8]

I also experimented with the *relative* regional unemployment rate, *i.e.* dividing the

[8] DPDD, DPAN, DPDE, DPDA, DPAV, IBWD, DPAU, DPDF, DPAQ, DPAT, DPAR, and DPAM: http://www.statistics.gov.uk/statbase/tsdintro.asp The data for one region (Northwest and Merseyside) are missing between September 1990 (the first 'BHPS month') and December 1993. I apply the time trend of all of England (time series VASS) to fill this gap, scaling the values so that my predicted value for Northwest and Merseyside for the first month in the data (January 1994) is equivalent to the actual one.

monthly claimant count rates for each region by that region's average over the time period under consideration (September 1990 to May 2004). This would eliminate the region fixed effects which might be useful since the region of residence could be deemed endogenous. The sign and statistical significance of the unemployment variable were not affected by this, and all the other variables had almost the same coefficients. I therefore went back to using the 'normal' (absolute) regional unemployment rate since this variable is easier to interpret.

'Earning power' and 'benefit entitlement' are predicted values of each person's individual incomes in and out of the labour market. They include all types of incomes applicable, but for simplicity I shall refer to them as earning power and benefit entitlement. They have a similar purpose as the wage distribution and replacement ratio variables typically used in the job search literature, as already mentioned.

I use predicted rather than actual values for two reasons: first, to overcome potential endogeneity; second, because at one point in time a person can only be either in employment or in non-employment. In line with the literature, I use *individual* rather than household or family income, so ignoring household effects. An interesting aspect would be to test in how far predicted household income plays a role for these transitions. Because of endogeneities, this might be particularly important when analysing the joint labour market behaviour of older couples which has not yet been done with UK data.

I use the FES rather than the BHPS to predict these two variables because information on earnings and especially on benefits is much more detailed in the FES than in the BHPS, which should make them more reliable. I drop FES respondents who are younger than 40 or older than 70, *i.e.* I use the same age group as for my transition regressions. I adjust for inflation using the Retail Price Index of National Statistics.[9] My two income variables are measured in January 2004 prices.

Separately for men and for women, I regress the natural logarithm of employees' total income on calendar time (in months since 1900), age, and occupational dummies to estimate the predicted earning power, and the natural logarithm of non-employed people's total income on calendar time and age to estimate the predicted benefit entitlement.[10] I run these auxiliary regressions using the natural logarithm of the two income variables as the dependent variables rather than levels to yield a better fit. I present the results of these regressions in Table 4 in the Appendix.

One problem might be a correlation of the variance matrices of the earnings and

[9] http://www.statistics.gov.uk/STATBASE/tsdataset.asp?vlnk=229 The exact time series used is CHAW which is a monthly index number which includes all items, *i.e.* also housing, indirect taxes, and mortgage interest payments. I change the base month from January 1987 to January 2004.

[10] There are two total income variables in the FES, 'normal total personal gross income' and 'current total personal gross income'; I use the former to smooth transitory income variation.

benefits regressions caused by households in which there are workers as well as non-workers. This could lead to biases in my predictions, and would complicate the Murphy-Topel standard error correction for generated regressor bias (cf. Appendix). To overcome this problem, I generate a pseudo-random variable for each 'mixed' household to decide whether to include its worker(s) or its non-worker(s) in my auxiliary regressions. This should eliminate the correlation of the two variance matrices.

Since this could introduce some non-randomness because of the different treatment of larger households, I also experimented with deciding for *each* household (*i.e.* also for households where there are only workers or only non-workers) whether to use it in the earnings or the benefits regression. This would of course decrease sample size a lot (because households which were assigned to be included in the earnings regression but which only have non-workers and households which were assigned to be included in the benefits regression but which only have workers would have to be dropped). The results between all three methods (include everybody; include only either workers or either non-workers in 'mixed' households; assign households randomly to either regression and then drop accordingly) were however so similar that I decided to employ the random dropping only for 'mixed' households (the correlations of earnings and benefits predicted in the three different ways were always higher than 0.99).

Since the predictors, calendar time, age, and occupation, are available in exactly the same way in the BHPS as they are in the FES, I use the coefficients of the auxiliary regressions to predict earning power and benefit entitlement for each BHPS respondent. In contrast to Blau and Meghir and Whitehouse, I use the natural logarithm of the income variables in the transition regressions (rather than levels). This has two major advantages. First, it simplifies making statements about the elasticities of making a certain transition with respect to the two income variables and second, it enables us to test the replacement ratio hypothesis, since the replacement ratio is defined as $R = \frac{benefits}{earnings} \therefore \ln(R) = \ln(benefits) - \ln(earnings)$. By testing whether $\beta_{benefits} = -\beta_{earnings}$ we can therefore test whether the replacement ratio hypothesis holds.

Missings in the occupational dummy variable in the BHPS (which is one of my predictors) are filled as described for the marital status dummy variables. For people for whom I never observe an occupation in the BHPS, *i.e.*, for people who never worked during the panel period, the earning power is set equal to the base category (non-skilled manual) which seems to be reasonable.

Predicting the two income variables gives rise to an identification issue similar to that of any instrumental variables approach, *i.e.* there have to be some predictors in the two income regressions which do not have a direct effect on the transition equations; only indirectly via the two income variables. That is why I do not use calendar time and

Variable	Men	Women
Failure	0.0073	0.0075
(Never married)	0.0885	0.0611
Married	0.7812	0.7064
Separated	0.0250	0.0234
Divorced	0.0854	0.1215
Widowed	0.0199	0.0876
Age in months minus 479	137.2959	141.3131
(Age in years)	51.3580	51.6928
Duration	47.4213	46.9202
Employment experience	0.9068	0.6227
Ever had child	0.7654	0.8198
Ever had OP	0.6497	0.4633
(Excellent health)	0.2579	0.2249
Good health	0.4502	0.4501
Fair health	0.1990	0.2216
Poor health	0.0752	0.0813
Very poor health	0.0177	0.0221
Unemployment rate	6.4662	6.5982
Earnings in Jan 04 GBP/wk	399.9163	202.8809
Benefits in Jan 04 GBP/wk	138.3641	58.1184
Person-months	310,287	369,193
Persons	3,828	4,533

Table 1: Sample means (person-averages rather than person-month averages for the three time-constant variables)

occupation in the transition equations; I am assuming that calendar time and occupation do not influence the transition hazards directly, only via earnings and benefits. I use age in the prediction as well as in the transition equations.

In Table 1, I present the averages of the dependent and explanatory variables for my final estimation sample (variables which are not used in my regressions are in parentheses). I am reporting person-month averages except for the three time-constant variables employment experience, ever had child, and ever had OP, for which I am reporting person-averages.

How does my specification compare to those of Blau and Meghir and Whitehouse? Overall, it is closer to the parsimonious specification of Meghir and Whitehouse rather than to the huge model which Blau estimates. However, compared to both Meghir and Whitehouse and Blau, I use more detailed information on health (four dummies vs two or one in Meghir and Whitehouse, one in Blau), marital status (four dummies vs none in Meghir and Whitehouse, one in Blau), and unemployment (regional rate vs national). Furthermore, I use three unique time-constant variables to capture people's occupational pension, parental, and work histories.

4 Results

4.1 Empirical Transition Hazard Rates by Various Characteristics

4.1.1 Methodology

This part of the paper reports conditional transition probabilities ('empirical hazard rates') between the two labour market states depending on the age of the person. I calculate them separately for different groups of people as classified by sex and occupational pension status.

Separating these calculations by sex will help to see whether the difference in the state pension age (65 for men, 60 for women) leads to differences in labour market behaviour. Differentiating between OP members and non-members allows for a check of whether the phenomenon of Meghir and Whitehouse that men who are OP members have a higher return probability than male non-members can still be observed in more recent data, and whether there is any difference for women in that respect.

These descriptive devices shall be used to point towards interesting features of the data which warrant further investigation in the multivariate parametric analysis which follows.

My definition of the hazard rate is analogous to the Kaplan-Meier hazard, with the important difference that age rather than duration is the analysis time. The annual EN hazard 'at' age 65 (denoted by $h^{EN}_{65,a}$, with the a subscript for 'annual') is the number of people who made a transition from employment to non-employment between the month directly before their 65^{th} birthday and twelve months thereafter ($EN_{780,792}$, as $65 \times 12 = 780$ and $780 + 12 = 792$), divided by the risk set. The risk set consists of the number of people in state E directly before their 65^{th} birthday (E_{780}) minus half of those censored between the two relevant points in time ($EC_{780,792}$). This is known as actuarial adjustment.[11] Algebraically,

$$h^{EN}_{65,a} = \frac{EN_{780,792}}{E_{780} - (EC_{780,792})/2}.^{12} \tag{8}$$

[11] Actuarial adjustment therefore deals with the fact that some people who are observed at the starting age are not observed anymore twelve months later. Subtracting 50% of the censored observations from the risk set assumes that the density of censoring over time is uniform.

[12] By analogy, the formula for the *monthly* age-dependent hazard 'at' age 65 ($h^{EN}_{65,m}$) uses $_{780,781}$ subscripts.

I also compute 95% confidence intervals for my age-dependent 'hazard' h. These are given by

$$ci_{95\%}(h) = \max[h \pm \Phi^{-1}(0.975)h\sqrt{\frac{1-(h/2)^2}{f}}, 0], \tag{9}$$

where, in analogy to Greenwood's (1926) formula, $\Phi^{-1}(\cdot)$ is the inverse cumulative standard normal probability distribution function and f the numbers of failures ($EN_{780,792}$ in the previous example).

I talk about differences in the hazard rates being 'statistically significant' if these 95% confidence intervals do not overlap.

4.1.2 A First Overview: Annual Transition Hazards for Men and Women

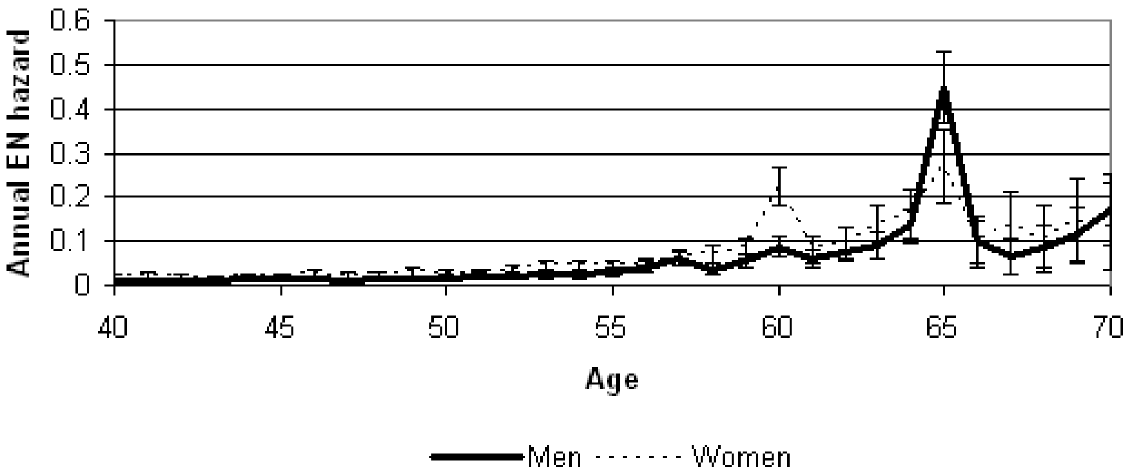

Figure 2: Annual EN hazard by sex

The exit hazard, in Figure 2, is very similar for men and for women, with three exceptions: the peak at age 65 is higher for men than for women, the peak at age 60 for women than for men (with these two differences being statistically significant), and in general women's exit hazard is slightly higher at older ages than men's (although this difference is not statistically significant). Men's age-60 peak is only very minor.

In numbers, the annual hazard of stopping to work 'at' age 65 is 44.97% for men (3.81 times larger than the average of the two surrounding years) and 26.99% (2.08 times larger) for women. At age 60, the corresponding figures are 8.79% (1.52 times larger) for men and 22.30% (2.68 times larger) for women.

The return hazard, in Figure 3, is statistically significantly higher for men than for women up to age 55 and almost completely the same for the two sexes thereafter.

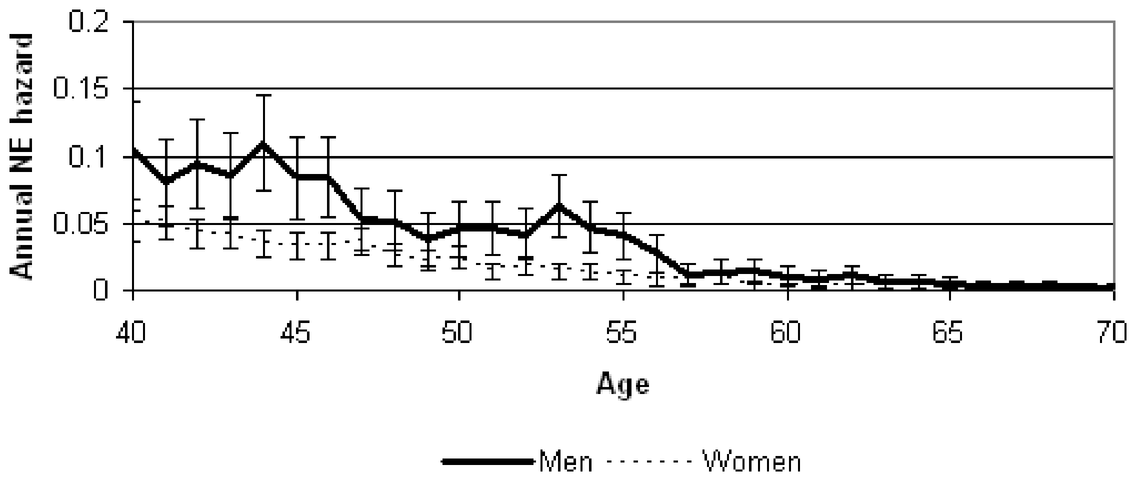

Figure 3: Annual NE hazard by sex

Both for men as well as for women, there are no differences in the transition hazards depending on whether or not you ever had children. This is interesting because one might have suspected that women who had children could have different employment patterns later during their life.

4.1.3 Assessing Birthday Effects: Monthly Transition Hazards for Men and Women

An interesting question is whether there is heaping of transitions at certain ages (so-called 'birthday effects'), or whether transitions are evenly spread over the year. This is because birthday effects at ages which play important roles in a country's pension system could indicate incentive effects of the pension system. Blau addressed this question using quarterly hazard rates to analyse US data while Meghir and Whitehouse for the UK only addressed it implicitly when referring to steep parts in the survivor function at age 65 (keep in mind that they analysed men only). We can spot such steep parts in Figure 1 at age 65 for men and at age 60 for women.

Examining the monthly hazard rates in my own data reveals that birthday effects are important in the UK. A man who is employed directly before his 65^{th} birthday has a 45.38% risk of stopping to work within one month (12.93 times larger than the average of the two surrounding months), a woman 19.49% (4.05 times larger). For the 60^{th} birthday these figures are 3.25% (6.08 times larger) for men and 11.78% (4.53 times larger) for women. This means that for British men, birthday effects are approximately twice as large as Blau found for the US.

4.1.4 Annual Transition Hazards by Occupational Pension Membership

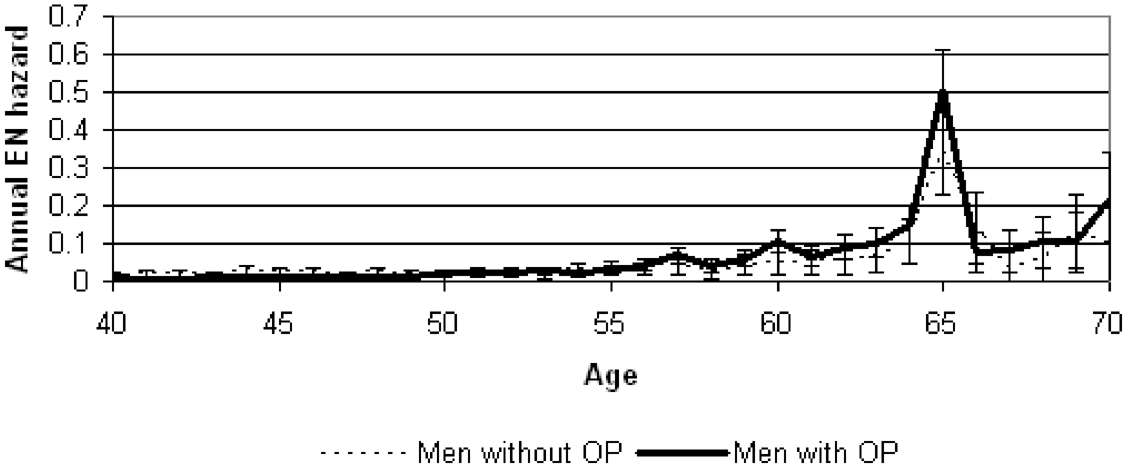

Figure 4: Annual EN hazard by OP status for men

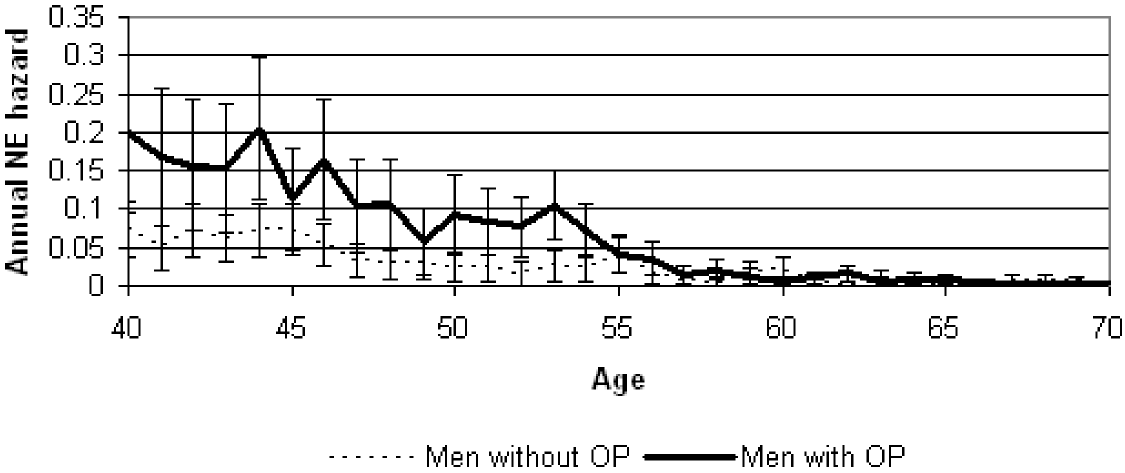

Figure 5: Annual NE hazard by OP status for men

Looking at men in Figures 4 and 5, having ever been member of an OP scheme is not associated with exit rates from employment, but it does increase your return hazard up to age 55. This is consistent with Meghir and Whitehouse (1997: 332). OP members also have a higher age-65 exit peak, but the difference is not statistically significant.

For women, in Figures 6 and 7, we observe a similar pattern. The return hazard for OP members is higher than for non-members up to age 55, and the difference is statistically significant (stronger than for men). Furthermore, female OP members have

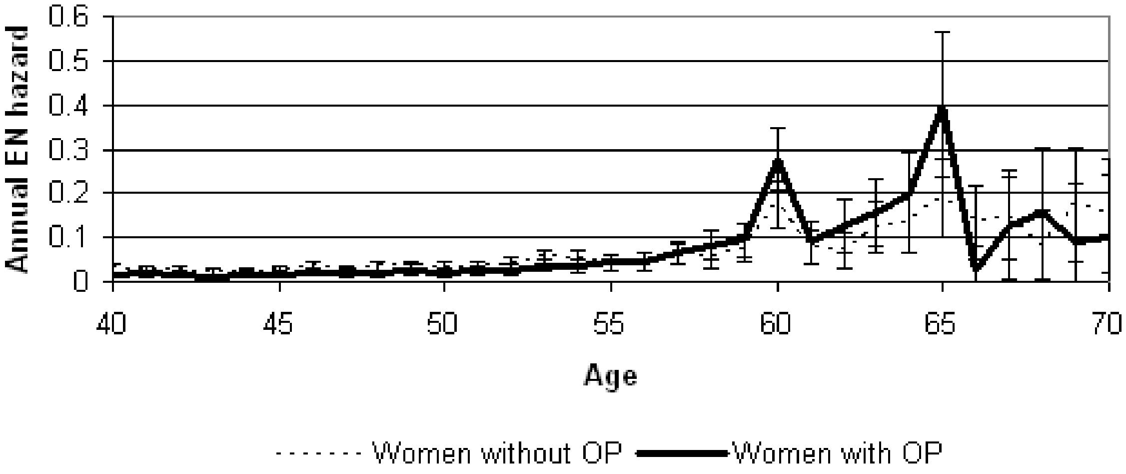

Figure 6: Annual EN hazard by OP status for women

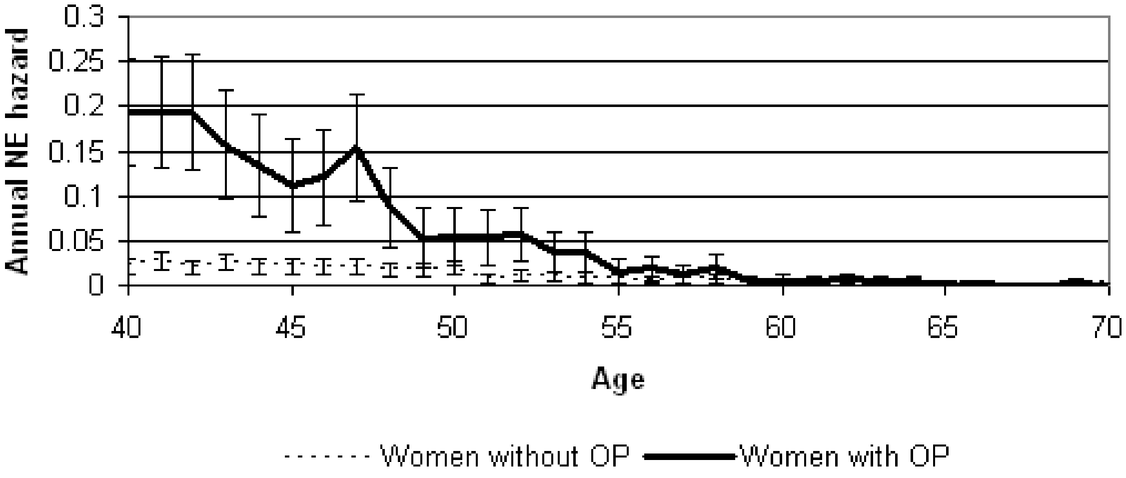

Figure 7: Annual NE hazard by OP status for women

Exit hazard coefficients	Men		Women	
	Without UH	**With UH**	**Without UH**	**With UH**
Married	−0.2501*	−0.1935	−0.3234*	−0.4133***
Separated	−0.2664	−0.2669	−0.0390	−0.2386
Divorced	−0.0719	0.0249	−0.2480	−0.3066*
Widowed	−0.2833	−0.2095	−0.2184	−0.2093
Age	0.0077***	0.0054**	0.0045***	0.0016
Age squared	−0.0136***	−0.0076	−0.0065	0.0033
Age = 60	0.9585***	1.0027***	2.1644***	2.1809***
Age = 65	3.0484***	3.0714***	2.1399***	2.1600***
Duration	−0.0545***	−0.0503***	−0.0623***	−0.0569***
Duration squared	0.0967***	0.0844***	0.1200***	0.1077***
Age x duration	0.0522***	0.0538***	0.0616***	0.0646***
Employment exp.	−1.3094***	−1.1686***	−0.8023***	−0.7567***
Ever had child	−0.0986	−0.1173	−0.3245***	−0.3061***
Ever had OP	0.2104***	0.2265***	−0.1680***	−0.1690**
Good health	0.1434*	0.1464**	0.0265	0.0453
Fair health	0.2560***	0.2510***	0.4133***	0.4330***
Poor health	1.0171***	1.0325***	0.9442***	1.0106***
Very poor health	1.2938***	1.4155***	1.1238***	1.1508***
Unemployment rate	−0.0658***	−0.0346	−0.0961***	−0.0672***
Earnings	−0.5899***	−0.5905***	−0.5633***	−0.5200***
Benefits	−0.0138	0.1495	−0.0128	−0.2615
# of person-months	232,741	see below	212,469	see below
Log likelihood	−7992.5608	see below	−8727.0997	see below
UH: π (cf. text)		0.0888***		0.8141***
UH: EN mass point		−1.4775***		1.4303***

Table 2a: Exit hazard regression coefficients without and with unobserved heterogeneity (UH). ***: Statistically significant at the 1% level; **: 5% level; *: 10% level (Murphy-Topel standard errors; cf. Appendix).

more pronounced exit peaks at ages 60 and 65 than non-members, but these differences are again not statistically significant.

4.2 Hazard Regression Analysis

Table 2a presents the coefficients of the exit regressions; Table 2b those of the return regressions. Each table contains four columns, the first for men without unobserved heterogeneity (a standard cloglog model with the standard errors adjusted for clustering by person, allowing for some adjustment to deal with the fact that we observe most people for more than one month), the second for men with unobserved heterogeneity, the third for women without unobserved heterogeneity, and the fourth for women with unobserved heterogeneity. All the standard errors have been corrected for generated regressor bias

Return hazard coefficients	Men		Women	
	Without UH	With UH	Without UH	With UH
Married	0.1985	0.2394	−0.5307**	−0.4362
Separated	−0.0489	−0.1727	−0.3227	−0.2821
Divorced	0.0161	0.0877	−0.1535	−0.0246
Widowed	−0.1910	−0.0003	−0.4575	−0.3294
Age	0.0071	0.0069	0.0056	0.0035
Age squared	−0.0422**	−0.0392*	−0.0412**	−0.0320**
Duration	−0.0540**	−0.0251	−0.0279*	0.0019
Duration squared	0.3298*	0.2537	0.2740***	0.2375**
Age x duration	−0.0056	−0.0314	−0.0280	−0.0614***
Employment exp.	1.3437***	1.2625***	0.8058***	0.7578***
Ever had child	0.0435	0.1164	0.3252***	0.3506
Ever had OP	0.0762	0.1838	0.4291***	0.4373***
Good health	0.0503	0.1047	0.0139	0.0105
Fair health	−0.3208**	−0.2759*	−0.2872*	−0.2903**
Poor health	−1.3127***	−1.2826***	−1.0996***	−1.1030***
Very poor health	−1.4884***	−1.3899***	−1.8703***	−1.8907***
Unemployment rate	−0.0966	−0.1076	−0.1995***	−0.1925***
Earnings	0.6268***	0.6638***	0.6219***	0.6573*
Benefits	−1.4606***	−1.6306***	−1.4256***	−1.3130***
# of person-months	77,546	310,287	156,724	369,193
Log likelihood	−3738.9565	−11630.479	−5634.803	−14226.431
UH: π (cf. text)		0.0888***		0.8141***
UH: NE mass point		1.3810***		−1.4009***

Table 2b: Return hazard regression coefficients without and with unobserved heterogeneity (UH).
***: Statistically significant at the 1% level; **: 5% level; *: 10% level (Murphy-Topel standard errors; cf. Appendix).

using the estimator by Murphy and Topel (1985) which is also discussed in Greene (2003) and Hardin (2002). This estimator is described in the Appendix.

My discussion of the results in this section will focus on the model with unobserved heterogeneity. There are not many differences between the two models but some exist, and since the unobserved heterogeneity parameters all turn out to be highly statistically significant these differences could be crucial.

For reasons which I have already mentioned I will refer to hazard ratios (cf. Section 2.2) rather than coefficients. I will also use Figures 10 to 21 which show predicted hazard rates, and to Table 3 which shows the median transition ages computed from the predicted survivor function in several alternative scenarios. The advantage of these figures and this table is that they make it possible to compare the importance of dummy and continuous variables; for continuous variables, the hazard ratio alone is not very helpful since it does not tell us what the mean and standard deviation of the corresponding variable are. The combination of hazard ratios, predicted hazard rates, and median exit and return ages will enable us to assess in a rigorous way which of the explanatory variables are the most important ones.

Let us now examine the regression results in detail.

Most of the marital status coefficients are statistically insignificant. The only two variables which are statistically significant are married and divorced in women's exit regression: married women have a 34% lower exit risk than never married, separated, and widowed women; divorced women a 26% lower exit risk.

The effects of age and duration are not so easy to assess since both variables appear in levels as well as squared, plus there is also the interaction term. One way of analysing the impact of age and duration is to keep one variable constant at a certain level and to vary the other. Doing this reveals an important difference between men and women: the effects of age and duration on exit are almost the same for the two sexes while their effects on return are very different.

Varying age while keeping duration constant at its sex-specific average shows that *ceteris paribus*, increasing age continuously increases the exit risk for both men and women. At age 40, the combined hazard ratio of the age and duration variables is 1.0 for both men and women, at age 50 it is 2.3 for men and 1.8 for women, at age 60 (without taking into account the birthday effect dummy variable which I will cover later on) 4.4 for men and 3.7 for women, at age 65 (again without taking into account the birthday effect dummy variable) 5.5 for men and 5.4 for women. Thereafter, the hazard ratio of women is higher than that of men.

Doing the same for the return to the labour market as shown in Figure 8 shows that the *ceteris paribus* impact of age on the return hazard (again holding duration constant

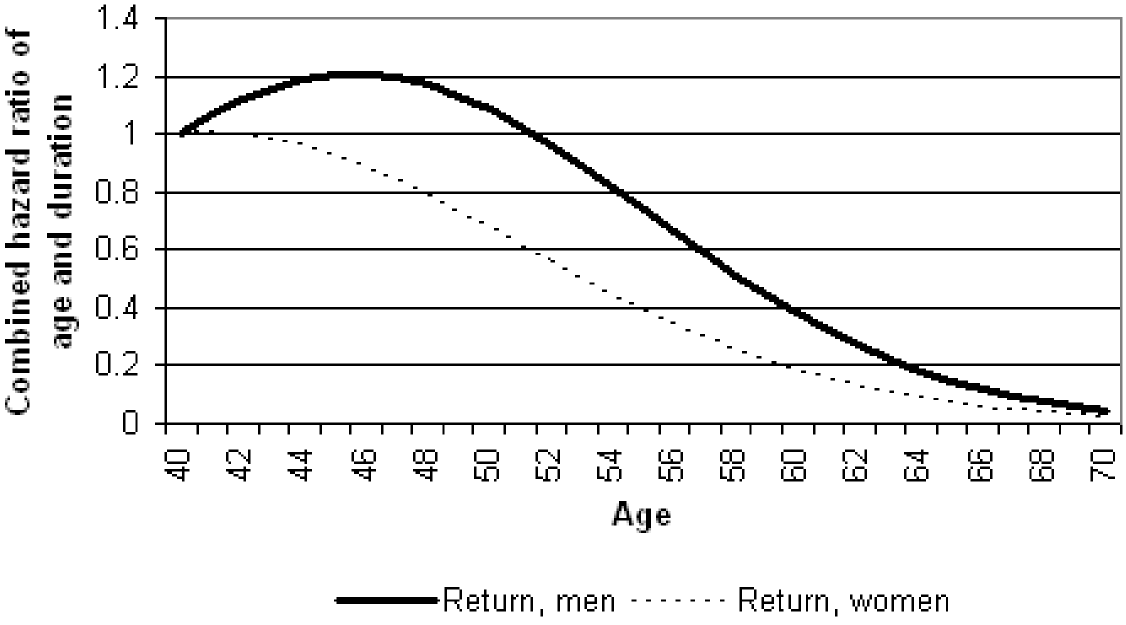

Figure 8: Combined *ceteris paribus* effect of age and duration on the predicted monthly NE hazard by age, holding duration constant at its sex-specific average

at its sex-specific average) is inversely U-shaped: the age effect on the return probability reaches its maximum at age 46 for men and at age 41 for women. At age 70 the sizes of the effect for men and for women meet, but up to that age, the age variable has a more positive effect on men than on women. The difference is largest at age 51 where men's combined hazard ratio of the age and duration variables is 1.0 and women's 0.6.

Varying duration while keeping age constant at its sex-specific average shows that *ceteris paribus*, the risk of exiting work falls as duration increases. The effect is virtually identical for men and women. When duration is 1 (*i.e.* you are in the first month of your employment spell), men's combined hazard ratio of the age and duration variables is 1.0, women's 0.9. After one more year in employment (*i.e.* a value of the duration variable of 13), this hazard ratio has already fallen to 0.5 for both sexes, after another additional year to 0.3 (men) and 0.2 (women), respectively.

Lastly, we can do the same for the effect of duration on return as shown in Figure 9. As for age, I again find that the difference between men and women is in the return behaviour, not in the exit behaviour: for men, the *ceteris paribus* effect of duration on return is falling, on women somewhat increasing. This means that the longer a man is out of employment, the lower his *ceteris paribus* probability to come back to work. For women, the opposite is true. When duration is 1, the combined hazard ratio of the age and duration variables is 1 for both men and women, one year later it is 0.7 for men and

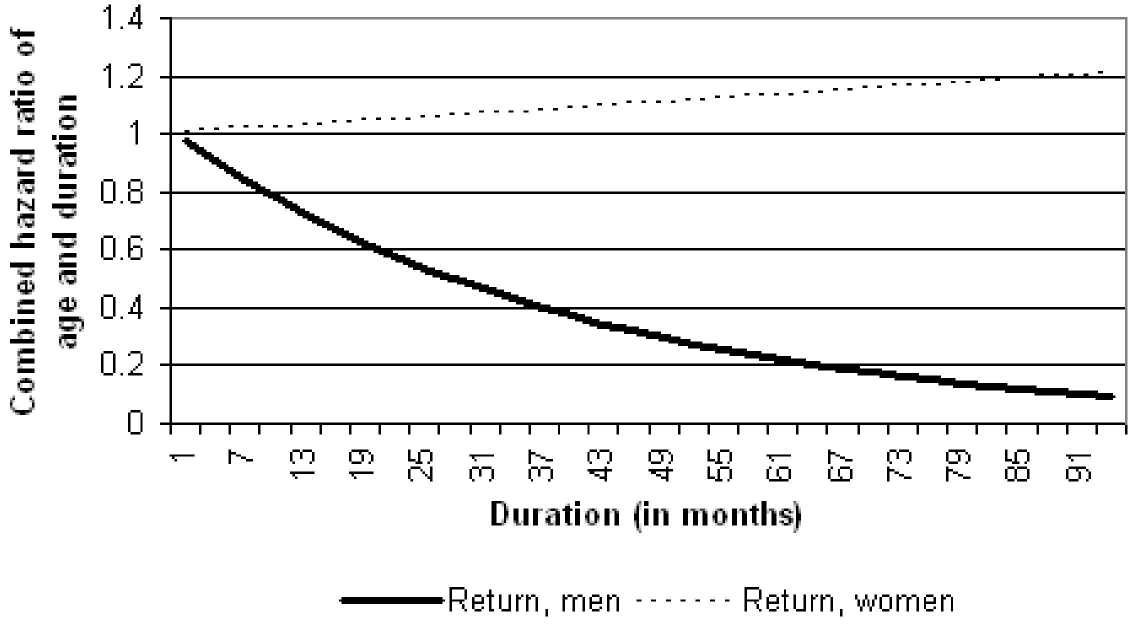

Figure 9: Combined *ceteris paribus* effect of age and duration on the predicted monthly NE hazard by duration, holding age constant at its sex-specific average

1.0 for women, another year later 0.5 for men and 1.1 for women.

So far, I have shown that the effects of age and duration on exit are very similar for the two sexes, but their effects on return are quite different. Now we also need to take a look at the two birthday effect dummy variables which appear in the exit regression, one for the month of the 60^{th} birthday or the month thereafter, and one for the month of the 65^{th} birthday or the month thereafter. The coefficient sizes are huge and markedly bigger than the size of any other dummy variable. At a man's 60^{th} birthday, the exit risk is 2.7 times higher than can be explained by all the other variables. At a man's 65^{th} birthday, the hazard ratio is 21.6 (sic). For women, as could already be suspected from Section 4.1.3 which discussed evidence for birthday effects in the empirical hazard rates, the effects of the two birthdays are much more similar than for men; the corresponding hazard ratios are 8.9 and 8.7, respectively. This shows that even after controlling for a number of important explanatory variables, birthday effects are an important phenomenon in the labour market behaviour of older Britons.

As expected, employment experience reduces the exit risk and increases the return probability for both men and women. The effects are statistically significant: comparing a man who has always worked when observing his work history between ages 15 and 40 to a man who never has, the former has a 69% lower exit risk than the latter, and a 3.5 times as large return probability. *Prima facie*, this looks huge, but one needs to take into account that the variation in work experience is not that large: as we know from Table 1, the mean value for men is 0.9068 (90.68%); the standard deviation is 0.2306. Table 3 shows that reducing men's employment experience by one third lowers the median exit age of those who were working at age 40 from 53.6 by 3.3 to 50.3, and increases the median return age of those who were not working at age 40 from 42.6 by 1.2 to 43.8.

Comparing a woman who always worked to one who never worked, the former has a 53% lower exit risk than the latter, and a 2.1 times larger return probability. The raw effects are therefore smaller for women than for men. Looking at Table 3 again shows that also the effects of a relative change in the employment experience variable (which takes into account that women's employment experience has a different, smaller, mean and a different, larger, standard deviation than men's) are smaller for women than for men.

Having ever had a child has a statistically significant effect on women's exit risk: women with a child or children have a 26% lower exit risk than women without. This effect is statistically significant at the 1% level. Table 3 shows that the median exit age of women who were working at age 40 is 51.3 if they ever had a child and 2.5 years lower (48.8) if they never had. Having ever had children does not have statistically significant effects on women's return probability or on men's transition probabilities.

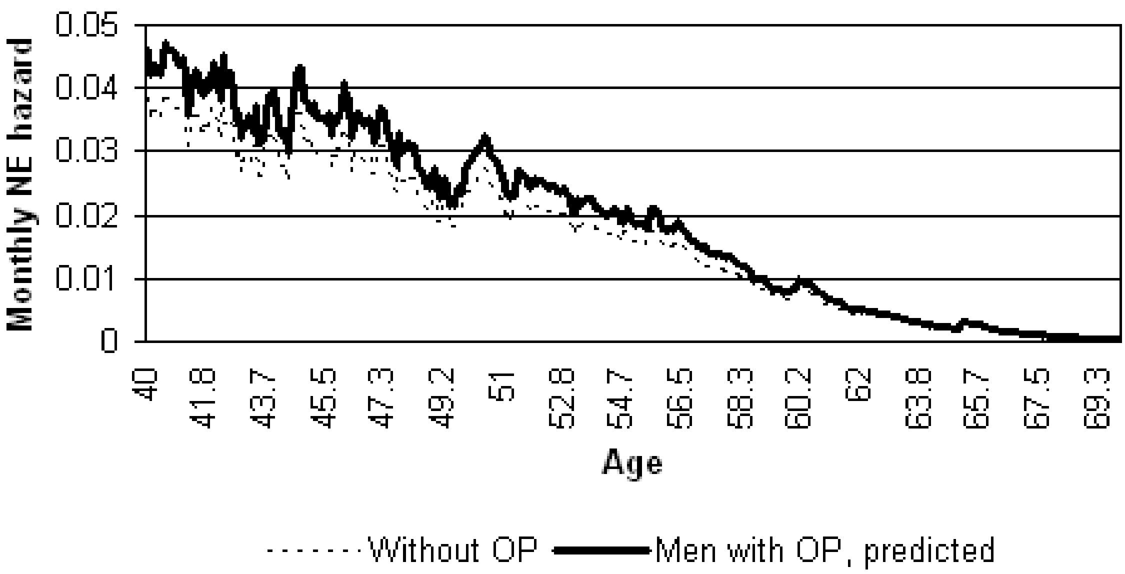

Figure 10: *Ceteris paribus* effect of not being an OP members compared to being one on the predicted monthly NE hazard for men by age

Having ever been member of an OP scheme shows the expected effects for women: it lowers their exit risk by 16% and it increases their return probability by 55% (both sta-

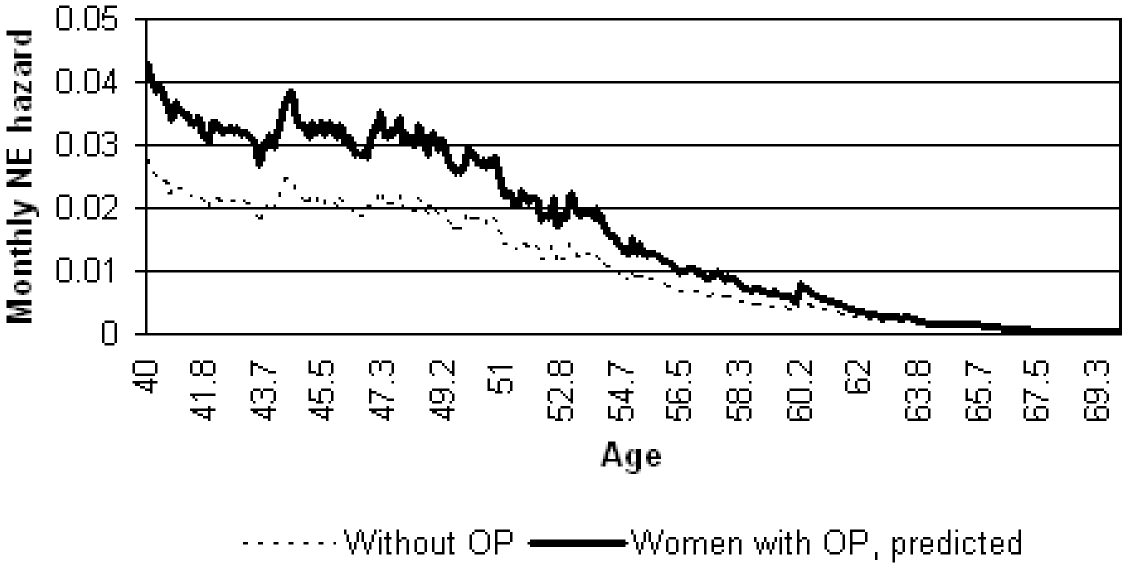

Figure 11: *Ceteris paribus* effect of not being an OP members compared to being one on the predicted monthly NE hazard for women by age

tistically significant). Figure 11 shows a hypothetical counterfactual related to women's return probability: the thick line represents the predicted return probability of women who are OP members, with the values of the other explanatory variables equal to those observed among women who have the corresponding age, are out of work, and OP members. The thin dashed line plots the hypothetical counterfactual of making these women non-OP members, keeping the values of all other explanatory variables the same. In terms of the median transition ages shown in Table 3, this means that making women who are OP members non-members decreases their median exit age by 1.5 years and increases their median return age by 0.9 years.

However, the effects of occupational pension membership are counterintuitive for men. The effect on the return probability is not statistically significant (cf. Figure 10), and moreover, men who are OP members have an exit risk 25% higher than non-members (statistically significant at the 1% level). This translates into the absurd prediction that men who are not OP members exit the labour market 2.1 years later than members (Table 3). This finding is different to what Meghir and Whitehouse found when plotting employment survivor functions of men with and without OP schemes and may warrant further examination.

Health status is found to be one of the most important determinants of older people's labour market transitions (Figures 12 to 15). In three out of the four equations, being in good or excellent health does not make any statistically significant difference. However,

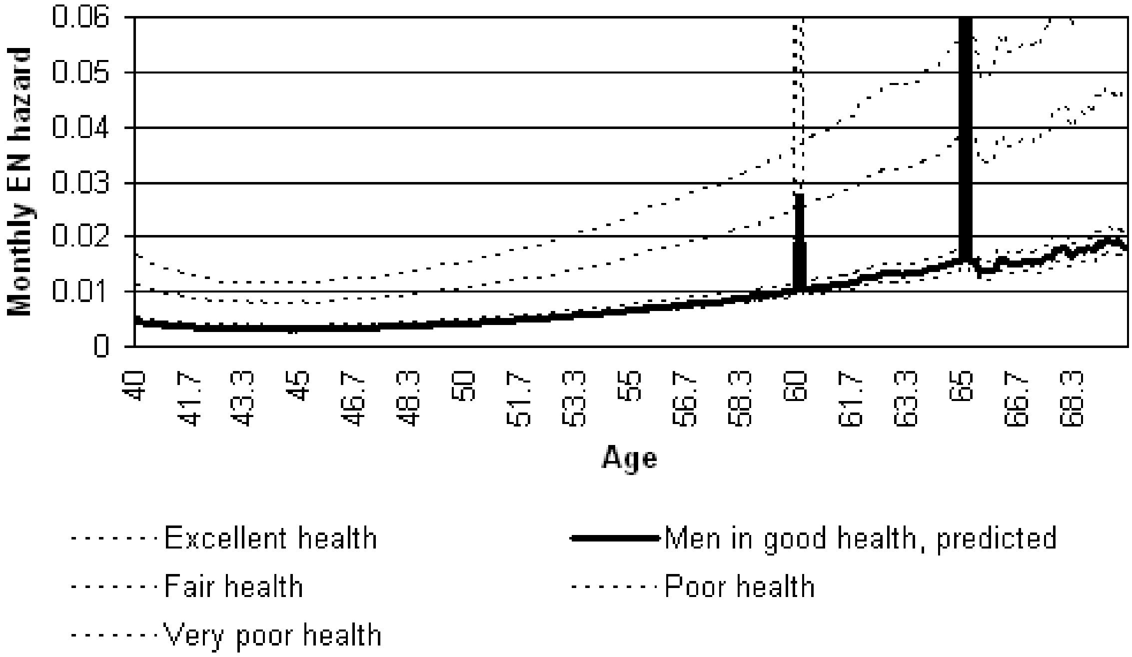

Figure 12: *Ceteris paribus* effects of health compared to being in good health on the predicted monthly EN hazard for men by age

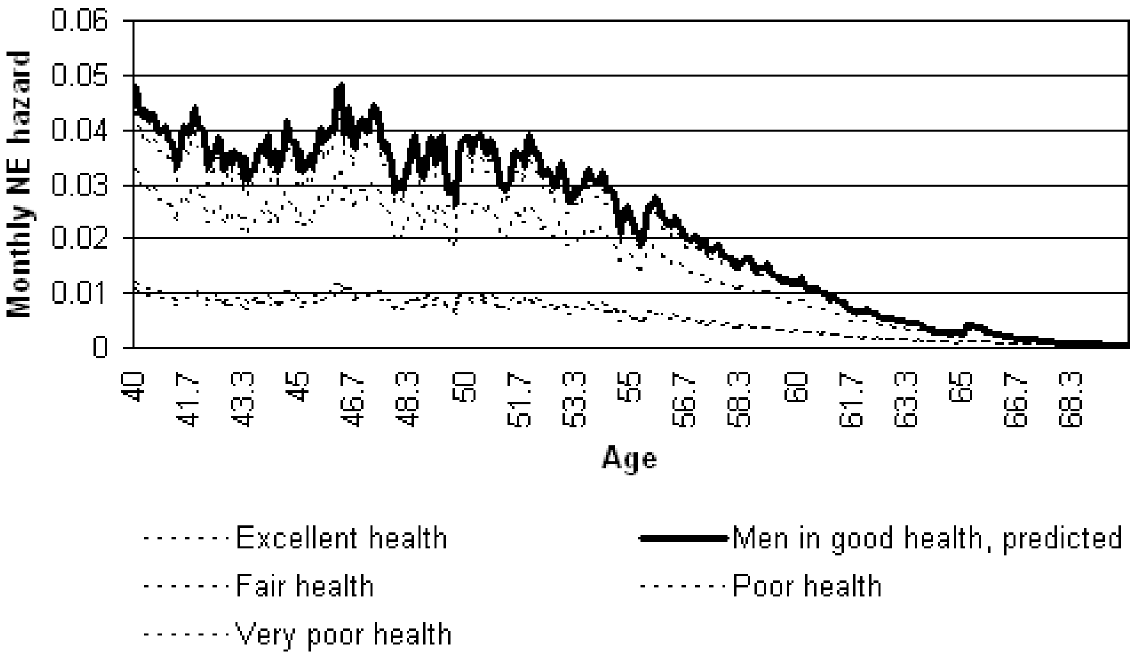

Figure 13: *Ceteris paribus* effects of health compared to being in good health on the predicted monthly NE hazard for men by age

Figure 14: *Ceteris paribus* effects of health compared to being in good health on the predicted monthly EN hazard for women by age

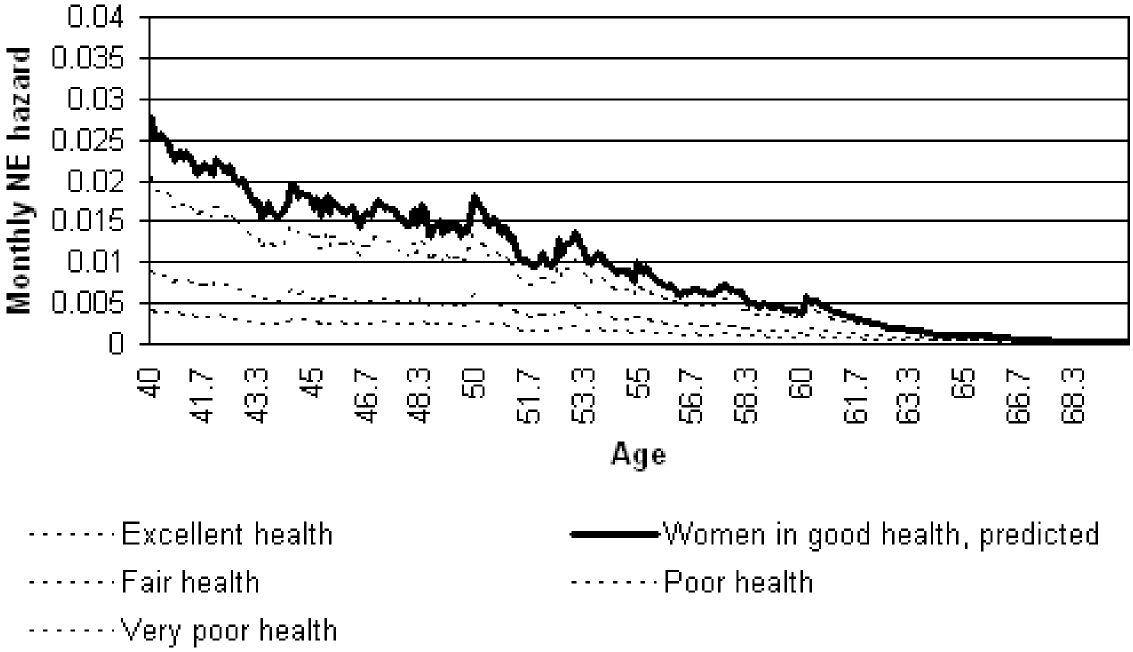

Figure 15: *Ceteris paribus* effects of health compared to being in good health on the predicted monthly NE hazard for women by age

as we can see from the figures, being in poor or very poor health does make a huge difference. Men in poor health have a 2.8 times larger exit risk than men in excellent health; for women, the corresponding hazard ratio is almost the same (2.7). Being in very poor health further increases this risk: for men, from 2.8 times larger to 4.1 times larger, for women, from 2.7 times larger to 3.2 times larger.

For the return, we see similarly huge effects. Men in poor health have a 72% lower return probability than men in good or excellent health; for women it is 67% lower. Being in very poor health further reduces the chances to return to work: men in very poor health have a 75% lower return probability than men in good or excellent health; for women it is even 85% lower.

Of course this also translates into correspondingly huge effects on the median transition ages shown in Table 3. One particularly striking result: for women who were not working at age 40 and continuously in very poor health, I predict only 39.9% to return to work by age 70. Therefore, I cannot even calculate a median return age for this group.

The regional unemployment rate has statistically significant effects on women's labour market transitions. Unfortunately, these are not so easy to interpret: when regional unemployment is high, women make less transitions in either direction. For men, there are no statistically significant effects.

As expected, higher earnings lower one's exit risk and increase one's return probability. This is also in accordance with theoretical results shown in microeconomic optimal retirement decision models such as the one already mentioned earlier on by Kingston (2000). The effects are statistically significant but not that large. As can be seen from Table 3, increasing predicted earnings by one third increases men's median exit age by 1.6 years (from 53.6 to 55.2) and women's by 1.2 years (from 51.2 to 52.4). It also reduces the median return age by 0.5 years for men (from 42.6 to 42.1) and by 0.6 years for women (from 43.3 to 42.7).

The effects of benefits are quite different: they do not have any statistically significant effects on exit, but they do have important effects on return, as shown in Figures 16 and 17 which plot the effects of reducing benefits by one third. The median return age is reduced by 1.3 years for men (from 42.6 to 41.3) and by 1.5 years for women (from 42.6 to 41.1). This means that such a reduction by one third halves the median duration out of work for men; the effect is similarly large for women. The impacts of benefits which I find are markedly bigger than in Blau for the US and Meghir and Whitehouse for the UK.

Putting together what I have said about the effects of earnings and benefits and looking at the confidence intervals of the coefficient estimates, we can already suspect that the replacement ratio hypothesis does not hold. Indeed formal tests of $\beta_{benefits} =$

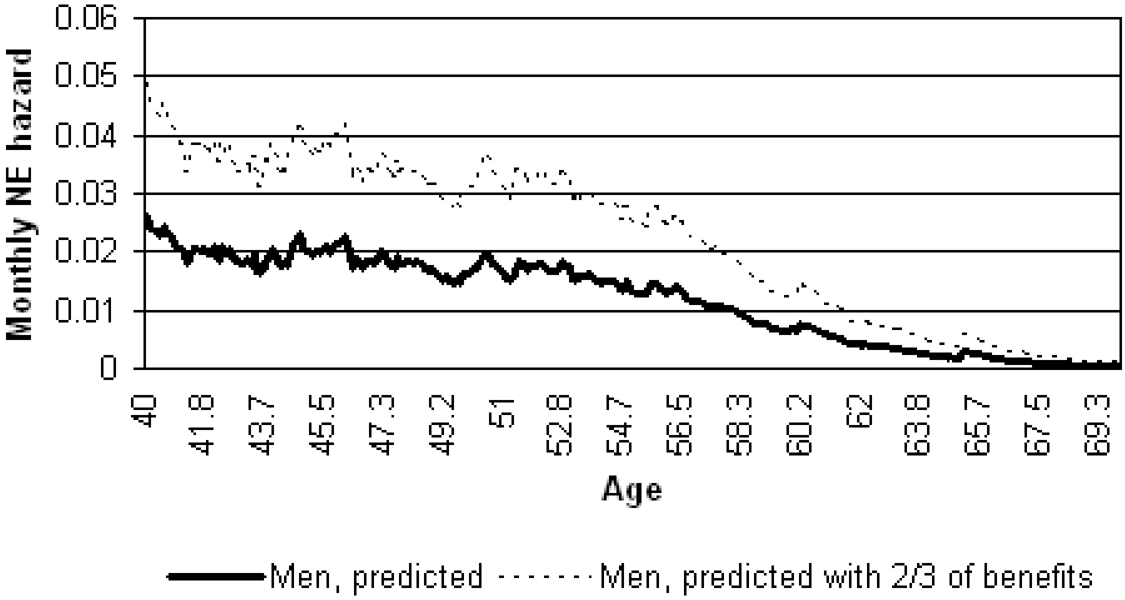

Figure 16: *Ceteris paribus* effect of reducing predicted benefits by one third on the predicted monthly NE hazard for men by age

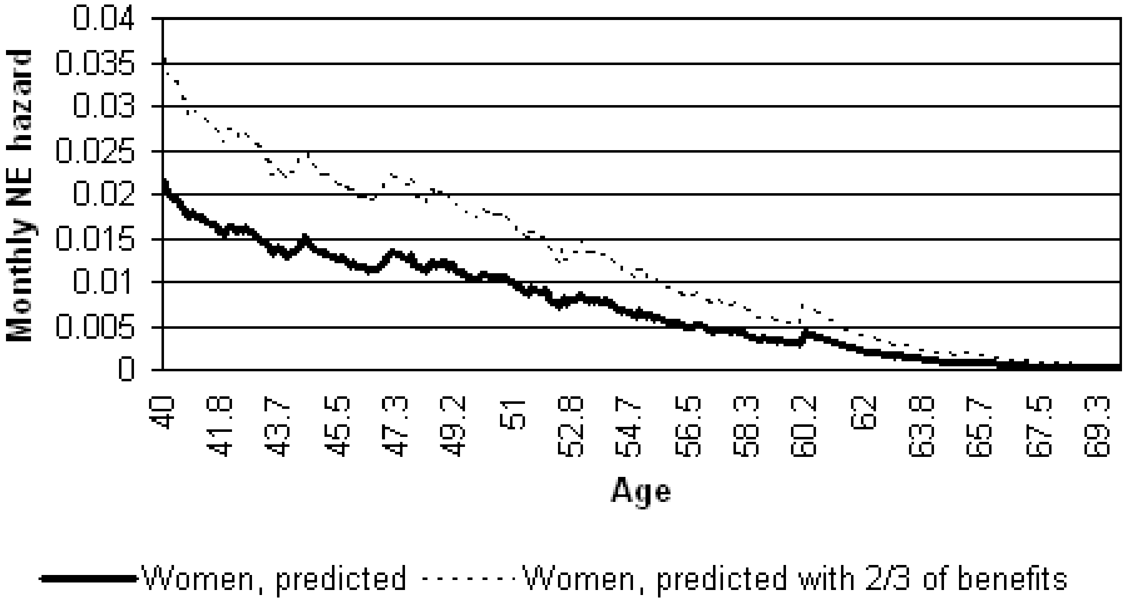

Figure 17: *Ceteris paribus* effect of reducing predicted benefits by one third on the predicted monthly NE hazard for women by age

$-\beta_{earnings}$ were always rejected, for men as well as for women and without UH as well as with UH.

Summing up, we can keep in mind that apart from age, where I find huge peaks at ages 65 (for both sexes) and 60 (for women), health and benefits seem to be the most important determinants of labour market transitions quantitatively. The impact of non-labour income on the return probability could well warrant further examination.

Similar to Meghir and Whitehouse, introducing unobserved heterogeneity does not lead to systematic changes in the coefficients, but there are a couple of changes with respect to statistical significance. I do not formally test for the presence of unobserved heterogeneity (Meghir and Whitehouse develop a score test of unobserved heterogeneity for multiple spell models), but the mass point probabilities and positions are always statistically significant at the 1% level.

The estimated positions of the mass points are virtually the same for the two sexes (they have opposite signs but the corresponding mass point probabilities are also almost the inverse of each other). In terms of hazard ratios, movers among men have a 4.4 times higher exit risk than male stayers, and a 4 times higher return probability.[13] Among women, the corresponding hazard ratios are 4.2 and 4.1.[14] 8.88% of men and 18.59% of women are movers, the remainder (91.12% of men and 81.41% of women) stayers. The fact that there are twice as many movers among women as among men, which is statistically significant at any level, may be interesting as a separate result since it shows that women are subject to more labour market transitions; women's employment patterns are less stable than men's.

To conclude this section, let us now take a look at Figures 18 to 21. These four figures compare actual hazard rates to those predicted by my regression model and can be used to assess the goodness of fit or within-sample prediction quality of the model. There is of course a lot of noise caused by small cell sizes (especially at older ages for the exit hazard and at younger ages for the return hazard) which is deliberately not modelled since it does not appear to be systematic, but the general shape of the hazard as well as the two genuine peaks at ages 60 and 65 are well predicted (although I have two smaller peaks at the birthday month and the subsequent month rather than one big peak at the birthday month which is due to the fact that, as already mentioned earlier on, I was not able to include single-month dummy variables in the full model). When plotting the corresponding survivor functions (not shown in the paper, although this result can be seen implicitly in Table 3), it can be seen that the actual and predicted survivor functions

[13] $1/\exp(-1.4775)$ gives the exit hazard ratio of male movers compared to male stayers, $\exp(1.3810)$ the return hazard ratio.

[14] $\exp(1.4303)$ gives the exit hazard ratio of female movers compared to female stayers, $1/\exp(-1.4009)$ the return hazard ratio.

Predicted median transition ages

Exit	Men	Women
Actual	53.4	50.5
REF: Predicted	53.6	51.2
Predicted with 2/3 of experience	−3.3***	−1.6***
Predicted with 4/3 of earnings	+1.6***	+1.2***
Predicted with 2/3 of benefits	+0.6	−0.9
REF: With OP	53.5	52.5
Without OP	+2.1***	−1.5**
REF: Good health	54.2	52.4
Excellent health	+1.4**	+0.4
Fair health	−1.0***	−3.2***
Poor health	−7.4***	−7.4***
Very poor health	−9.8***	−8.1***
REF: With child/children	54.1	51.3
Without child/children	−1.1	−2.5***
Return	**Men**	**Women**
Actual	42.3	43.3
REF: Predicted	42.6	43.3
Predicted with 2/3 of experience	+1.2***	+0.6***
Predicted with 4/3 of earnings	−0.5***	−0.6*
Predicted with 2/3 of benefits	−1.3***	−1.5***
REF: With OP	41.2	41.5
Without OP	+0.3	+0.9***
REF: Good health	41.3	42.4
Excellent health	+0.2	±0.0
Fair health	+0.7*	+1.1**
Poor health	+4.8***	+7.2***
Very poor health	+5.4***	>+27.6***
REF: With child/children	42.9	43.3
Without child/children	+0.4	+1.9

Table 3: 'How do selected explanatory variables change the predicted age by which 50% of those who were employed (upper panel) or non-employed (lower panel) at age 40 make a transition to the other state?' Predictions based on the model with unobserved heterogeneity (UH).

According to my predictions, 60.1% of women not working at age 40 and continuously in very poor health have not returned to work by age 70, therefore the '>+27.6***' entry for this group.

***: Statistically significant at the 1% level compared to the reference category ('REF') in the corresponding subpanel; **: 5% level; *: 10% level (Murphy-Topel standard errors; cf. Appendix).

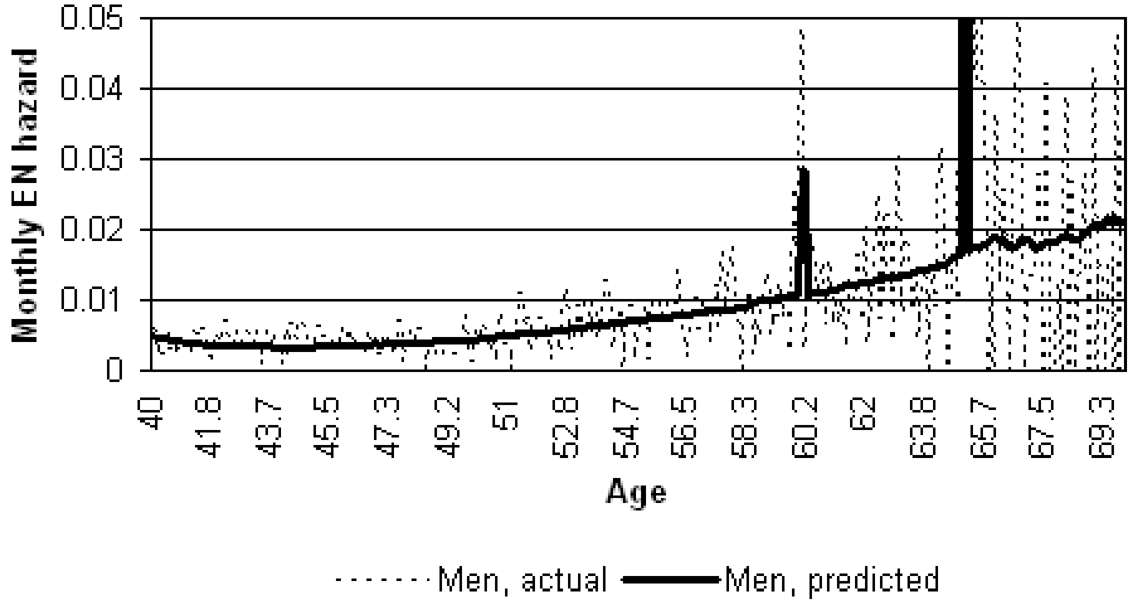

Figure 18: Goodness of fit: actual vs predicted monthly EN hazard for men by age

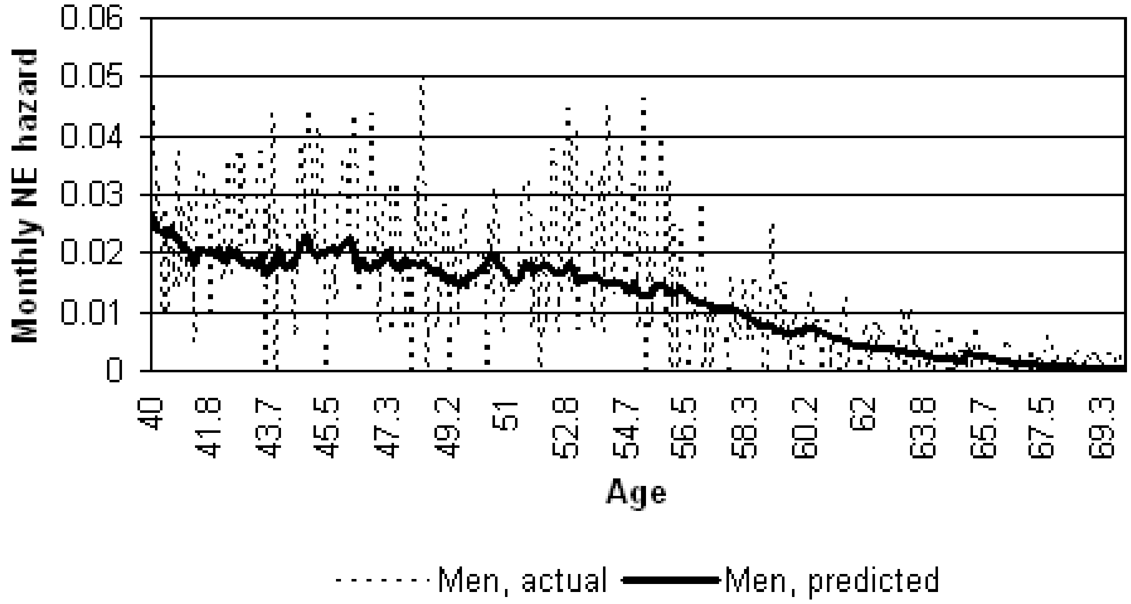

Figure 19: Goodness of fit: actual vs predicted monthly NE hazard for men by age

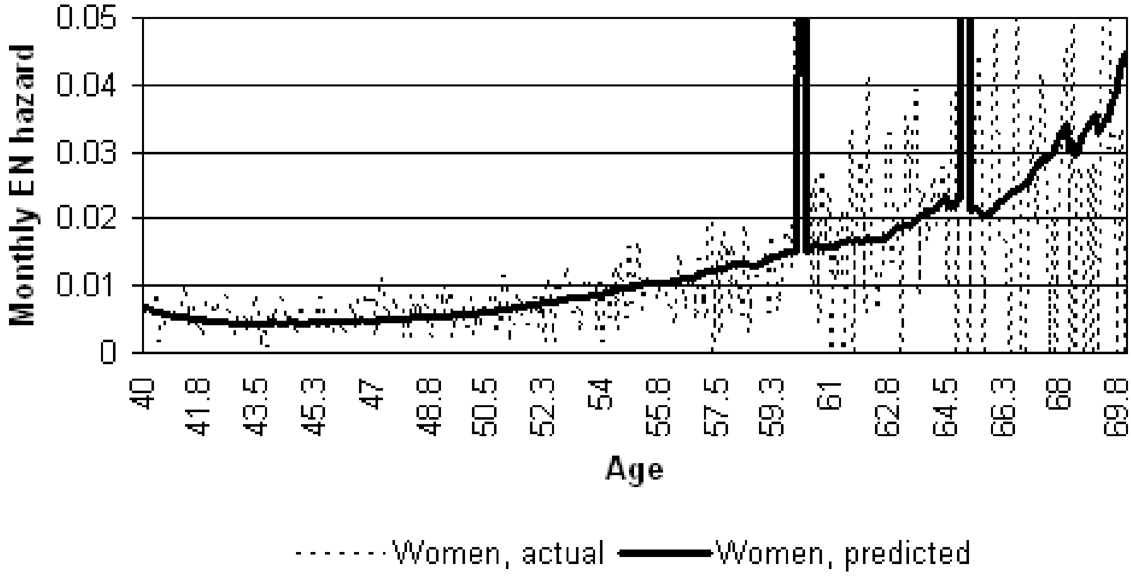

Figure 20: Goodness of fit: actual vs predicted monthly EN hazard for women by age

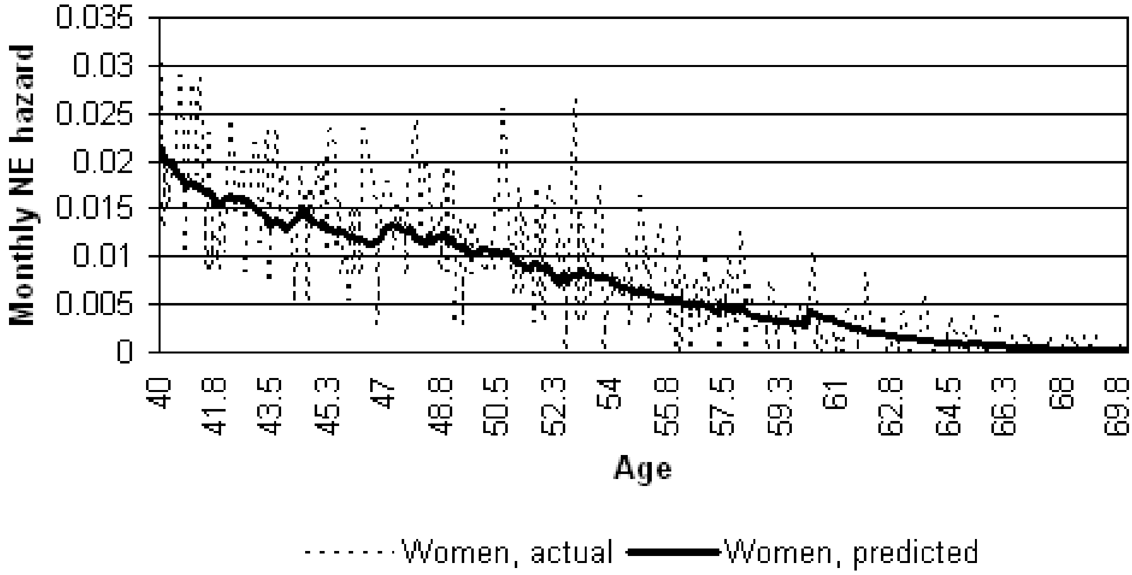

Figure 21: Goodness of fit: actual vs predicted monthly NE hazard for women by age

are almost identical, also in replicating the discrete downward jumps at ages 60 and 65. This is a further indication that while unsystematic noise is not picked up by the model, the systematic component of labour market behaviour is well predicted.

5 Conclusions

In this paper, I presented the first comprehensive analysis of older men and women's labour market transitions in the United Kingdom. It is the first paper to study older UK women's labour market transitions and the effects of occupational pension membership in detail, and also the first to assess the exact size of retirement heaping at certain ages.

When looking at the empirical hazard rates out of and back to employment by age, I found very strong evidence of birthday effects or age heaping, again particularly for men but also, only somewhat less, for women. The peaks at ages 65 and 60 are markedly bigger than in the literature, both absolutely as well as relatively (*i.e.* divided by the average of the surrounding two time intervals). For women, both peaks are important (and less marked) while men basically have one major peak at age 65. The importance of these peaks seems to have risen over time also within my own data. Whether one ever had children does not seem to be associated with the empirical transition hazard. Members of occupational pension schemes are more likely to return to work than non-members.

In my regression analysis, I present several interesting findings. First, marital status does not seem to be a very important determinant of older people's labour market transitions in the UK. Second, employment experience in younger years decreases people's exit risk and increases the return probability, but the effects are not that large. Third, women seem to face a generally lower transition risk in times of high regional unemployment. This may call for further investigation. Fourth, the impact of health problems is very important and bigger than found in previous studies. Fifth, OP membership has the expected effects for women, but the size of the effects is not that large. For men, the effects of OP membership are either difficult to explain (exit) or not statistically significant (return). Sixth, having ever had a child is associated with a lower labour market exit risk for women. Seventh, in line with the literature, the earning potential is found to be not that important in influencing one's transition probabilities, but the benefit entitlement is found to have an important impact on the return hazard which is larger than in previous analyses.

The parameters of my specification of unobserved heterogeneity are all highly statistically significant, but, similar to earlier results, it does not seem to matter a lot for the general picture of the coefficient estimates. Twice as many women as men are in the group with less stable employment patterns ('movers').

A lot of research remains to be done, and there are many open questions. How do older men and women who live together in the United Kingdom coordinate their retirement decisions? What can we say about the importance of part-time work and self-employment among older people in the UK? What is the role of earnings, benefits, and assets of other household members? How could we decompose the effect of non-labour income into genuine benefit income and asset income? How can we explain the odd effect of occupational pension membership on men's exit—for instance with a wealth effect? Why do women have fewer transitions in and out of work when regional unemployment is high? Why does the small group of women who never had a child have a higher risk to exit the labour market? Future research will have to address these questions.

Natural logarithm of normal total personal gross income	Men Working	Men Not working	Women Working	Women Not working
Calendar time	0.0012	0.0020	-0.0038***	0.0064
Calendar time squared	-0.0022	-0.0074	0.0326***	-0.0145
Before	-0.3351***	-1.5116***	0.0757**	-0.3791***
Age	0.0002	0.0031***	-0.0013***	-0.0017**
Age squared	-0.0020***	-0.0041**	0.0039***	0.0090***
Occ: employers/managers (large)	1.1976***		1.5537***	
Occ: employers/managers (small)	0.8519***		1.1182***	
Occ: professional (self and emp)	1.1339***		1.5522***	
Occ: int. non-manual, work	0.8337***		1.1512***	
Occ: int. non-manual, foreman	0.6793***		0.9450***	
Occ: junior non-manual	0.3748***		0.5605***	
Occ: personal service wkr	-0.1499***		0.1902***	
Occ: foreman manual	0.6502***		0.8182***	
Occ: skilled manual wkrs	0.4807***		0.7606***	
Occ: semi-sk manual/own acc wkrs	0.3591***		0.5447***	
Occ: farmers – employers	0.6252***		0.9359**	
Occ: farmers – own acc/agric wkrs	0.2556***		0.3022***	
Occ: HM Forces	1.0651***		1.2334***	
Constant	5.3016***	4.5087***	4.7774***	3.6356***
Number of obs	11,854	6,800	11,669	10,128
R-squared	0.3594	0.1881	0.3218	0.0480

Table 4: FES income predictions for working and non-working men and women.
***: Statistically significant at the 1% level; **: 5% level; *: 10% level (standard errors clustered by household ID).

Appendix A: Income Regressions

In Table 4, this Appendix presents the results of the FES income regressions. These coefficients have then been used on the BHPS to predict income in and out of work for each BHPS respondent.

- Time is measured in months since 1900.
- 'before' is a dummy variable to deal with a structural break in the FES *normalginc* time series between 1993 (up to which the FES referred to calendar years) and 1994–95 (from which on it referred to financial years). In my income predictions, I do not use the coefficient of *before*, *i.e.*, I am predicting on a post-1993 basis.
- Age is measured in months.
- The occupational base category is *unskilled manual*.

- These predictions use data which I adjusted for inflation as described in Section 3 of this paper.

Appendix B: Standard Error Correction

When estimating a regression model where some of the explanatory variables have been generated through other (so-called first stage) regressions, the standard errors of the so-called second stage will be underestimated. There are several estimators of corrected standard errors. Greene (2003) presents the estimator by Murphy and Topel (1985):

$$V_2^{corr} = V_2 + V_2(CV_1C^T - RV_1C^T - CV_1R^T)V_2 \tag{10}$$

where V_2 is the asymptotic variance matrix of the second stage, V_1 the asymptotic variance matrix of the first stage, and C and R matrices based on derivatives of the first- and second-stage log likelihood function with respect to the first- and second-stage parameter vectors θ_1 and θ_2:

$$C = \left\{ \left(\frac{\partial \ln L_2}{\partial \theta_2}\right) \quad \left(\frac{\partial \ln L_2}{\partial \theta_1^T}\right) \right\} \tag{11}$$

$$R = \left\{ \left(\frac{\partial \ln L_2}{\partial \theta_2}\right) \quad \left(\frac{\partial \ln L_1}{\partial \theta_1^T}\right) \right\} \tag{12}$$

If the first and second stage models are estimated on different data sets, $R = 0$ will be zero, leading to the shorter expression

$$V_2^{corr} = V_2 + V_2CV_1C^TV_2 \tag{13}$$

In my case, the second-stage models are the hazard regressions (cloglog). I have two OLS first-stage regressions, one for earnings and one for benefits. If there are two generated variables and assuming that the two first-stage regressions are independent from each other,

$$CV_1C^T = \begin{bmatrix} C_1 & C_2 \end{bmatrix} \begin{bmatrix} V_{11} & 0 \\ 0 & V_{12} \end{bmatrix} \begin{bmatrix} C_1^T \\ C_2^T \end{bmatrix} = C_1V_{11}C_1^T + C_2V_{12}C_2^T \tag{14}$$

where C_1 and V_{11} belong to the earnings regression, C_2 and V_{12} to the benefits regression.

I present the Murphy-Topel standard errors for both models, for the one without unobserved heterogeneity as well as for the one with. It is important to note that in the latter case these standard errors are not consistent but since there is no commonly used consistent estimator of the standard errors for models with generated regressors and

Heckman-Singer type unobserved heterogeneity (cf. Meghir and Whitehouse 1997: 341), I still present them. Since the correction lead to somewhat larger standard errors but not to any major changes we can be reasonably confident about them.

References

[1] Abbring, J. H. and van den Berg, G. J. (2003). 'The Unobserved Heterogeneity Distribution in Duration Analysis', Unpublished manuscript, Amsterdam: Department of Economics, Vrije Universiteit Amsterdam. Internet: `http://staff.feweb.vu.nl/jabbring/papers/mphl0312.pdf`

[2] Arellano, M. and Meghir, C. (1992). 'Female Labour Supply and On-the-Job Search: An Empirical Model Estimated Using Complementary Data Sets', *Review of Economic Studies*, vol. 59(3), pp. 537–559.

[3] Bardasi, E. (2002). BHPS Work History File, Computer Files, Colchester: Institute for Social and Economic Research, University of Essex.

[4] Blau, D. M. (1994). 'Labor Force Dynamics of Older Men', *Econometrica*, vol. 62(1), pp. 117–156.

[5] Campbell, N. (1999). 'The Decline of Employment Among Older People in Britain', CASEpaper CASE/19, London: Centre for Analysis of Social Exclusion, London School of Economics and Political Science.

[6] Disney, R., Meghir, C., and Whitehouse, E. (1994). 'Retirement Behaviour in Britain', *Fiscal Studies*, vol. 15(1), pp. 24–43.

[7] Elias, P. (1997). 'Who Forgot They Were Unemployed?', Working Paper 97-19, Colchester: ESRC Research Centre on Micro-social Change, University of Essex.

[8] Emmerson, C. (2003). 'Pension reform in the United Kingdom: Increasing the role of private provision?', in: Clark, G. L. and Whiteside, N., *Pension Security in the 21st Century*, Oxford: Oxford University Press, pp. 168–192.

[9] Flinn, C. J. and Heckman, J. J. (1983). 'Are Unemployment and Out of the Labor Force Behaviorally Distinct Labor Force States?', *Journal of Labor Economics*, vol. 1(1), pp. 28–42.

[10] Greene, W. H. (2003). *Econometric Analysis*, Upper Saddle River, NJ: Prentice-Hall.

[11] Greenwood, M. (1926). 'The "Errors of Sampling" of the Survivorship Tables', in: *Reports on Public Health and Medical Subjects No. 33: A Report on The Natural Duration of Cancer*, Appendix I, London: His Majesty's Stationery Office.

[12] Hardin, J. W. (2002). 'The robust variance estimator for two-stage models', *The Stata Journal*, vol. 2(3), pp. 253–266.

[13] Heckman, J. and Singer, B. (1984). 'A Method for Minimizing the Impact of Distributional Assumptions in Econometric Models for Duration Data', *Econometrica*, vol. 52(2), pp. 271–320.

[14] Humphrey, A. *et al.* (2003). 'Factors affecting the labour market participation of older workers', Research Report No 200, Department for Work and Pensions. Internet: `http://www.dwp.gov.uk/asd/asd5/rports2003-2004/rrep200.asp`

[15] Jenkins, S. P. (2004). *Survival Analysis*, Unpublished manuscript, Colchester: Institute for Social and Economic Research, University of Essex. Internet: `http://www.iser.essex.ac.uk/teaching/degree/stephenj/ec968/pdfs/ec968lnotesv5.pdf`

[16] Jones, S. R. G. and Riddell, W. C. (1999). 'The Measurement of Unemployment: An Empirical Approach', *Econometrica*, vol. 67(1), pp. 147–161.

[17] Kingston, G. H. (2000). 'Efficient Timing of Retirement', *Review of Economic Dynamics*, vol. 3(4), pp. 831–840.

[18] Lancaster, T. (1992). *The Econometric Analysis of Transition Data*, Cambridge, UK: Cambridge University Press.

[19] Lazear, E. P. (1986). 'Retirement from the Labor Force', in: Ashenfelter, O. C. and Layard, R., *Handbook of Labor Economics*, vol. I, Amsterdam: North-Holland.

[20] Mastrobuoni, G. (2000). 'Public pensions in Italy and Germany: A comparison based on panel data', *Rivista di Politica Economica*, vol. 12, pp. 37–78.

[21] Meghir, C. and Whitehouse, E. (1997). 'Labour market transitions and retirement of men in the UK', *Journal of Econometrics*, vol. 79(2), pp. 327–354.

[22] Morgenstern, R. D. and Barrett, N. S. (1974). 'The Retrospective Bias in Unemployment Reporting by Sex, Race and Age', *Journal of the American Statistical Association*, vol. 69(346), pp. 355–357.

[23] Murphy, K. M. and Topel, R. H. (1985). 'Estimation and inference in two-step econometric models', *Journal of Business and Economic Statistics*, vol. 3(4), pp. 370–379.

[24] Neumann, G. R. (1999). 'Search Models and Duration Data', in: Pesaran, M. H. and Schmidt, P., *Handbook of Applied Econometrics*, vol. II, Oxford: Blackwell.

[25] OECD (2000). *Reforms for an Ageing Society*, Paris: Organisation for Economic Co-operation and Development.

[26] Pagan, A. (1984). 'Econometric Issues in the Analysis of Regressions with Generated Regressors', *International Economic Review*, vol. 25(1), pp. 221–247.

[27] Paull, G. (2002). 'Biases in the Reporting of Labour Market Dynamics', Working Paper 02/10, London: Institute for Fiscal Studies.

[28] Peracchi, F. and Welch, F. (1994). 'Trends in Labor Force Transitions of Older Men and Women', *Journal of Labor Economics*, vol. 12(2), pp. 210–242.

[29] Prentice, R. L. and Gloeckler, L. A. (1978). 'Regression analysis of grouped survival data with application to breast cancer data', *Biometrics*, vol. 34, pp. 57–67.

[30] Smeaton, D. and McKay, S. (2003). 'Working after State Pension Age: Quantitative Analysis', Research Report No 182, Department for Work and Pensions. Internet: `http://www.dwp.gov.uk/asd/asd5/rrep182.asp`

[31] Taylor, M. F. (ed.) with Brice, J., Buck, N., and Prentice-Lane, E. (2005). *British Household Panel Survey User Manual Volume A: Introduction, Technical Report and Appendices*, Colchester: University of Essex. Internet: `http://www.iser.essex.ac.uk/ulsc/bhps/doc/vola/`