Empirical Analysis of Household Savings Decisions in Context of Uncertainty: A cross-sectional approach

Paolo Lucchino & Justin van de Ven*

Abstract

Most empirical studies of savings behaviour that explicitly take account of the influence of uncertainty consider for identification data that describe the evolution of circumstances observed during an appreciable period of the life-course. Here we report results obtained for a dynamic programming model that has been adapted to permit identification of preference parameters using data observed at a point in time for a given population cross-section. We discuss the advantages of this approach, and our empirical results demonstrate its feasibility in context of contemporary computing technology.

Key Words: Dynamic Programming, Savings, Labor Supply

JEL Classifications: C51, C61, C63, H31

1 Introduction

Empirical analysis of intertemporal decision making is complicated by the effects of uncertainty on incentives. Where uncertainty is considered sufficiently important to warrant a central place in a structural model, then dynamic programming methods are now commonly employed. Studies of savings behaviour in this vein often limit the computational burden by focussing upon the evolving circumstances of individual birth cohorts. The computational advantage that is gained by limiting a dynamic programming model to focus on a single birth cohort in this context is, however, off-set by attendant complications associated with the empirical specification. This paper explores the benefits for empirical analysis of relaxing a dynamic programming model of savings and labour supply to project forward from a population cross-section. Matching a standard life-cycle model of the consumption/savings and labour/leisure margins to British survey data, we conclude that the advantages of projecting forward from a population cross-section can out-weigh the associated computational costs in context of contemporary computing technology.

A complex two-dimensional relationship exists between time, cohort, and age effects that characterise differences between heterogeneous population subgroups. Focussing upon the evolving circumstances of a single birth cohort is a useful way for empirical studies to cut through this complexity, as age, time and cohort effects are then described by a single dimension. Such a simplification is particularly appealing where the central subject of interest is non-trivially complex, as is often the case for decision problems that have no closed form solution. Hence, the dynamic programming literature that explores

*NIESR, 2 Dean Trench Street, London SW1P 3HE, UK. Corresponding author: jvandeven@niesr.ac.uk.

savings behaviour in the context of uncertainty has focussed predominantly upon empirical analysis of cohort-specific structural models, following the seminal study by Gourinchas & Parker (2002).

Alternative data options exist for empirical analyses that focus on cohort-specific structural models of savings behaviour. An obvious choice is to parameterise a cohort-specific model by matching it to data observed for a single birth cohort (e.g. Attanasio et al. (2005)).[1] There are, however, at least two key drawbacks of adopting this approach: it is usually difficult to obtain an adequate description of the evolving policy context; and it remains uncertain how far results obtained for a single birth cohort can be generalised to the wider population.

These two drawbacks stem from fundamental features of the empirical problem in relation to savings behaviour. An empirical analysis of household savings decisions in context of uncertainty requires for identification data observed over an appreciable period of life. The longer is the period from which data for analysis are drawn, the greater is the scope for non-trivial variation of the policy environment underlying observed behaviour. The greater is variation of the policy environment over multiple dimensions, the stronger is the proposition that such variation is likely to be an important determinant of observed behaviour.

Aspects of the policy environment that typically exhibit substantial variation with time, and which are likely to influence savings decisions of the household sector, include taxes and benefits, (pre-transfer) rates of return, variation of employment opportunities, and the changing nature of family demographics. Obtaining comprehensive (pseudo) panel data regarding all of these factors usually represents a significant challenge, and integrating them into a structural model in a coherent fashion is more challenging still. Furthermore, allowing for such variation can work to off-set any computational advantage that is derived from focussing on the circumstances of a single birth cohort. We are not aware, for example, of any dynamic programming model of household sector savings that includes an explicit account of reforms to tax and benefits policy implemented during the period of estimation.[2] Such complications hamper efforts to reflect adequately the incentives faced by individual birth cohorts, and limit the scope for associated sensitivity analysis of results obtained.

One popular way to address the question of how far results obtained for a single birth cohort generalise to the wider population is to consider the sensitivity to data reported for alternative birth cohorts (as in Attanasio et al. (2008)). Such an approach does not, however, help to mitigate difficulties associated with adequately describing the evolving policy context. Alternatively, empirical techniques can be used that are designed to estimate age-specific moments which control for time and cohort effects

[1] French (2005) applies a similar procedure, but uses regression techniques to improve estimated age profiles for his reference cohort by drawing upon data observed for near-by cohorts.

[2] An explicit allowance for evolving tax and benefits policy has, however, been implemented in microsimulation models based on functional assumptions for behaviour; see, e.g., Nelissen (1998).

(e.g. Sefton et al. (2008)). Collinearity between age, cohort and time effects requires an additional restriction to permit identification. One common restriction, suggested by Deaton (1997), is to assume that time effects average out over the long run. This assumption produces estimated age profiles that represent an average taken over all cohorts included in the panel data used for estimation. The averaging that such methods apply obscures the nature of the underlying policy environment, so that it is difficult – if not impossible – to ensure that the assumed structural specification provides an adequate representation of the incentives underlying observed behaviour.

A third approach that has been applied in the literature is designed to simplify identification of the incentives that underly observed behaviour, which is the principal drawback associated with the two alternatives referred to above. In this case, empirical analysis is based upon cross-sectional data that are adjusted to reflect assumptions about the relationship between the characteristics of the population cross-section and those of a single birth cohort (e.g. van de Ven (2010)). Focussing on cross-sectional data limits the incentives underlying observed behaviour to those that applied at a single point in time, which are relatively easy to document. The drawback of this approach, however, is that it requires strong assumptions to be made to derive a stylised relationship between the characteristics of the population cross-section and those of a single birth cohort; assumptions that are unlikely to hold in practice.[3]

The basic premise of this paper is that an overlapping-generations (OLG) model of savings behaviour can help to mitigate the two weaknesses of the existing empirical literature that are identified above. Such a model can project the circumstances of a population cross-section through time, and is therefore well adapted to considering implications for a broad segment of society. Furthermore, this modelling approach is capable of describing behaviour observed throughout the life-course at a single point in time, albeit for individuals drawn from different birth cohorts. As such, an OLG modelling approach permits preferences to be identified on cross-sectional survey data, which considerably simplifies the task of describing the incentives underlying observed behaviour. In matching such a model to contemporary survey data observed for Britain, we have found that this estimation strategy is both computationally feasible in context of contemporary computing technology, and facilitates the process of parameterisation.

Although OLG models of savings in context of uncertainty are not new (e.g. Livshits et al. (2007), Hansen & Imrohoroglu (2008), Feigenbaum (2008)), most of the associated studies focus on implications of theory, rather than the empirical task of matching models to survey data. We are not aware of any other study that exploits the empirical advantages of an OLG model of savings behaviour that we discuss above.

[3] Adjusting age profiles of income and consumption by trend growth, for example, rests upon the assumption that the economy is in a steady-state equilibrium, characterised by a stable growth path. This assumption is highly unlikely to hold for any modern economy.

In Section 2 we provide an overview of the model that is used to undertake the analysis. Details regarding the analytical mechanics that underly our empirical approach are described in Section 3. The data are described in Section 4, and results are reported in Section 5. Emphasis in this section is placed on drawing out the ways that preference parameters influence alternative observable margins, which are crucially important for parameter identification. A concluding section provides a summary and directions for further research.

2 The Structural Model

The decision unit is the nuclear family, defined as a single adult or partner couple and their dependant children. The model divides the life course of a sample of reference adults into annual increments, to consider the evolving circumstances of their families, including endogenous decisions regarding consumption, labour supply, and the decision to participate in personal pensions. These decisions are made to maximise expected lifetime utility, given a family's prevailing circumstances, its preference relation, and beliefs regarding the future. Heterogeneous circumstances of reference adults are limited to the following eight characteristics:

- year of birth	- age	- relationship status	- wage potential
- liquid wealth	- private pension wealth	- timing of pension access	- time of death

The principal innovations to the related literature are the inclusion of 'year of birth' and 'relationship status' in the list of model characteristics – year of birth is required to project forward from a population cross-section, and relationship status is a useful characteristic to capture the savings and labour supply incentives with which we are concerned. Preferences are described by a nested Constant Elasticity of Substitution specification, and expectations are rational in the sense that they are consistent with the intertemporal processes that govern individual characteristics. Three characteristics are uncertain from one year to the next (relationship status, wage potential, and time of death), and the remaining five are non-stochastic throughout the life course. This section describes the key features of the model; technical details can be found in our companion paper Luccino & van de Ven (2013).

2.1 Preference relation

Expected lifetime utility of reference adult i at age a is described by the time separable function:

$$\begin{aligned} U_{i,a} &= \frac{1}{1-\gamma}\left\{ u\left(\frac{c_{i,a}}{\theta_{i,a}}, l_{i,a}\right)^{1-\gamma} + \right. \\ &\left. +E_a\left[\sum_{j=a+1}^{A}\delta^{j-a}\left(\phi^b_{j-a,a}u\left(\frac{c_{i,j}}{\theta_{i,j}}, l_{i,j}\right)^{1-\gamma} + (1-\phi^b_{j-a,a})\left(\zeta_0+\zeta_1 w^+_{i,j}\right)^{1-\gamma}\right)\right]\right\} \end{aligned} \quad (1)$$

where $\gamma > 0$ is the (constant) coefficient of relative risk aversion; E_a is the expectations operator; A is the maximum potential age; δ is an exponential discount factor; $\phi^b_{j-a,a}$ is the probability of reference adult from birth year b surviving to age j given survival to age a; $c_{i,a} \in R^+$ is discretionary composite (non-durable) consumption; $l_{i,a} \in [0,1]$ is the proportion of family time spent in leisure; $\theta_{i,a} \in R^+$ is adult equivalent size based on the "revised" or "modified" OECD scale; the parameters ζ_0 and ζ_1 reflect the "warm-glow" model of bequests; and $w^+_{i,a} \in R^+$ is net liquid wealth when this is positive and zero otherwise.

Each adult family member has three discrete labour supply alternatives, representing full-time, part-time, and non-employment. Although our central interest is in savings behaviour, we also include labour supply to reflect the endogeneity that exists between these two decision margins.

The modified OECD scale assigns a value of 1.0 to the family reference person, 0.5 to their spouse (if one is present), and 0.3 to each dependent child. Although spouses are modelled explicitly, the number of dependent children in each family is modelled as a deterministic function of the age, relationship status, and birth year of the reference person (discussed in Section 2.5). The OECD scale is currently the standard scale for adjusting before housing costs incomes in European Union countries, and is included here to reflect the impact that family size has been found to have on the timing of consumption (e.g. Attanasio & Weber (1995) and Blundell et al. (1994)).[4]

The warm-glow model of bequests simplifies the associated analytical problem, relative to alternatives that have been considered in the literature.[5] Including a bequest motive in the model raises the natural counter-party question of who receives the legacies that are left. We focus only upon decisions regarding bequests made, and do not consider bequests received as this would add to both the state space of the decision problem, and the dimensionality of modelled uncertainty.

A Constant Elasticity of Substitution function was selected for within period utility,

$$u\left(\frac{c_{i,a}}{\theta_{i,a}}, l_{i,a}\right) = \left(\left(\frac{c_{i,a}}{\theta_{i,a}}\right)^{(1-1/\varepsilon)} + \alpha^{1/\varepsilon} l_{i,a}^{(1-1/\varepsilon)}\right)^{\frac{1}{1-1/\varepsilon}} \tag{2}$$

where $\varepsilon > 0$ is the (period specific) elasticity of substitution between equivalised consumption $(c_{i,a}/\theta_{i,a})$ and leisure $(l_{i,a})$. The constant $\alpha > 0$ is referred to as the utility price of leisure. This specification of preferences is standard in the related literature.

[4] An empirical study by Fernandez-Villaverde & Krueger (2006) of US data from the Consumer Expenditure Survey suggests that roughly half of the variation observed for lifetime household consumption can be explained by changes in household size, as described by equivalence scales. See Balcer & Sadka (1986) and Muellbauer & van de Ven (2004) on the use of this form of adjustment for household size in the utility function.

[5] See, for example, Andreoni (1989) for details regarding the warm-glow model.

2.2 The wealth constraint

Equation (1) is maximised, subject to an age specific credit constraint imposed on liquid net worth, $w_{i,a} \geq D_a$ for reference adult i at age a. In context of income uncertainty, and a preference relation where marginal utility approaches infinity as consumption tends toward zero, rational individuals will never choose to take on debt equal to or greater than the discounted present value of the minimum potential future income stream (however unlikely that stream might be). This rule is used to define D_a, subject to the additional constraint that all debts be repaid by age $a_D \leq A$. Intertemporal variation of $w_{i,a}$ is, in most periods, described by the simple accounting identity:

$$w_{i,a} = w_{i,a-1} - c_{i,a-1} + \tau_{i,a-1} \tag{3}$$

where τ denotes disposable income net of non-discretionary expenditure. The only potential departures from equation (3) occur when a reference adult is identified as getting married or encountering a marital dissolution, in which case wealth is adjusted by a fixed factor to reflect the proportion of family assets accruing to the spouse. We treat cohabitating partners symmetrically throughout the analysis, so that marriage doubles liquid wealth, and marital dissolution halves it.

The tax function assumed for the model can be represented by:

$$\tau_{i,a} = \tau(l_{i,a}, x_{i,a}, n_{i,a}, n_{i,a}^c, r_{i,a}w_{i,a}, pc_{i,a}, b, a) \tag{4}$$

which depends on labour supply, $l_{i,a}$; private non-capital income, $x_{i,a}$; the number and age of adults, $n_{i,a}$, a; the number of dependent children, $n_{i,a}^c$; the return to liquid assets, $r_{i,a}w_{i,a}$ (which is negative when $w_{i,a} < 0$); private pension contributions, $pc_{i,a}$; and birth year, b.

Non-capital income $x_{i,a}$ is equal to labour income $g_{i,a}$ plus pension annuity income. Non-capital income is split between adult family members to reflect the taxation of individual incomes in the UK.

The interest rate, $r_{i,a}$, is assumed to depend upon whether $w_{i,a}$ indicates net investment assets, or net debts. Where $w_{i,a}$ is (weakly) positive, then r takes a fixed value r^I. When $w_{i,a}$ is (strictly) negative, then r is designed to vary from r_l^D at low measures of debt to r_u^D when debt exceeds the value of working full time for one period $\left(g^{ft}\right)$:

$$r^s = \begin{cases} r^I & \text{if } w^s \geq 0 \\ r_l^D + \left(r_u^D - r_l^D\right) \min\left\{\frac{-w^s}{g^{ft}}, 1\right\}, r_l^D < r_u^D & \text{if } w^s < 0 \end{cases} \tag{5}$$

Specifying $r_l^D < r_u^D$ reflects a so-called 'soft' credit constraint in which interest charges increase with loan size.

2.3 Private Pensions

Private pensions are modelled as defined contribution schemes administered at the family level. In each year, a family with earnings exceeding a minimum threshold, g_l^P, can choose whether to make fresh contributions to its pension scheme. If a family chooses to contribute to its pension, then a fixed share of its total pre-tax labour income, π^P, is added to its accumulated pension fund. Contributing families also receive an employer contribution to their pension fund, which is specified as a fixed share of pre-tax labour income, π_{ec}^P. Eligible employer contributions to a family's pension fund in any given year are lost if the family chooses not to contribute to its scheme in the respective year. Wealth held in a private pension fund, $w_{i,a}^P$, is assumed to be illiquid, and attracts a fixed rate of return r^P. In most periods prior to pension receipt, pension wealth follows the accounting identity:

$$w_{i,a}^P = r^P w_{i,a-1}^P + \left(\pi^P + \pi_{ec}^P\right) g_{i,a-1} \phi_{i,a-1}^P \tag{6}$$

where $\phi_{i,a-1}^P$ is an indicator variable, equal to one if the family of reference adult i at age a contributes to its pension, and zero otherwise. As with liquid wealth, the only departure from equation (6) is following relationship transitions, when pension wealth is adjusted by a fixed factor to reflect the proportion of family assets accruing to the spouse. As with liquid wealth, relationship formation doubles pension wealth, while relationship dissolution halves it.

A family can choose to access its accumulated pension fund at any time between ages 55 and 70, at which time 25% of the fund is taken as a tax-free lump-sum and the remainder used to purchase an inflation adjusted life annuity (which is assessable for income tax as discussed above). The annuity rates assumed for analysis are calculated with reference to the survival rates assumed for individual birth cohorts, an assumed return to capital, and an assumed transaction cost levied at the time of purchase.

This specification of private pension opportunities is designed broadly to reflect the terms of occupational pension schemes administered in the UK. Although the terms of most occupational pensions – regarding contribution rates, investment strategies, and draw-down options – have been highly opaque and poorly understood, rates of (voluntary) participation have traditionally been quite high, so that occupational pensions accounted for **% of all private pension wealth in 2006. One reason for this is that employers have subsidised participation in their sponsored schemes, represented in the model by π_{ec}^P.

2.4 Labour income dynamics

Wages are modelled at the family level, which helps to simplify the dynamic programming problem. The labour income of the family of reference adult i at age a, $g_{i,a}$, depends upon the family's latent wage, $h_{i,a}$, the wage offer received, $o_{i,a}$, its labour supply decision, $l_{i,a}$, and whether the family has started to draw on its private pension wealth, $ret_{i,a}$:

$$\begin{aligned} g_{i,a} &= \lambda_{i,a} h_{i,a} \\ \lambda_{i,a} &= \lambda^{wo}(o_{i,a})\,\lambda^{emp}(l_{i,a})\,\lambda^{r}(ret_{i,a}, a) \end{aligned} \tag{7}$$

We describe each of these factors in turn.

Wage offers

Wage offers are included in the model to allow for the possibility of involuntary unemployment, which we have previously found to be important in matching the model to rates of employment during peak working years. After experimenting with alternative options, our analysis considers only wage offers received by reference adults. Receipt of a wage offer is modelled as uncertain between one period and the next, subject to age and relationship specific probabilities $p^{o}(n_{i,a}, a)$. If a wage offer is received by a reference adult, $o_{i,a} = 1$, then the family's labour income is an increasing function of the reference adult's labour supply. If a wage offer is not received by a reference adult, $o_{i,a} = 0$, then the family's labour income depends only on the labour supply decision of the reference adult's spouse (if one exists), implying non-employment of the reference adult where working incurs a leisure penalty.

Labour supply decisions

Each adult in a family chooses one of three labour options corresponding to full-time, part-time, and non-employment. Increased labour supply results in a lower fraction of time enjoyed as leisure, and a higher fraction of the family's latent wage received as labour income (assuming that a wage offer is received). Each adult in a couple is treated symmetrically, and the effects of part-time employment on leisure and earnings are specified as fractions of full-time employment, π^{l}_{pt} for leisure and π^{g}_{pt} for earnings.

Accessing pension wealth

The model is defined to permit a wage penalty to be imposed on families that have previously accessed their private pension wealth, where the wage penalty can be made age dependent: $\lambda^{r}(ret_{i,a}, a)$. This feature was added to the model after observing a substantive propensity in its absence for families to access their pension wealth and draw down their labour supply prior to state pension age, only to return to work following state pension age. The wage penalty that is discussed here is designed to reflect the reduced labour market opportunities that often accrue to people who choose to take early retirement.

Latent wages

A family's latent wage during most of the life-course is assumed to follow the stochastic process described by equation (8):

$$\log\left(\frac{h_{i,a}}{m\left(n_{i,a},b,a\right)}\right) = \psi\left(n_{i,a-1}\right)\log\left(\frac{h_{i,a-1}}{m\left(n_{i,a-1},b,a-1\right)}\right) + \kappa\left(n_{i,a-1}\right)\frac{\left(1-l_{i,a-1}\right)}{\left(1-l_W\right)} + \omega_{i,a-1} \quad (8a)$$

$$\omega_{i,a} \sim N\left(0,\sigma^2_{\omega,n_{i,a-1}}\right) \quad (8b)$$

where the parameters $m\left(.\right)$ account for wage growth, $\psi\left(.\right)$ accounts for time persistence in earnings, $\kappa\left(.\right)$ is the return to another year of experience, and $\omega_{i,a}$ is a family specific disturbance term. All model parameters vary by relationship status ($n_{i,a}$), and the wage growth parameters also depend upon the age and birth year of the reference adult. The only exceptions to this specification are when a reference adult enters or departs a cohabitating relationship, in which case we assume equal latent wages between spouses (so that h doubles following marriage formation and halves following marital dissolution).

The form of equation (8) has a number of desirable properties that motivated its selection. First, it is a parsimonious wage specification that has been explored at length in the literature (e.g. Creedy (1985)). It requires the addition of just two state variables to the decision problem (h,ω), only one of which is uncertain (ω). Secondly, the appearance of the $m\left(.\right)$ terms on both sides of equation (8) helps to simplify parameterisation of the model. Increasing $m\left(a\right)$ by 3 percentage points, for example, will *ceteris paribus* increase $h_{i,a}$ by 3 percentage points without also feeding through to increase $h_{i,a+1}$. This property is lost if the $m\left(.\right)$ terms are replaced by a single factor on the right-hand-side of equation (8). And thirdly, we have found that the addition of an experience effect to the wage equation helps to match the model to the age profile of labour supply (e.g. Sefton & van de Ven (2004)). This issue is taken up in Section 3.2.

2.5 Allowing for family demographics

Family demographics depend upon the survival probabilities of reference adults, the intertemporal evolution of the relationship status of reference adults, and the allowance made for dependent children.

Modelling survival

The model focusses upon survival with respect to reference adults only; the mortality of the spouses of reference adults is aggregated with divorce to obtain the probabilities of a relationship dissolution (discussed below). Survival in the model is governed by age and year specific mortality rates, which are commonly reported components of official life-tables.

Modelling relationship status

A 'relationship' is defined as a cohabitating partnership (including formal marriages and civil partnerships). The relationship status of each reference adult in each prospective year is considered to be uncertain. The transition probabilities that govern relationship transitions depend upon a reference adult's existing relationship status, their age, and birth year. These probabilities are stored in a series of 'transition matrices', each cell of which refers to a discrete relationship/age/birth year combination.

The model does not distinguish family units on the basis of characteristics that are specific to the partners of reference adults. This approach assumes that any difference between families that is due to spouse-specific attributes is captured by other aspects of the model. The age distribution of partners of reference adults who are currently 70 years old, for example, will be implicit in the assumed probabilities that 70 year old reference adults experience a marital dissolution (due to the associated influence on spousal mortality). Similarly, the wage dispersion of spouses will be reflected in the variance of the innovation term assumed for the latent wages of couples in the model ($\sigma^2_{\omega,n_{i,a-1}}$ in equation 8).

Modelling children

Children are modelled as deterministic functions of a reference adult's age, relationship status and birth year. Non-parametric functions are assumed for dependent children, with a separate dummy variable representing each relationship/age/birth year combination. Hence, all reference adults with the same birth year, age, and relationship status are also assumed to have the same number of dependent children, which may take a non-integer value.

3 Basic Mechanics of the Empirical Approach

In common with the existing dynamic programming literature, a two stage procedure was used to identify parameters that match our structural model to survey data.[6] The first stage identified a subset of parameters exogenously from the model structure. These parameters were calculated using empirical techniques that have changed little since the advent in the 1960s of 'classical' dynamic microsimulation models. Given the model parameters evaluated in the first stage, remaining model parameters were adjusted in a second stage so that selected 'simulated moments' implied by the structural model matched to 'sample moments' estimated on survey data. Conceptually, the second stage of the procedure involves adjusting unobserved model parameters to ensure that the behaviour implied by the assumed theoretical framework best reflects behaviour observed in practice.

The principal departure between the analysis reported in this paper and the related literature is

[6] This two-step procedure is well adapted to the non-trivial computation times required to determine the implications of a given parameter combination and the large number of parameters involved.

that the model described in Section 2 is designed to consider the decisions of a population cross-section, rather than of a single birth cohort. Our analysis is motivated by the proposition that this cross-sectional approach facilitates evaluation of the second stage of the procedure that is referred to above. In this section we describe how we have implement this second stage, with emphasis on the relative advantages of taking a cross-sectional perspective.

3.1 Evaluating simulated moments

The population moments implied by the assumed theoretical framework for any given parameter combination were calculated on micro data obtained by: (i) solving the lifetime decision problem for any conceivable combination of individual specific circumstances; and (ii) using the solutions obtained in (i) to project a reference population through time.

Solving the decision problem

An analytical solution to the utility maximisation problem described in Section 2 does not exist, and numerical solution routines were consequently employed. These solution routines are structured around a 'grid' that over-lays all feasible combinations of individual specific characteristics (the state space).[7] As noted in Section 2.1, the model assumes that there is a maximum potential age to which any individual may survive, A. At this age, the decision problem is deterministic, and trivial to solve. The solution routine that we employed starts by solving for utility maximising decisions at all intersections of the grid that correspond to this final period of life, and stores both the maximising decisions and optimised measures of utility (the value function). These solutions at grid intersections for age A are used to approximate solutions at age A more generally, via the linear interpolation routine that is described in Keys (1981).

Given results for age A, the solution routine that we used then considers decisions at intersections corresponding to the penultimate age, $A-1$. Here, expected lifetime utility is comprised of the utility enjoyed at age $A-1$, and the impact that decisions taken at age $A-1$ have on circumstances – and therefore utility – at age A. Given any decision set at age $A-1$, d_{A-1}, the solution routine projects forward the set of individual specific characteristics at age A, z_A, that is implied by the processes assumed to govern intertemporal transitions (e.g. equation 3 for wealth or equation 8 for wage potential). If characteristics at age A are uncertain, then each potential characteristic vector z_A^p is projected forward with an assigned probability pr_A^p. Uncertainty in the model is either between a

[7] The grid assumed for analysis has the following dimensions: 26 points for liquid wealth between ages 20 and 74, and 151 points between ages 75 and 130; 26 points for earnings potential between ages 20 and 74; 16 points for private pension rights from age 20 to 74, and 151 points between ages 75 and 130; 2 points for wage offers between ages 20 and 74; 2 points for pension receipt from age 55 to 75; 2 points for relationship status from age 20 to age 89. Hence, the grid considered for analysis comprised 4,909,273 individual cells. This problem was solved in 7.8 minutes on a desktop workstation purchased in 2011.

discrete set of alternatives (relationship status, wage offers, and death), or over a continuous normal distribution (wage potential). Expectations over continuous normal distributions were approximated at 5 discrete points, using weights and abscissae implied by the Gauss-Hermite quadrature (implemented following Press et al. (1986)). These terms, combined with a von Nueuman Morgenstern preference relation, allow the expected lifetime utility associated with any decision set d_{A-1} to be evaluated. A numerical routine (described below) was used to search over the set of feasible decisions to maximise expected lifetime utility at each intersection of the grid corresponding to age $A-1$. These solutions, and the associated measures of optimised utility are stored, and the solution routine then considers the next preceding age. Repeated application of this procedure obtained a numerical approximation of the solution to the lifetime decision problem at all intersections of the grid spanning the feasible state space.

The numerical search routine that was employed for this study is adapted to the decisions that are considered for analysis. As described in Section 2, families are assumed to decide over one continuous domain relating to the consumption/savings margin, and a series of decisions over discrete alternatives relating to labour supply, pension participation, and the take-up of pension benefits. The model uses Brent's method, as described in Press et al. (1986), to search over the eligible consumption domain for a local maximum to expected lifetime utility, for each potential combination of discrete decision alternatives. Of all feasible alternative solutions, the one associated with the maximum numerical approximation of expected lifetime utility was taken as the solution to the lifetime decision problem. A routine that guards against the identification of sub-optimal local maxima (described in Appendix A) was also run to test the robustness of the results that are reported here.

Calculating simulated moments

The simulated moments used to guide adjustment of the model's parameters were calculated on data projected for a population of reference adults drawn from a nationally representative cross-sectional survey. The circumstances of each reference adult described by the survey were used to locate them within the grid structure that is referred to above. Given their respective grid co-ordinates, the linear interpolation methods that are also mentioned above were used to approximate each reference adult's utility maximising decision set, as implied by the numerical solutions identified at grid intersections. Given a family's characteristics (state variables) and behaviour, its characteristics were projected through time following the processes that are considered to govern their intertemporal variation. Where these processes depend upon stochastic terms, random draws were taken from their defined distributions in a process that is common in the microsimulation literature (sometimes referred to as Monte Carlo simulation).

3.2 Adjusting model parameters

The second stage of the model parameterisation involved identifying the parameters of the assumed preference relation and a selected set of parameters governing intertemporal evolution of latent wages. Preference parameters are unobservable, and are consequently prime candidates for the second stage of the parameterisation. Although wages are observable, our allowance for an experience effect made behavioural responses to the wage specification an important consideration for associated identification.

This stage of the empirical analysis is commonly conducted either by manual calibration or optimisation of a loss function using an econometric criterion.[8] Our focus in this paper is to explore how model parameters influence implied moments for a population cross-section. The results reported here were consequently obtained via a series of manual adjustments of model parameters, guided by graphical representations and sums of squared errors for a set of age specific population moments, following the approach by Sefton et al. (2008). Although this approach sacrifices some objectivity in the specification of model parameters, it also facilitates a detailed understanding of the behavioural implications of alternative parameter combinations with which we are principally concerned, relative to a numerical "black-box".

The assumed preference (see Section 2.1) relation includes six parameters: relative risk aversion, γ; an exponential discount factor, δ; two parameters for the warm-glow model of bequests, ζ_0 and ζ_1; the intra-temporal elasticity, ε; and the utility price of leisure, α. In contrast, the specification adopted for wages (see Section 2.4) includes a very large number of parameters. The persistence of latent wages, ψ, and the factor effects of alternative labour supply decisions, λ^{emp}, were identified in the first stage of the model parameterisation. This left the parameters governing wage growth, $m(.)$, experience $\kappa(.)$, earnings volatility $\sigma^2_{\omega,n_{i,a}}$, and the factor effects of pension take-up λ^r to be identified in the second stage of the parameterisation. Following extensive experimentation, we settled upon the following step-wise procedure to identify these various parameters.

We divided the calibrated model parameters into two sets; set A comprising the parameters governing wage growth and earnings volatility, and set B comprising all other calibrated parameters. We began by setting all wage growth parameters $m(.) = 1$, and made initial guesses for earnings volatility, $\sigma^2_{\omega,n_{i,a}}$. Given these assumptions for set A parameters, and the model parameters identified exogenously from the model structure in the first stage of the analysis (discussed above), we adjusted the parameters in set B to reflect behaviour observed at a single point in time for a reference population cross-section. Having obtained first approximations for set B parameters, we then adjusted the parameters $m(.)$ and $\sigma^2_{\omega,n_{i,a}}$ to reflect historical earnings data. This procedure was then repeated until convergence between

[8] Econometric methods include Simulated Minimum Distance (Lee and Ingram, 1991), Method of Simulated Moments (Stern, 1997), Indirect Estimation (Gourieroux et al., 1993) and Efficient Method of Moments (Gallant and Tauchen, 1996).

the two sets of parameters A and B was obtained. We found that it was not necessary to iterate between these two sets of model parameters more than two times to obtain convergence, a property that is attributable to the invariance of the cross-sectional population characteristics considered for adjusting parameters in set A, which we discuss further below.

Parameters identified on cross-sectional data

All six preference parameters of the model (see Section 2.1), the experience effects on earnings, κ, and the factor effects of pension take-up, λ^{ret} (see Section 2.4), were identified on behaviour observed for a single population cross-section. This is notable, given that preference parameters are often a central focus of interest in the related literature. It is also extremely useful because it simplifies specification of the policy context underlying the behaviour considered for identification, and omits feed-back effects that can otherwise complicate adjustment of parameters.

The feed-back effects mentioned above complicate any empirical analysis that refers to dynamic behaviour observed over time. Suppose, for example, that were were interested in matching a structural model of savings and retirement to data observed during an appreciable period of the life-course. If a given set of model parameters implied savings early in the life course that over-stated observed data, then this might suggest that preferences should reflect greater impatience. Adjusting preferences in this way would imply lower wealth later in life, which might then influence model implications for the timing of retirement. Such feed-back effects can be ignored in an empirical analysis of household sector savings when the analysis is based on a structural model of the population cross-section, which helps to simplify the identification problem considerably.

We started by adjusting preference parameters to match the model to age-specific moments of labour supply. Increasing the utility price of leisure α tends to decrease labour supply throughout the working lifetime. Increasing the effect of experience on earnings (κ in equation 8) increases the price of leisure early in the working lifetime, relative to later in life. And reducing the factor effect of pension take-up tends to decrease labour supply late in the working lifetime. These three model parameters consequently provide a high degree of control over the employment profile throughout the life-course.

Having obtained an approximate match to moments of employment, we next considered the level of employment income. The human capital of each reference adult is exogenously given by the cross-sectional data upon which the specification of preference parameters was based. Two key factors were consequently the focus of attention in this stage of the analysis. The first is the imputation of employment income for adults in the reference cross-sectional data who were not observed to be employed, and the second is the distribution of labour supply implied by the structural model.

Latent wages for people who were not reported to be working full-time in our reference data-set were imputed in the first stage of the model parameterisation, based on a regression specification that controls for sample selection.

The distribution of employment implied by our structural model is influenced by the intra-temporal elasticity. If the utility maximisation problem with which we are concerned were separable, then the preference relation defined by equation (2) would imply the following relationship in the region of the optimum:

$$\frac{l}{c} = \left(\frac{\hat{h}}{h}\right)^{\varepsilon} \left(\frac{\hat{l}}{\hat{c}}\right) \tag{9}$$

where population average terms are indicated by $\hat{*}$ (hats). Thus for high income households $(\hat{h}/h) < 1$, increasing the intra-temporal elasticity tends to decrease the demand for leisure relative to consumption, and vice versa for low income households. Raising ε consequently tends to lift average wages of the employed population subgroup implied by the model.

Finally, γ, δ, and the two bequest parameters $\zeta_{0/1}$, were jointly adjusted to reflect moments of consumption and pension scheme participation. Increasing the discount factor δ makes families more patient, and consequently tends to decrease consumption and increase rates of pension scheme participation throughout the working lifetime. In contrast, exaggerating the bequest motive, determined by $\zeta_{0/1}$, tends to lower consumption late in the life course when the probability of imminent mortality becomes appreciable, and reduces the incentive to participate in private pensions, as pension annuities do not feature in the bequest w^{+}. Taken together δ and $\zeta_{0/1}$ provide a high degree of control over the age profile of consumption implied by the structural model, and of pension scheme participation, when each of these behavioural margins is considered in isolation.

The addition of relative risk aversion γ is required to ensure that the model can match to moments of consumption and pension scheme participation at the same time. Put another way, the addition of pension scheme participation to the margins considered of analysis is required to identify γ. Whereas raising δ tends to imply lower consumption and higher pension scheme participation, *ceteris paribus*, raising γ tends to imply lower consumption *and* lower pension scheme participation, by exaggerating the precautionary savings motive. We consequently started with a high value of γ (=5), and adjusted δ and $\zeta_{0/1}$ to match the age profile of consumption. If the associated rates of pension scheme participation were then too low, we reduced γ and readjusted δ and $\zeta_{0/1}$, repeating this process until both consumption and rates of pension scheme participation matched at the same time.

Parameters identified on historical earnings data

The drift parameters, $m(.)$, and the dispersion parameters, $\sigma^2_{\omega,n_{i,a-1}}$, were calibrated against historical data by projecting the reference population cross-section backward through time. The drift parameters were adjusted to reflect age and year specific geometric means of employment income calculated separately for singles and couples from survey data. The model includes a separate drift parameter for each age, year, and relationship combination, so that a close match could be obtained to the associated sample moments. Given the large number of model parameters involved, this stage of the parameterisation was undertaken using an automated procedure. First, age, year, and relationship specific means of log employment income implied by the model were calculated from the simulated panel data projected back in time for the reference population cross-section. These simulated moments were subtracted from associated sample moments estimated from survey data. The differences so obtained were then multiplied by a 'dampening factor' of 0.4, and the exponent of the result was multiplied by the associated drift parameter to obtain a revised value for the parameter. The model was then re-run assuming the revised drift parameters. This procedure was repeated until the average absolute variation of parameters within any year fell below 5 percentage points.

Similarly, the variance parameters were adjusted to reflect age, year, and relationship specific variances of log employment income calculated from survey data. Unlike the drift parameters, however, only two parameters – one for singles and another for couples – were adjusted to reflect the dispersion of employment income. These two model parameters were subject to manual adjustment.

4 Survey Data

We begin by defining the cross-sectional data selected for analysis before describing the sample moments used to conduct the second stage of the model calibration.

4.1 The reference population cross-section

Data for the reference population cross-section were drawn from wave 1 of the Wealth and Assets Survey (WAS). This survey is designed to provide representative data for households and individuals in Great Britain describing demographics, income, and the distribution and type of assets and debts. As such, the survey is ideally suited for the analysis that is undertaken here. Wave 1 of the survey was drawn from the Postcode Address File, specified to reflect the population accommodated in private households in Great Britain, excluding Scotland north of the Caledonian Canal, the Scottish Islands and the Isles of Scilly. The survey was designed to over-sample from high wealth households, and information was solicited from all individuals aged 16 or over in each responding household (excluding full-time students between 16 and 18 years of age). Data were collected continually between July 2006 to June 2008, and the survey achieved a response rate of 55 percent, providing information for 71,268 individuals in 30,595 households.[9]

[9] Although data from wave 2 of the Wealth and Assets Survey were not publicly available when this analysis was conducted, these were released in the summer of 2012.

The survey data reported by the WAS were subject to three key adjustments: the sample was constrained; the units of analysis were altered; and missing earnings data were imputed. All of these adjustments were made using the Stata statistical package, and the associated code is available from the authors upon request.

Data limits imposed

The WAS sample covers the period July 2006 to June 2008. Although unemployment remained reasonably stable throughout this period, and the stock market did not crash until October 2008, measures of consumer confidence did fall substantially in the second year of the sample.[10] To avoid contaminating the calibration with behaviour observed in context of unusually pronounced negative sentiment, we therefore limit our calibration to the population cross-section observed during the year between July 2006 and June 2007.

We further restrict the WAS cross-sectional sample to omit any household with a member reported to be self-employed, due to well-recognised difficulties in evaluating their financial circumstances. Households with a member who was recorded as having a non-contributory pension scheme were also excluded to limit the heterogeneity in savings incentives described by our sample (primarily constituting public-sector employees).

Family units

The WAS organises data by survey household, which is a broader unit of analysis than the nuclear families that we focus on here. To account for this mis-match, we began by identifying family units, defined as married couples and their dependent children (under age 18), from the micro-data reported by the WAS. Most family-level statistics required for the model could be obtained by summing over the individual specific data reported by the WAS within each family unit. The notable exception was home ownership, in which case the value of the main home was allocated to the family unit of the household reference person (identified by the survey).

Gender neutrality is a guiding principal adopted for our analysis, and we consequently elected to represent each adult aged 20 or over reported by the WAS separately in our reference population cross-section. Model characteristics specific to the reference adult, including year of birth and age, were equated to the characteristics of the respective adult reported by the WAS. All other model characteristics were set equal to their values reported for each adult's family.[11] This approach meant that family characteristics of couples reported in the WAS were represented twice in our data – once for

[10] In the two years to July 2007, the GfK Consumer Confidence Index was approximately stable at -4 (negative being an indication of pessimism on average). The index fell between August 2007 and July 2008, from -4 to -39.

[11] The characteristics obtained for each family unit from the WAS were: age, relationship status (single/couple), total net non-pension wealth, full-time earnings (imputed if not reported), total private pension wealth, whether income received from private pension, and a population weighting factor.

each partner. To avoid over-sampling couples in our empirical analysis, we divided the weighting factor attached to each couple by two. The total sample size considered for analysis, following the sample selection described above, was 22,143 family units.

Imputing missing earnings

We are fortunate that the WAS provides almost all of the detail that we require for the reference population cross-section. The most notable exception applies to adults who are not reported as working full-time in the WAS, in which case the survey does not report their full earnings potential. We imputed the missing data wherever necessary, using reduced-form regressions estimated on the wider WAS data. Associated regression results are reported in Appendix C.

4.2 Sample moments

Our calibration strategy is described in Section 3.2. This strategy was implemented with reference to the following sample moments:

1. The proportion of adult family members employed full-time, part-time and not at all, by age and relationship status, estimated on data for the population cross-section observed in 2006.

2. The geometric mean of family employment income, by age and relationship status, estimated on data for the population cross-sections observed from 1971 to 2006.

3. The variance of family log employment income, by age and relationship status, estimated on data for the population cross-sections observed from 1971 to 2006.

4. The geometric mean of family consumption, by age and relationship status, estimated on data for the population cross-section observed in 2006.

5. The proportion of families reporting to contribute to private pensions, by age and relationship status, estimated on data for the population cross-section observed in 2006.

Theses sample moments were estimated on survey data from the Family Expenditure Survey (FES) and the Family Resources Survey (FRS). In common with the WAS, the FES and FRS are conducted by the Office for National Statistics, use similar sampling frames and methods, and typically achieve similar response rates to the WAS. The most significant departures between the sampling approaches implemented by the three surveys are the over-sampling of high wealth households by the WAS, and the time period covered by the respective sampling frames: while we focus on the WAS data reported for the year extending between July 2006 and June 2007, the FES reports data for the calendar year,

Table 1: Calibrated model parameters adjusted to match behaviour reported for the British population cross-section in 2006

relative risk aversion (gamma)	2.2500
intratemporal elasticity (epsilon)	0.8580
utility price of leisure (alpha)	1.7000
discount factor (delta)	0.9600
bequest motive constant (zeta0)	8.17E-09
bequest motive slope (zeta1)	4.08E-08
experience effect singles (kappa)	0.0425
experience effect couples (kappa)	0.0125
factor effects of pension take-up	
at age 55	0.6000
from age 65	0.0000

and the and FRS reports data for the UK financial year (starting in April). We ignore the mismatch between the time frames covered by these alternative data sources.

The FES is the principal source of micro-data for domestic expenditure in the UK. In addition to expenditure, it provides detailed information regarding family demographics, employment, and earnings, and covers a relatively long time-series, reporting at annual intervals from 1971. Most of the sample moments used for calibrating the model parameters were consequently estimated on FES data. The exception is participation rates in private pensions, which are more adequately described by the FRS than the FES.

5 Calibrated Preference Parameters

Here we report the calibrated model parameters that were adjusted endogenously to the structural model, and which were identified on behaviour observed for a reference population cross-section. As discussed in Section 3.2, this includes all of the parameters of the assumed preference relation, experience effects on earnings and associated factor effects of pension take-up. Specification of all other model parameters are reported in Appendices B to D.

5.1 Calibrated parameters

Calibration of the model parameters to behaviour observed for the population cross-section required testing over 185 alternative parameter combinations. Our preferred parameter set is reported in Table 1.

The calibrated value for the parameter of relative risk aversion – at 2.25 – is within the broad range identified by the associated literature. Simulations undertaken by Auerbach & Kotlikoff (1987), for example, are based upon a coefficient of risk aversion of 4, while Cooley & Prescott (1995) consider a value of 1. Grossman & Shiller (1981) and Blundell et al. (1994) report estimates just over 1.0, while

Hansen & Singleton (1983), Mankiw et al. (1985), and Ziliak & Kniesner (2005) report estimates of approximately 1. Values of the coefficient of risk aversion required to explain the equity premium puzzle (Mehra & Prescott (1985)) are high by comparison, supported by econometric estimates reported by Grossman & Shiller (1981), Mankiw (1985) and Hall (1988). Nevertheless, evidence from attitudinal surveys suggest that the value is unlikely to be greater than 5 (Barsky et al. (1997)).

The value obtained for the intra-temporal elasticity of substitution implies that consumption and leisure are direct substitutes.[12] The utility price of leisure is in the region of 1.0 by construction[13], and the discount factor implies a higher rate of time discounting than either the return to positive balances of net liquid wealth (1.5% p.a.) or pension assets (2.5% p.a.). The bequest motive increases as the bequest parameters decrease, so that the relatively small parameter values arrived at for ζ_0 and ζ_1 indicate the importance of bequests in matching the model to survey data. The experience effects, which are assumed to be age independent, imply that earnings potential increases by 4.3% for each year of full-time employment by single adults, and by 1.3% for each year of full-time employment for adults in couples. The factor effects of pension take-up apply a 40% discount to earnings at age 55, and increase linearly to 100% from age 65.

Numerical simulations indicate that the calibrated model parameters imply an inter-temporal elasticity of consumption of 0.512 measured at the population means.[14] This is in contrast to the controversial finding by Hall (1988) that the inter-temporal elasticity may not be very different from zero (see also Dynan (1993), Grossman & Shiller (1981), and Mankiw (1985)). Other studies have, however, found evidence to support substantially higher inter-temporal elasticities than reported here. Attanasio & Weber (1993), for example, find that focussing upon cohort data for individuals who are less likely to be liquidity constrained than the wider population obtains an estimate for the inter-temporal elasticity of consumption of 0.8 on UK data, and Attanasio & Weber (1995) report estimates between 0.6 and 0.7 for the US. Other empirical studies that support higher rates for the inter-temporal elasticity include Blundell et al. (1993) (0.5), Blundell et al. (1994) (0.75), Engelhardt & Kumar (2007) (0.75), Hansen & Singleton (1983) and Mankiw et al. (1985) (just over 1).

[12] The preference relation described by equations (1) and (2) implies that $U_{cl} = (1/\varepsilon - \gamma) U_c U_l / u^{1-\gamma}$, which is negative when $1/\varepsilon < \gamma$.

[13] the equivalence scale is re-scaled by a factor of 470

[14] This statistic was estimated by numerically calculating the derivative $d(\Delta \ln c_{i,t}) / d \ln r_{i,t}$, where $\Delta \ln c_{i,t} = \ln c_{i,t} - \ln c_{i,t-1}$, for the reference population cross-section. The derivative was taken by perturbing interest rates up by 0.5% (giving an elasticity estimate of 0.549), and down by 0.5% (giving an estimate of 0.474). The average between these two estimates is reported here.

5.2 Identification

Employment rates

Summary statistics describing the match between simulated and sample moments for employment are reported in Table 2, and associated dis-aggregated statistics are reported in graphical form in Appendix E. As the simulated moments tend to exhibit less variation between ages than the sample moments, we report average differences by age band between these two sets of statistics in the table, along with root-square-mean errors as an overall measure of fit. The age bands that are reported here were selected to separate out the three key life stages of work-force entry, prime working (and child-rearing) ages, and retirement.

The statistics reported for our preferred parameter combination in Table 2 indicate that the model has a systematic tendency to understate rates of part-time employment. In respect of single adults, the counterpart to lower part-time employment is higher rates of full-time and non-employment. In respect of couples, the preferred parameter set results in persistently high rates of full-time employment, and low rates of part-time and non-employment. The results reported for singles consequently suggest that some improvement may have been obtained by increasing the preference for leisure (e.g. increasing the utility price of leisure), even if this required an allowance for different parameters for singles and couples. Associated analysis in this regard indicated that increasing the preference of couples for leisure produced low geometric mean earnings early in the simulated lifetime, with associated distortions to consumption and pension scheme participation.

There are very slight differences between the measures of fit obtained with and without experience effects. This, however, obscures the fact that the experience effect helps to mitigate excessively high rates of full-time employment generated by the model for couples late in the simulated lifetime. That is, raising the experience effect allows a lower utility price of leisure to be assumed while maintaining rates of employment early in the simulated lifetime, resulting in lower rates of employment later in life. This is indicated by the relatively pronounced increase in the RMSE associated with full-time employment of singles when the experience effects are omitted from the analysis.

As expected, a high utility price of leisure tends to lower rates of full-time employment implied by the model, in favour of part-time and non-employment. The increase from 1.7 to 2.2 in this parameter fully reverses the over-statement of full-time employment and under-statement of non-employment implied by the model for couples, relative to the associated sample statistics.

The bottom panel of Table 2 reveals the important role played by wage effects of pension take-up in our model. Omitting the effects of pension take-up on wages results in a substantial increase in rates of full-time employment and decrease in rates of non-employment late in life, and a pronounced increase

Table 2: Measures of fit between simulated and sample rates of employment, by relationship status and age band

	single			couple		
age band	full-time	part-time	not emp	full-time	part-time	not emp
preferred parameter combination						
20-29	-0.085	-0.022	0.106	0.118	-0.036	-0.082
30-54	0.071	-0.060	-0.011	0.106	-0.063	-0.042
55+	0.003	-0.074	0.071	0.090	-0.090	0.000
RMSE	0.114	0.103	0.124	0.138	0.105	0.073
no experience effects						
20-29	-0.080	-0.025	0.105	0.161	-0.054	-0.107
30-54	0.097	-0.074	-0.023	0.160	-0.083	-0.077
55+	-0.006	-0.078	0.084	0.092	-0.090	-0.002
RMSE	0.123	0.109	0.133	0.162	0.106	0.090
high utility price of leisure (2.2)						
20-29	-0.137	0.021	0.115	-0.102	0.049	0.052
30-54	-0.105	0.050	0.055	-0.130	-0.057	0.187
55+	-0.042	-0.069	0.111	0.000	-0.083	0.083
RMSE	0.131	0.112	0.135	0.124	0.104	0.147
no wage effects of pension take-up						
20-29	-0.085	-0.022	0.106	0.118	-0.036	-0.082
30-54	0.071	-0.060	-0.011	0.096	-0.061	-0.035
55+	0.165	0.018	-0.183	0.423	-0.033	-0.390
RMSE	0.154	0.089	0.161	0.288	0.096	0.268

Notes: simulated less sample moments averaged over age bands

RMSE = root-mean-square-error for all ages; "not emp" = not employed

in the associated RMSEs of the model fit.

Moments of employment income

Summary statistics of the match obtained to moments of employment income are reported in Table 3; see Appendix E for disaggregated statistics in graphical form. This table indicates that the largest disparities between our preferred parameter combination and sample moments occur toward the end of the working lifetime, when the sample of employed individuals becomes small. The bottom panel of the table reports the influence of reducing the intra-temporal elasticity (ε in equation 2) on the fit obtained, which was the principal parameter adjusted to match this aspect of the data (as discussed in Section 3.2).

Moments of consumption and pension scheme participation

As discussed in Section 3.2, four key parameters were adjusted to match the model to moments of consumption and pension scheme participation: the discount factor (δ), the parameter of relative risk aversion (γ), and the parameters governing the bequest motive $\left(\zeta_{0/1}\right)$. The fit obtained to the geometric mean of consumption and rates of pension scheme participation is reported in Table 4, and disaggregated statistics are provided in Appendix E.

Our preferred parameter combination obtains a close match to the sample moments that are re-

Table 3: Measures of fit between simulated and sample moments of employment income, by relationship status and age band

	single		couple	
age band	geo mean	variance	geo mean	variance
preferred parameter combination				
20-29	-38.114	0.594	-12.394	0.016
30-54	-11.212	0.448	-12.522	-0.020
55-64	-35.990	0.293	-54.451	0.612
65+	-152.533	-0.461	52.493	-0.928
RMSE	95.497	0.752	133.880	0.638
low intra-temporal elasticity (0.8)				
20-29	-44.244	0.612	-45.428	0.091
30-54	-31.743	0.451	-66.977	0.003
55-64	-46.462	0.275	-88.199	0.676
65+	-152.533	-0.461	49.880	-0.933
RMSE	99.452	0.752	143.219	0.647

Notes: simulated less sample moments averaged over age bands
RMSE = root-mean-square-error for all ages
"geo mean" = geometric mean; "variance" = variance of log income
earnings reported in £2006 p.w.

ported here, which is notable given the relatively few model parameters that were involved. Whereas moments of consumption are relatively insensitive to the assumed discount factor, projected rates of pension scheme participation are not: decreasing the discount factor from 0.96 to 0.94 alters the RMSE associated with geometric means of consumption by less than £1 per week, at the same time as it reduces rates of pension scheme participation by approximately 5% for singles and 10% for couples (increasing the RMSE by 0.04 for each). Decreasing the parameter of relative risk aversion from 2.25 to 1.5 is shown to increase consumption and pension scheme participation appreciably. The increase of the RMSE on moments of consumption is in the region £50 per week for singles, and twice that for couples. Meanwhile, age specific rates of pension scheme participation increase on average by 0.10 for singles and 0.07 for couples. Omitting the bequest motive is shown to have a similar impact on rates of pension scheme participation as reducing the parameter of relative risk aversion from 2.25 to 1.5 discussed above. But the impact of omitting the bequest motive on consumption is shown to be appreciably higher than the other parameter perturbations that are considered here, raising age specific geometric means of consumption for cohabitating couples by £280 per week, with large effects projected later in life.

The heterogenous nature of the effects that each of these parameters have on consumption and pension scheme participation is both the source of parameter identification, and the reason for the close match obtained to the behavioural margins that are discussed here. In the absence of matching the model to rates of pension scheme participation, for example, we would have had in-adequate detail to distinguish the parameter of relative risk aversion from the parameters governing the bequest motive.

Table 4: Measures of fit between simulated and sample moments of consumption and pension scheme participation, by relationship status and age band

	single		couple	
age band	gm cons	pens part	gm cons	pens part
preferred parameter combination				
20-29	4.910	-0.089	12.790	-0.103
30-54	8.699	-0.008	-7.108	0.097
55+	-5.782	-0.041	6.265	-0.111
RMSE	27.003	0.069	39.693	0.128
low discount factor (0.94)				
20-29	0.035	-0.141	7.343	-0.220
30-54	7.908	-0.068	5.493	-0.021
55+	-8.227	-0.064	4.848	-0.125
RMSE	26.758	0.106	41.275	0.164
low relative risk aversion (1.5)				
20-29	67.724	-0.016	125.109	-0.012
30-54	82.089	0.134	117.258	0.167
55+	67.834	0.024	146.547	-0.066
RMSE	80.702	0.121	136.865	0.147
no bequest motive				
20-29	111.232	-0.031	211.093	-0.025
30-54	143.592	0.140	240.902	0.201
55+	188.038	0.022	366.694	-0.071
RMSE	162.885	0.128	296.151	0.173

Notes: simulated less sample moments averaged over age bands
RMSE = root-mean-square-error for all ages
"gm cons" = geometric mean of consumption (£2006 p.w.)
"pens part" = rates of private pension scheme participation

Put another way, our results may be understood as the product of an exactly identified system. We cannot, for example, conclude – because we have obtained a close match to rates of pension scheme participation using a model that omits associated decision costs – that such costs are unimportant in practice. Rather, if we expanded our model to include decision costs regarding pension scheme participation, then we would need to augment our calibration strategy to include a decision margin that is well suited to identify such costs. That remains a subject for another study.

6 Conclusions

This paper reports the results of a calibration conducted for a structural model of savings and labour supply to cross-sectional survey data for Britain. The structural model is specifically designed to consider behaviour for a population cross-section, with the objective of testing the conjectures that such a framework is computationally feasible on contemporary technology, and helps to address empirical complications that are discussed for cohort-specific models of behaviour. Importantly, our analysis indicates that preference parameters for a structural model of savings and employment – including the parameter of relative risk aversion – can be identified on behavioural margins observed for a population cross-section at a single point in time. We argue here that the additional complications involved in extending a dynamic programming model of savings to allow for heterogenous birth cohorts are more than off-set by the conceptual advantages derived when bringing such a model to survey data.

The model that we consider here is based upon a preference relation that is standard in the literature, and we set out a concise calibration strategy that focuses upon specific and important behavioural margins. Our preferred parameter specification is shown to match to observed behaviour over employment, income, consumption, and pension scheme participation. It is notable that we match to all of these behavioural margins through the adjustment of just 10 model parameters, six of which describe model preferences. Furthermore, the preference parameters obtained are broadly in line with those calculated in the associated empirical literature. One notable departure from stylised facts implied by the existing literature concerns the intertemporal elasticity of consumption: whereas empirical studies following the seminal work by Hall suggest that this elasticity may be close to zero, we find a value at population averages in the region of 0.5.

Parameterising a structural dynamic programming model of savings on data observed for a population cross-section at a point in time opens up a range of exciting empirical possibilities. One such possibility is to consider whether the intertemporal elasticity of substitution exhibits systematic variation with the economic cycle. This might help to explain the wide diversity of estimates that have previously been reported for this important preference parameter, with important behavioural and policy implications. Improvements in computing technology, and advancements in empirical methods will hopefully offer a wealth of opportunities during the next few decades to further our understanding of the decisions that people make.

References

Andreoni, J. (1989), ‘Giving with impure altruism: applications to charity and Ricardian equivalence’, *Journal of Political Economy* **97**, 1447–1458.

Attanasio, O., Low, H. & Sanchez-Marcos, V. (2005), ‘Female labor supply as insurance against idiosyncratic risk’, *Journal of the European Economic Association* **3**, 755–764.

Attanasio, O., Low, H. & Sanchez-Marcos, V. (2008), ‘Explaining changes in female labor supply in a life-cycle model’, *American Economic Review* **98**, 1517–1552.

Attanasio, O. P. & Weber, G. (1993), ‘Consumption growth, the interest rate and aggregation’, *Review of Economic Studies* **60**, 631–649.

Attanasio, O. P. & Weber, G. (1995), 'Is consumption growth consistent with intertemporal optimization?', *Journal of Political Economy* **103**, 1121–1157.

Auerbach, A. J. & Kotlikoff, L. J. (1987), *Dynamic Fiscal Policy*, Cambridge University Press, Cambridge.

Balcer, Y. & Sadka, E. (1986), 'Equivalence scales, horizontal equity and optimal taxation under utilitarianism', *Journal of Public Economics* **29**, 79–97.

Barsky, R. B., Kimball, M. S., Juster, F. T. & Shapiro, M. D. (1997), 'Preference parameters and behavioral heterogeneity: an experimental approach in the health and retirement survey', *Quarterly Journal of Economics* **112**, 537–579.

Blundell, R., Browning, M. & Meghir, C. (1994), 'Consumer demand and the life cycle allocation of household expenditure', *Review of Economic Studies* **61**, 57–80.

Blundell, R., Pashardes, P. & Weber, G. (1993), 'What do we learn about consumer demand patterns from micro data', *American Economic Review* **83**, 570–597.

Cooley, T. F. & Prescott, E. C. (1995), Economic growth and business cycles, *in* T. F. Cooley, ed., 'Frontiers of Business Cycles Research', Princeton University Press, Princeton.

Creedy, J. (1985), *Dynamics of Income Distribution*, Basil Blackwell, Oxford.

Deaton, A. (1997), *The Analysis of Household Surveys*, John Hopkins University Press, Baltimore.

Dynan, K. (1993), 'How prudent are consumers?', *Journal of Political Economy* **101**, 1104–1113.

Engelhardt, G. V. & Kumar, A. (2007), The elasticity of intertemporal substitution: new evidence from 401(k) participation. Federal Reserve Bank of Dallas Working Papers 0812.

Feigenbaum, J. (2008), 'Information shocks and precautionary saving', *Journal of Economic Dynamics and Control* **32**, 3917–3938.

Fernandez-Villaverde, J. & Krueger, D. (2006), Consumption over the life cycle: facts from Consumer Expenditure Survey Data. Working Paper, University of Pennsylvania.

French, E. (2005), 'The effects of health, wealth, and wages on labour supply and retirement behaviour', *Review of Economic Studies* **72**, 395–427.

Gallant, A. R. & Tauchen, G. E. (1996), 'Which moments to match?', *Econometric Theory* **12**, 657–681.

Gourieroux, C., Monfort, A. & Renault, E. (1993), 'Indirect inference', *Journal of Applied Econometrics* **8**, S85–S118.

Gourinchas, P. & Parker, J. (2002), 'Consumption over the Life Cycle', *Econometrica* **70**, 47–89.

Grossman, S. J. & Shiller, R. J. (1981), 'The determinants of the variability of stock market prices', *American Economic Review* **71**, 222–227.

Hall, R. E. (1988), Intertemporal substitution in consumption. NBER Working Paper W0720.

Hansen, G. D. & Imrohoroglu, S. (2008), 'Consumption over the life cycle: The role of annuities', *Review of Economic Dynamics* **11**, 566–583.

Hansen, P. & Singleton, K. (1983), 'Stochastic consumption, risk aversion and the temporal behaviour of asset returns', *Journal of Political Economy* **91**, 249–266.

Keys, R. G. (1981), 'Cubic convolution interpolation for digital image processing', *IEEE Transactions on Acoustics, Speech, and Signal Processing* **29**, 1153–1160.

Lane, Clark & Peacock (2008), Pension Buyouts 2008. www.lcp.uk.com.

Lee, B. S. & Ingram, B. (1991), 'Simulation estimation of time-series models', *Journal of Econometrics* **47**, 197–205.

Livshits, I., MacGee, J. & Tertilt, M. (2007), 'Consumer bankruptcy: A fresh start', *American Economic Review* **97**, 402–418.

Luccino, P. & van de Ven, J. (2013), Empirical analysis of household savings decisions in context of uncertainty; a cross-sectional approach. National Institute Disscussion Paper XXX.

Mankiw, N. G. (1985), 'Consumer durables and the real interest rate', *Review of Economics and Statistics* **62**, 353–362.

Mankiw, N. G., Rotemberg, J. J. & Summers, L. H. (1985), 'Intertemporal substitution in macroeconomics', *Quarterly Journal of Economics* **100**, 225–251.

Mehra, R. & Prescott, E. C. (1985), 'The equity premium: a puzzle', *Journal of Monetary Economics* **15**, 145–161.

Muellbauer, J. & van de Ven, J. (2004), Estimating Equivalence Scales for Tax and Benefits Systems. NIESR Working Paper 229.

Nelissen, J. H. M. (1998), 'Annual versus lifetime income redistribution by social security', *Journal of Public Economics* **68**, 223–249.

Press, W. H., Flannery, B. P., Teukolsky, S. A. & Vetterling, W. T. (1986), *Numerical Recipes: the art of scientific computing*, Cambridge University Press, Cambridge.

Sefton, J. & van de Ven, J. (2004), 'Simulating household savings and labour supply: an application of dynamic programming', *National Institute Economic Review* **188**, 56–72.

Sefton, J., van de Ven, J. & Weale, M. (2008), 'Means testing retirement benefits: fostering equity or discouraging savings?', *Economic Journal* **118**, 556–590.

Stern, S. (1997), 'Simulation based estimation', *Journal of Economic Literature* **35**, 2006–2039.

van de Ven, J. (2010), An empirical investigation of quasi-hyperbolic discounting. National Institute of Economic and Social Research Discussion Paper 355.

Ziliak, J. P. & Kniesner, T. J. (2005), 'The effect of income taxation on consumption and labor supply', *Journal of Labor Economics* **23**, 769–796.

A Routine to Guard Against Sub-optimal Local Maxima

The value function in this problem is neither smooth, nor concave (though it is increasing and continuous). This is because the labour supply decision is between discrete alternatives and the budget set is non-convex (due to means testing of welfare benefits). Non-concavities of the value function imply that the optimisation problem can have local maxima, a problem that can be addressed in a number of alternative ways. French (2005), for example, adopts a brute-force method by searching across the grid that defines all possible choice combinations of the assumed control variables. This is very time consuming, which complicates calibration and associated analysis. We consequently use a solution routine that takes advantage of the reasonably well-behaved nature of our analytical problem.

As discussed in Section 3.1, for each feasible discrete decision alternative our numerical search routine is designed to identify the measure of consumption that maximises expected lifetime utility. This routine identifies an initial candidate solution by testing over repeated consumption alternatives, based upon Brent's method. As the nature of our decision problem tends to cluster the local maxima of the decision problem, we perform a limited grid test above and below the initial candidate solution to identify alternative candidate solutions.[15] If an alternative local maximum is identified, then the solution routine extends the grid search domain above or below, depending on the location of the new

[15] The discrete nature of the labour supply decision, for example, leads to a well-recognised 'saw-tooth' profile for the first order condition, which may pass through zero at multiple – localised – points.

candidate solution, relative to the initial candidate. This process is continued until no further local maxima are found. Of all feasible solutions, the one that maximises the value function is then selected.

B Exogenously Identified Model Parameters

The model parameters for which exogenous estimates were obtained are concerned with four key issues: private income, taxation, private pensions, and demographics.

B.1 Private income

The model requires three separate (real) interest rates: the return to liquid assets, and lower and upper bounds for interest charges on debt. The return to liquid assets was set equal to 1.5% per annum, which is the average yield on long-term Gilts between 1970 and 2010 (reported by the Bank of England). The lower bound interest charge on debt was set to 8.4% per annum, which is the average interest paid on personal loans between 1995 and 2010, and the upper bound was set to 15.4% per annum, which is the average interest paid on credit card debt between 1995 and 2010 (also reported by the Bank of England).[16]

As noted in Section 3.2, the persistence of latent wages, ψ, and the factor effects of alternative labour supply decisions, λ^{emp}, were identified exogenously (see Section 2.4 for formal definitions of these parameters). The specification of latent wages was defined as a random walk with drift, so that $\psi = 1.0$. Full-time employment of all adult family members was assumed to reduce family leisure time by 40%. Furthermore, part-time employment was assumed to be equivalent to 40% of a full-time job, reducing both leisure time and earnings in the same proportion, $\lambda^{emp} = 0.4$.

B.2 Tax and benefits policy

Taxes and benefits are formally modelled on the transfer system that applied in the UK in 2006/7. Simulated transfer policy distinguishes between two periods of the life-course, subject to an age threshold set equal to state pension age $t_{SPA}(b)$. The program code adopted to simulate taxes and benefits in the model can be obtained from the authors upon request.

Taxes and benefits during the working lifetime

Prior to t_{SPA}, taxes and benefits are based upon the *Tax Benefit Model Tables* (TBMT) produced by the Department for Work and Pensions, which are designed to capture the key elements of the transfer system that applied to the healthy working-aged population. These include income taxes, national

[16] All financials reported in this study are specified in real terms, discounted to 2006 prices using the National Accounts final consumption deflator (ONS code YBGA).

insurance contributions, the working tax credit, the child tax credit, the child benefit, housing benefit, council tax benefit, Jobseeker's allowance, healthy start allowances, and free school meals.

The allowance made for child tax credit requires assumptions to be made about the child-care costs to which a family is subject. Similarly, the allowance made for housing benefit and council tax benefit require assumptions to be made about housing and council tax costs. These costs are all assumed to be non-discretionary, and are based on the assumptions reported in the April 2006 edition of the TBMT. Beyond the assumptions made by the TBMT, it was necessary to assume that child-care costs are incurred by any family with at least one dependent child, and where all adult household members work full-time.

As a brief overview, the disposable income of a family is calculated by:

1. evaluating aggregate *take-home pay* from the taxable incomes of each adult family member – this reflects the taxation of individual incomes in the UK

2. calculating *benefits* receipt (excluding adjustments for child care and housing costs) from aggregate household take-home pay – this reflects the fact that benefits tend to be provided at the level of the family unit

3. calculating non-discretionary *net child care costs* (after adjusting for child care related benefits) from aggregate take-home pay

4. calculating non-discretionary *net housing costs* (after adjusting for relevant benefits receipt) from aggregate take-home pay plus benefits less child care costs – this reflects the fact that housing benefit and council tax benefit in the UK are means tested with respect to income net of most other elements of the tax and benefits system

5. household *disposable income* is then equal to aggregate take-home pay, plus benefits, less net child care costs, less net housing costs.

Taxes and benefits from state pension age

A similar approach was taken to model taxes and benefits from state pension age as described above for the working lifetime. Unlike for the working lifetime, however, the specification of transfer payments from state pension age could not be based on the TBMT, as these do not cover retirement benefits. Rather, we referred to official rates and thresholds of the transfer schemes that applied in practice to specify this aspect of the transfer system.

Fiver transfer schemes are explicitly taken into account by the transfer system considered for analysis from state pension age. Income taxes take a step-wise rate structure similar to those applied in the working lifetime (but subject to a different tax-free minimum income threshold). The pension credit is a means-tested benefit scheme, which is withdrawn at the rate of £1 for every £1 of private income up to a minimum threshold, and then at the rate of £0.40 for every £1 of private income thereafter, until the benefit is exhausted. Housing benefits and council tax benefits are modelled in the same way as described for the working lifetime, including the associated assumption regarding the incidence of non-discretionary housing and council tax costs.

Finally, allowance is made for state contributory pension schemes. In practice, two schemes were applied in the UK in 2006; the basic state pension was subject to a maximum benefit value payable in respect of a minimum contributions history; and the state second pension was an income related benefit, rights to which were accrued in respect of national insurance contributions paid during the working lifetime. To avoid adding two state variables to the decision problem we represented both of these schemes by a single flat-rate state pension paid from state pension age, and set equal in value to the full basic state pension. This stylisation reflects the intention of policy reforms set out in the May 2006 Pensions White Paper published by the DWP, which appreciably relaxed the conditions required to obtain the full basic state pension, and severed the earnings link of the state second pension.

Transfer policy through time

Although much of the empirical analysis with which this study is concerned focusses on behaviour observed at a single point in time, the structural model upon which the analysis is based requires transfer policy to be described over an extensive time period. The transfer policy described above for 2006 is assumed to vary through time in two ways. First, the evolution of benefit values and income thresholds are assumed to describe constant growth rates. After experimenting with various alternatives, we settled upon the assumption that most of the associated growth rates equal 1.5% per annum, reflecting real earnings growth and the real return to long term government debt observed in the UK between 1970 and 2010. The key motivation for this assumption is that it ensures that the transfer system maintains pace with wages, omitting marginalisation of welfare provisions or extensive tax bracket creep. Sensitivity analysis indicated that our results are not qualitatively sensitive to the reasonable alternative of setting growth rates for benefits values and tax thresholds to reflect historical trends.

The only departure to the growth rates referred to above is the inter-temporal treatment of the rates and thresholds assumed for the Pension Credit, which is applied from state pension age (as discussed above). The Pension Credit is comprised of two elements; the Guarantee Credit, which is withdrawn at a rate of 100% in respect of private income, and the Savings Credit, which is withdrawn at a rate of 40%. We assume that the Guarantee Credit grows at 1.7% per annum and that the Savings Credit is held fixed in real terms. These assumptions are designed to reflect proposals for reform put forward by the Pensions Commission in 2005, and associated policy reforms set out in the May 2006 Pensions White Paper.

The second aspect of the policy environment that is subject to change through time concerns state pension age. The model reflects reforms reported in the May 2006 Pensions White Paper that increased the state pension age from 65 in 2020 to 68 in 2046.

B.3 Private pensions

Private pensions in the model depend upon the following parameters: the rate of return to pension wealth r^P, the minimum earnings threshold for pension contributions g_l^P, the rate of private contributions to pensions out of employment income π^P, the rate of employer pension contributions π_{ec}^P, the return assumed for calculating the price of pension annuities, and the fixed capital charge associated with purchasing a pension annuity.

There is a great deal of diversity in private pension arrangements in the UK, and in the details of occupational pensions in particular. Panel A of Figure 1 reveals that – although not universal – a sizeable majority of employees were offered some form of contribution in respect of participation in an employer sponsored pension. Eligibility to an employer sponsored pension is reported to increase from a low of between 30 and 40 percent among individuals on less than half of median earnings (increasing by age group), to between 75 and 85 percent among individuals on more than one and a half times median earnings. Approximately 60% of employees on median earnings are reported as being eligible to some form of employer pension contribution, with this rising to just over 80% for employees with very high earnings. The figure also indicates that eligibility to an employer pension contribution exhibited a stronger relationship with employee earnings than it did with age. Following these observations, we set g_l^Pequal to 75% of median earnings.

Panel B of Figure 1 indicates that, for employees who received an employer pension contribution, the distribution of employer pension contributions was dominated by a single mode between 12.5 and 15 per cent of employee wages. Bearing in mind that the decision by an employee not to participate in their employer's sponsored pension plan would usually result in the forfeiture of any matching employer pension contributions on offer, the scale of the employer contributions reported in Panel B provides an indication of how important these contributions were in supporting the UK system of private sector pension provisions. Panel B of Figure 1 also reveals that there was very little difference between the distributions of employer pension contributions offered in low-pay industries and the wider labour market, with the principal disparity being that employer contributions in excess of the mode were less frequent among employees in low pay industries. We consequently set the rate of employer contributions to 14%; the rate of private contributions to pension wealth was set to the 'normal' contribution rate stated in the guidance to interviewers for the FRS, equal to 8%.

We set the return assumed to pension wealth during the accrual phase, r^P, to 2.5% per annum, which is between the long-run real return to government debt (1.5%) and the return to equities (4.7%) observed between 1970 and 2010. The capital return assumed for calculating the price of pension annuities was set equal to 1.5%, reflecting the average rate of return to long-term government debt observed between 1970 and 2010, and the associated capital charge was set to 4.7% based on "typical" pricing margins reported in the pension buy-outs market (see Lane et al. (2008), p. 22).

B.4 Demographics

Three sets of demographic parameters were estimated exogenously from the model structure: life expectancy, described by a set of age and birth year specific probabilities of death for reference adults; relationship status, described by age and year specific transition probabilities for reference adults; and numbers of dependent children, described as age, year, and relationship specific averages for reference adults.

Life expectancy

The model requires age and birth year specific survival rates to simulate the risk of mortality. At the time of writing, the Office for National Statistics (ONS) reports period mortality rates for the UK that are distinguished by sex and age, at annual intervals between 1951 and 2060 inclusive, and between ages 0 and 100. The rates to 2010 are based on observed survival rates, and are projections thereafter. Three series of projections are reported by the ONS; a principal projection, a high life expectancy variant, and a low life expectancy variant. We focus on the principal projections here.

We assume a maximum potential age of 130 for the simulations.[17] The age specific mortality rates reported by the ONS were extended beyond age 100, using a smooth sigmoidal function to equal 1.0 (certain death) at age 130. Furthermore, the time series dimension of the age specific mortality rates reported by the ONS was extended to all age and year combinations feasible for any modelled birth cohort by assuming a constant exponential growth factor of 0.975 from the most approximate year described by the ONS data to exogenously assumed age and sex specific asymptotes for the distant past and future.

The model specification does not distinguish reference adults by their gender. The gender specific mortality rates that are reported by the ONS were consequently combined into a single series based on implied gender weights. Consider, for example, the cohort born in 1960. Assuming zero migration and

[17] The oldest age to which a human is documented to have survived is for Jeanne Calment of France, who died in 1997 at age 122 years, 164 days.

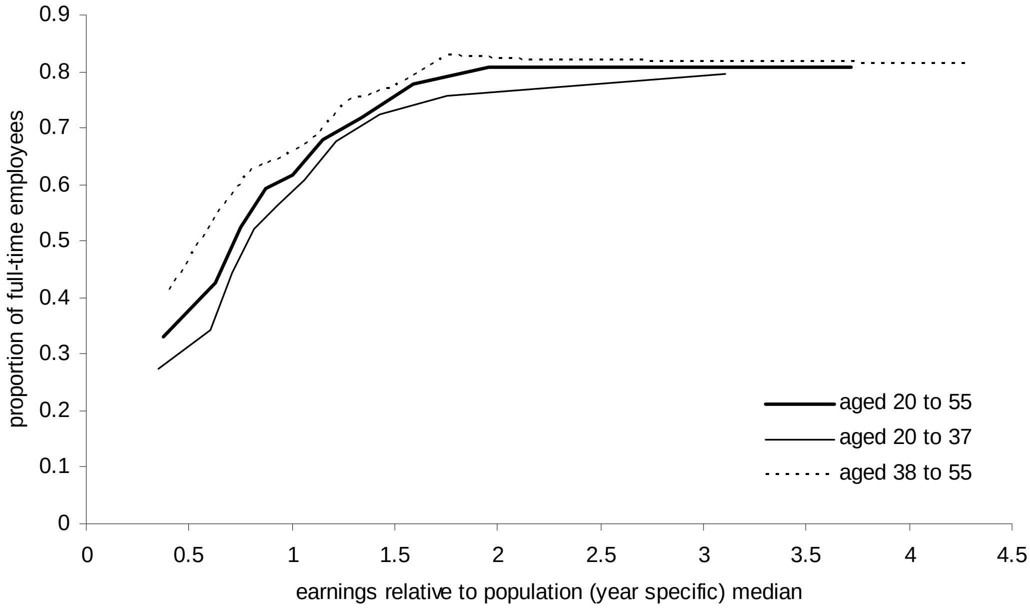

Panel A: eligibility to any employer pension contribution

Panel B: employer contribution where some contribution was made

Source: Panel A: author's calculations on individual level data from the 2006/07 and 2007/08 waves of the FRS

Panel B: author's calculations using data from waves 2005 to 2009 of the Annual Survey of Hours and Earnings

Notes: Earnings deciles defined within survey waves, and averaged across waves

Excludes employees in the public sector and the self-employed

low pay industries as defined by the Low Pay Commission (2010), Appendix 4

Figure 1: Eligibility rates of full-time employees to employer sponsored pensions by age and earnings

equal numbers of males and females at age 16, the gender specific mortality rates reported for this birth cohort by the ONS can be used to project the ratio of men to women through time. This ratio was used to obtain a weighted average of the gender specific mortality rates reported by the ONS for each modelled birth cohort. To avoid imposing unwarranted structure on the parameters, the mortality rates were stored in the form of a transition matrix, comprised of 111 rows (representing ages 20 to 130), and 112 columns (representing years 1951 to 2060, with two additional rows to represent the distant past and future). The transition probabilities used can be obtained from the authors upon request.

Relationship status

The model requires rates of marriage formation and dissolution. At the time of writing, the ONS reports historical data for the number of marriages in England and Wales by age, sex and calendar year at annual intervals between 1851 and 2009. The ONS also makes available for modelling purposes the component factors that underlie its population projections, which describe official estimates for the number of marriages by age and sex at annual intervals between 2008 and 2033. Furthermore, ONS population estimates by age, sex and marital status are available for England and Wales at annual intervals between 1971 and 2033. These statistics permit age and gender specific marital rates to be calculated for England and Wales at annual intervals between 1971 and 2033 inclusive.

Marriage dissolution in the model accounts for both divorce and death of a spouse. The ONS reports age and sex specific divorce rates for England and Wales at annual intervals between 1950 and 2010, which can be extended to 2032 by the component factors of the ONS population projections that are referred to above. Combined with the mortality statistics that are referred to in the preceding subsection, these two series of data provide sufficient information to compile age, sex and year specific marital dissolution rates between 1951 and 2032.

The rates of marriage formation and dissolution that are described above are imperfect for modelling purposes in (at least) three important respects. First, the marriage rates calculated on historical data do not account for marriages that are performed abroad. Secondly, it is well recognised that mortality rates are correlated with marital status, and the required detail to take this into account is not provided by the information that is referred to above. And thirdly, the majority of the statistics that are reported by the ONS focus on legal marital status, and do not extend to include civil partnerships or cohabitation.

The first and second problems identified above were addressed by adjusting marriage rates to age 44, and marital dissolution rates from age 45, to align age, sex, and year specific proportions of the population identified as married in the model to population estimates reported by the ONS (which are structured primarily around Census data). The focus of ONS statistics on legal marriage is problematic for modelling purposes due to the rise of civil partnerships and cohabitation, and the fact that couples who share the same address often engage in some pooling of consumption and income. This pooling of financial resources is recognised by the system of social security in the UK, which treats cohabitating couples in the same way as registered married couples when determining eligibility to most benefits (excluding state pensions and bereavement allowances). We consequently applied a final set of adjustments to account for this issue.

The Family Expenditure Survey (FES) provides detailed micro-data that can be used to determine age and sex specific proportions of the population married between 1971 and 1989, and married or cohabitating between 1990 and 2009. Starting from the rates of marriage and marital dissolution calculated for registered marriages (described above), it is possible to compute implied age, sex, and year specific proportions of the population married on the assumption of zero migration. These proportions of the population married by age, year, and sex were compared against the associated proportions calculated on FES data that allow for cohabitation. Marriage rates were then adjusted to age 44, and marital dissolution rates were adjusted from age 45, to match the implied proportions of the population in a cohabitating relationship to the proportions calculated on FES data. As the associated adjustments were exactly identified (involving the same number of model parameters as fitted moments), a precise fit to the sample moments was obtained.

The gender specific marital and marriage dissolution rates derived via the above procedure were aggregated into a gender neutral series in the same way as described above for mortality rates. Similarly, like mortality rates, the probabilities upon which change in relationship status depend were stored in two transition matrices, one defining the probabilities of marriage for individuals who were single in the preceding year, and one defining the probabilities of marital dissolution. The transition matrix for marriage is comprised of 65 rows (representing ages 20 to 84) and 41 columns (representing years 1971 to 2009 with two additional rows to represent the distant future and distant past). The transition matrix for marital dissolution is comprised of 86 rows (representing ages 16 to 101; all adults are assumed to be single from age 101) and 41 columns.

Number of dependent children

The number of dependent children is modelled as a deterministic function of age, year, and relationship status. This function is stored in the form of a matrix over these three dimensions, with dimensions 59 (representing ages 20 to 78) by 41 (representing years 1971 to 2009 with two additional elements to represent the distant past and future) by 2 (representing singles and couples). The elements of this matrix were set equal to averages reported in the FES.

C Regression Models Used to Input Missing Earnings Data

Four regression equations were estimated on the cross-sectional WAS data, which were used to impute earnings for individuals who were not reported to be working full-time in the sample; separate equations for men and women, and separate equations for those aged under age 50 and those aged 50 years or over. The specifications adopted for this analysis were constrained only by the information reported by the WAS, which includes a high degree of financial detail. After experimenting with various alternatives, regression results for the assumed earnings equations are reported in Table 5.

The parameter values reported in Table 5 indicate that earnings are positively correlated with education, lower for students, and tend to vary positively with health and socio-economic status. Self-reported savers tend to earn more than non-savers, and earnings are positively related to aggregate wealth, home ownership, and mortgage value. Part-time employment tends to imply an earnings penalty of 60% for women, and 70% for men; these compare with the model assumption that part-time work returns 40% of an individual's full-time wage (described in Appendix B.1).

The estimates obtained for rho – the correlation between the residuals of the target and selection equations – are interesting in their own right. These coefficients suggest that censoring tends to be more likely for low income individuals early in life, and more likely for high income individuals later in life, where the effects are not insignificant at the 90% confidence interval for women under age 50 and for men over age 49. Comparing the estimates obtained for sigma with the standard deviations of the associated dependent variables indicate that the regression models selected for analysis help to explain around 30% of the observed variation between individuals.

D Endogenously Identified Wage Parameters

As noted in Section 3.2, two sets of wage parameters were calibrated by matching moments of employment income implied by the structural model against associated moments estimated on time-series data: the drift parameters $m(.)$, and earnings volatility, $\sigma^2_{\omega,n_{i,a}}$. We calibrated the drift parameters to match the model to age, year, and relationship specific geometric means of employment income estimated on data reported by the Family Expenditure Survey between 1971 and 2010. These geometric means, averaged over all years are reported in Figure 2. As the earnings profiles reported in Figure 2 fall precipitously at higher ages, we calibrated the associated earnings parameters to moments calculated for singles to age 60 and for couples to age 64. The large number of parameters involved make reporting here impractical, and associated statistics can therefore be obtained from the authors upon request.

E Simulated vs Sample Moments – Disaggregated Statistics

	women aged 18-49		women aged 50+		men aged 18-49		men aged 50+	
	est	std error	est	std error	est	std error	est	std error
part-time work	-0.9515	0.0330	-0.9631	0.0472	-1.1867	0.0693	-1.4185	0.1095
no qualifications recorded	-0.1523	0.0406	-0.1014	0.0575	-0.1340	0.0294	-0.0702	0.0557
graduate qualifcations	0.3199	0.0288	0.2950	0.0666	0.1804	0.0252	0.1944	0.0631
student	-0.0973	0.0609	0.1083	0.1096	-0.1887	0.1112		
student under age 23	-0.1699	0.1252			-0.0292	0.1525		
self-reported health vgood	0.0490	0.0315	-0.1217	0.0776				
self-reported health good	0.0574	0.0311	-0.1477	0.0712				
self-reported health fair			-0.2029	0.0858				
self-reported health bad			-0.6148	0.3467				
self-reported health vbad			-0.2515	0.2358				
SEC 1 of 3	0.2316	0.0264	0.2922	0.0537	0.1987	0.0247	0.1890	0.0456
SEC 2 of 3	0.1174	0.0336	0.1810	0.0609	0.0603	0.0268	0.0134	0.0675
self-reported saver	0.0504	0.0250	0.1275	0.0481	0.0823	0.0197		
parner not working	0.0083	0.0212			-0.0511	0.0171		
wealth under £10000							-0.1990	0.0739
own accommodation (£)	0.1168	0.0271			0.0910	0.0211		
total net wealth (£)	4.76E-07	8.51E-08	1.14E-07	7.11E-08	8.25E-07	8.72E-08	2.92E-07	8.59E-08
total private pension wealth (£)							1.08E-07	4.47E-08
outstanding mortgage (£)			1.47E-06	3.46E-07			1.72E-06	3.14E-07
constant	9.4671	0.0553	8.7632	0.1775	9.8266	0.0397	9.7207	0.2211
rho	0.0582	0.0303	-0.0477	0.0334	0.0110	0.0383	-0.0960	0.0407
sigma	0.6145	0.0288	0.8033	0.0434	0.5193	0.0216	0.7282	0.0435
lambda	0.0358	0.0189	-0.0383	0.0274	0.0057	0.0200	-0.0699	0.0298
Sample	5916		6084		5389		5300	
Censored observations	1976		4468		1324		3561	
mean dependent variable	9.4115		9.2189		9.8577		9.7363	
std dev of dependent variable	0.8939		1.0902		0.7785		1.0756	

Regression estimates for target equations controlling for sample selection

Regression estimates calculated on Wealth and Assets Survey data, using the "heckman" command in Stata

Table omits age specific dummy variables; SEC = Socio Economic Class

Robust standard errors reported

All statistics are dummy variables, except for the financials indicated by the (£) symbol

Table 5: Regression Estimates for log Earnings, Controlling For Sample Selection

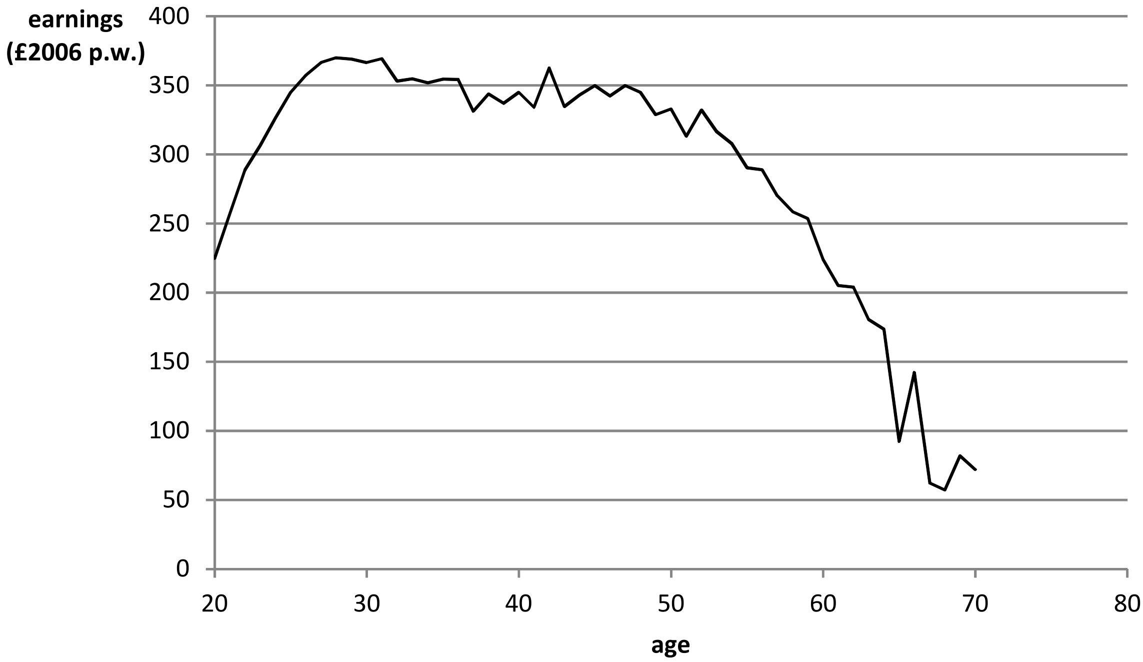

Panel A: Single adults

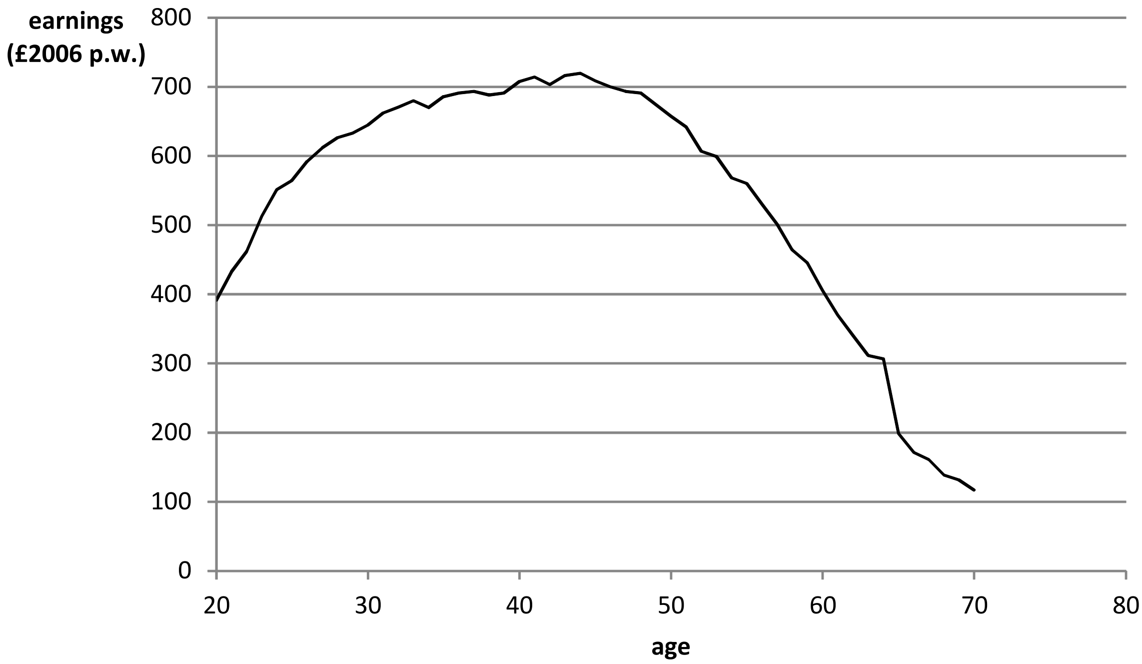

Panel B: Couples

Figure 2: Age profile of geometric mean earnings by relationship status; averaged over annual cross-sections 1971-2010

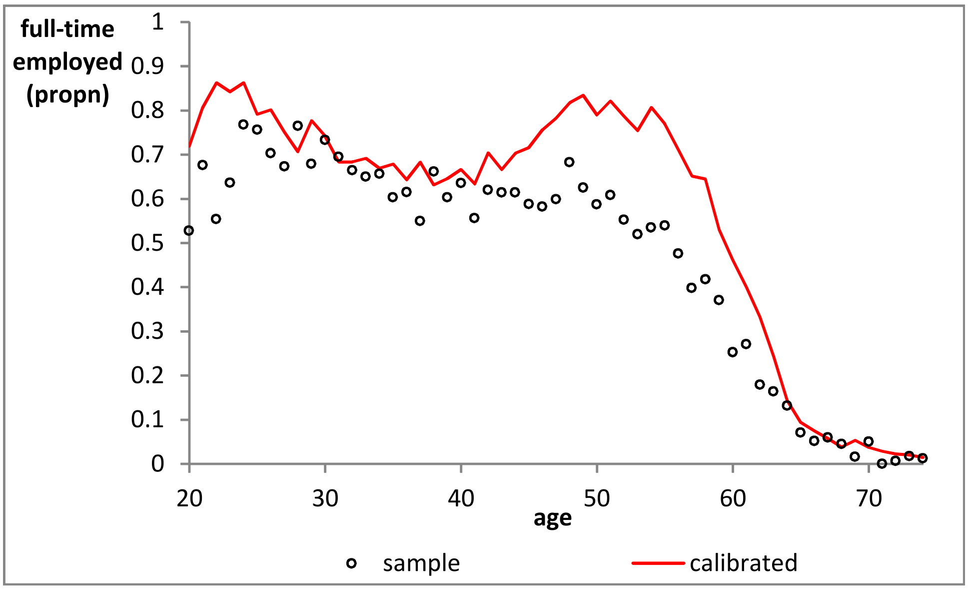

Panel A: rates of full-time employment

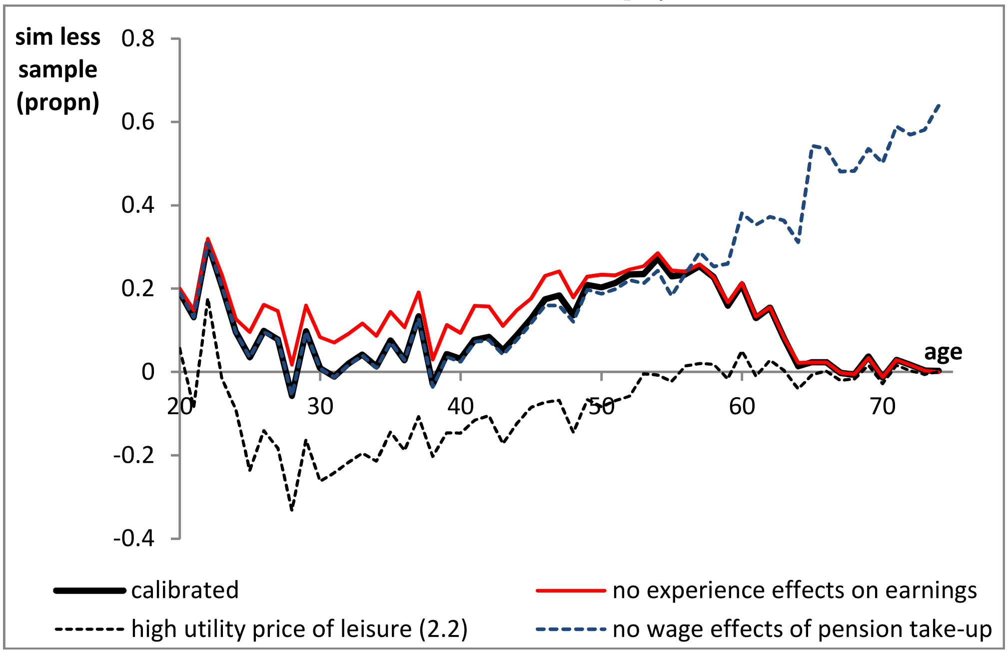

Panel B: simulated moments less sample moments

Figure 3: Rates of full-time employment of adults in a cohabitating relationship

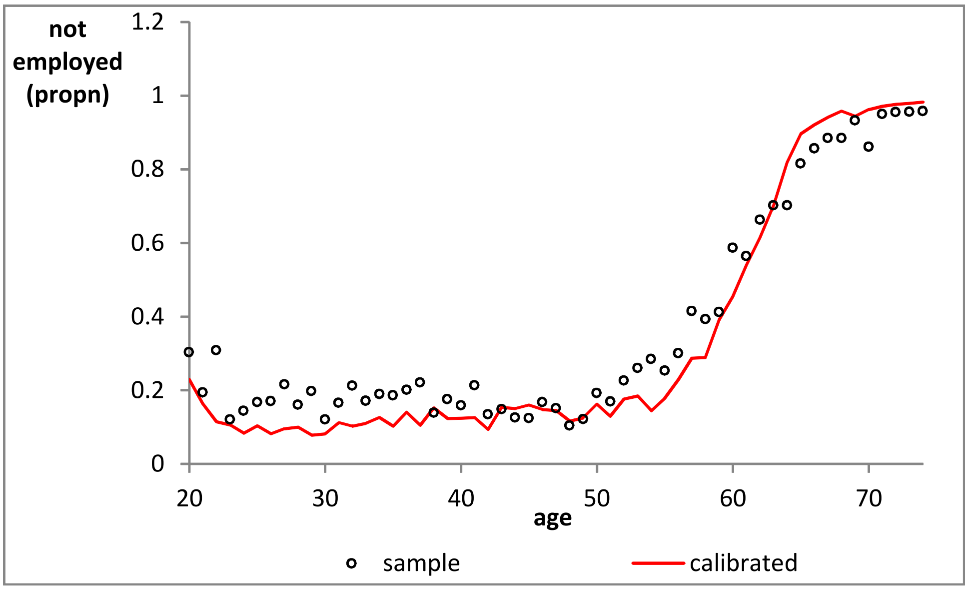

Panel A: rates of non-employment

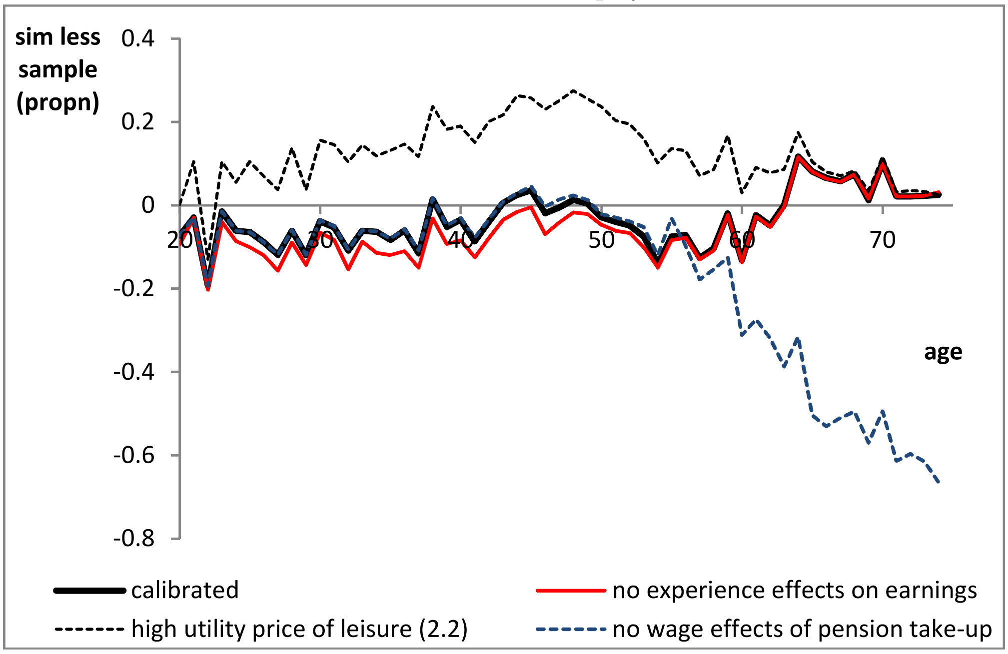

Panel B: simulated moments less sample moments

Figure 4: Rates of non-employment of adults in a cohabitating relationship

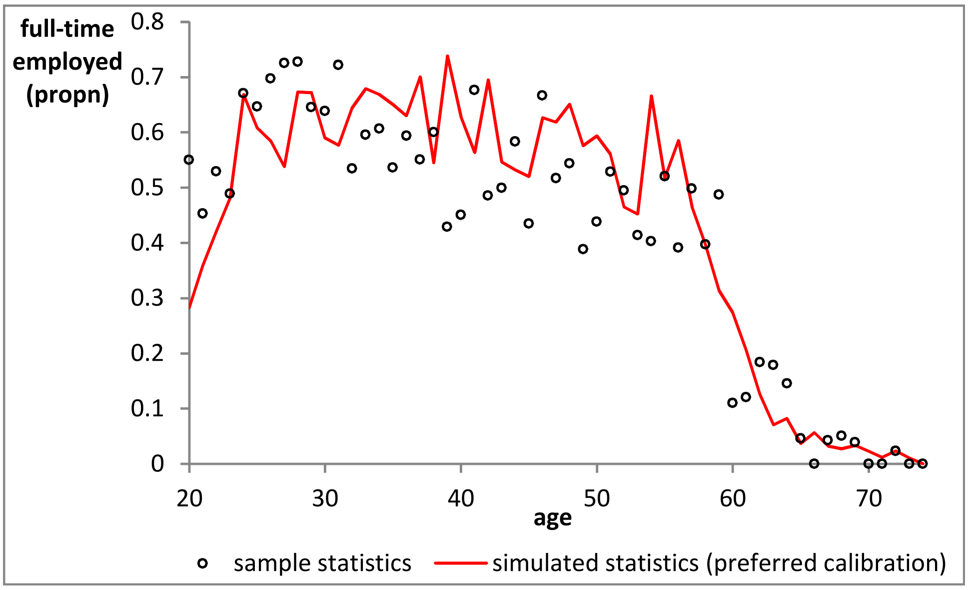

Panel A: rates of full-time employment

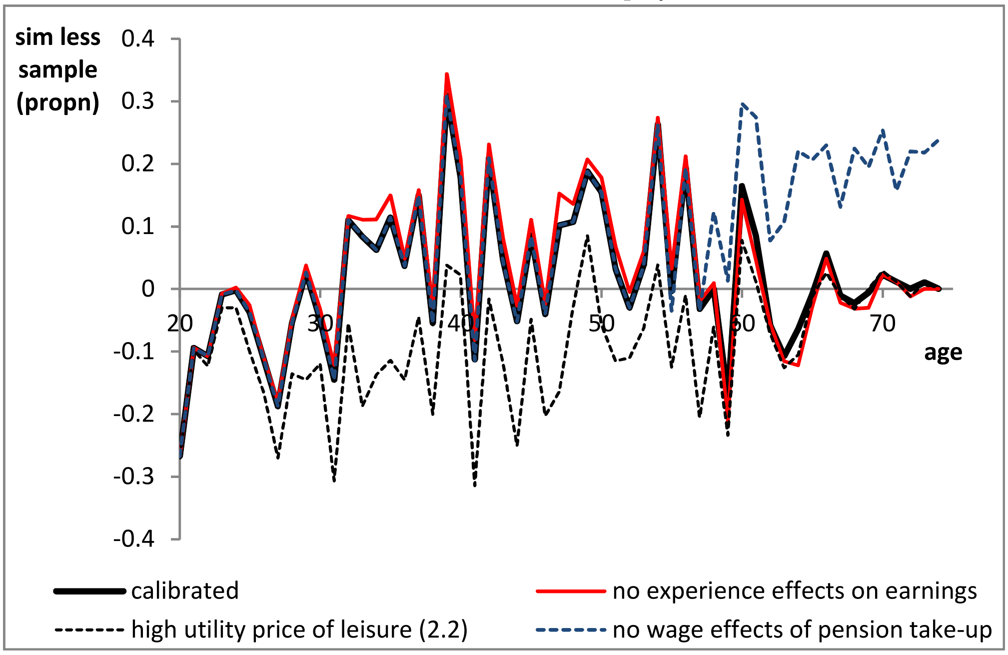

Panel B: simulated moments less sample moments

Figure 5: Rates of full-time employment of single adults

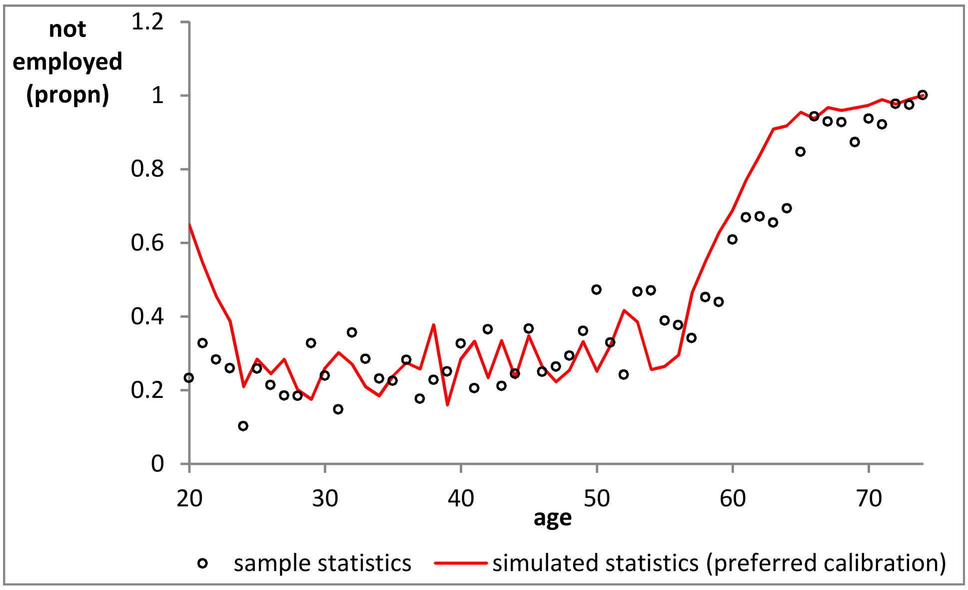

Panel A: rates of non-employment

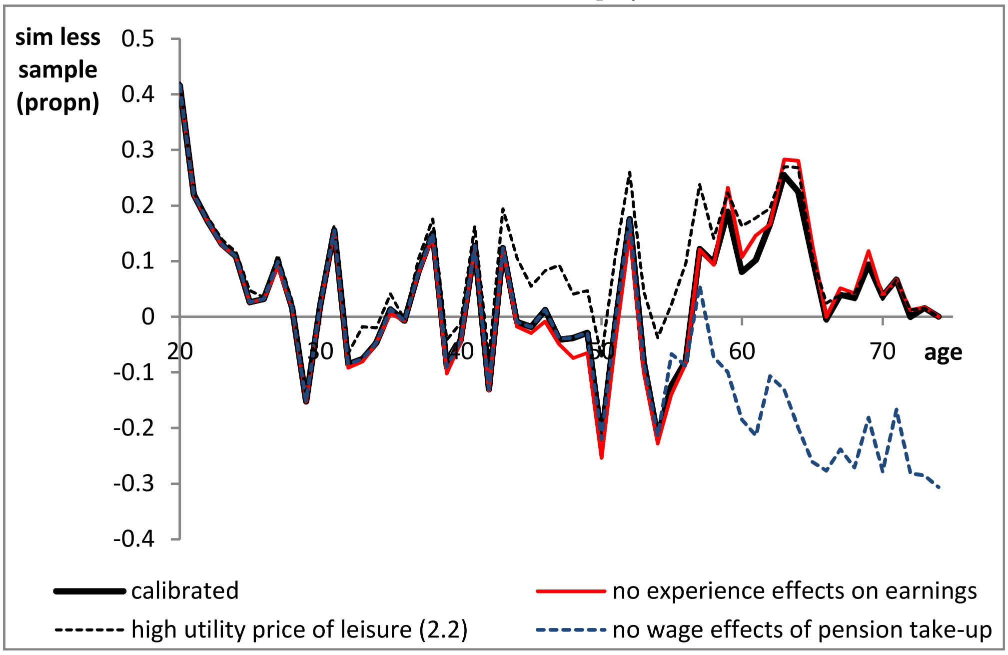

Panel B: simulated moments less sample moments

Figure 6: Rates of non-employment of single adults

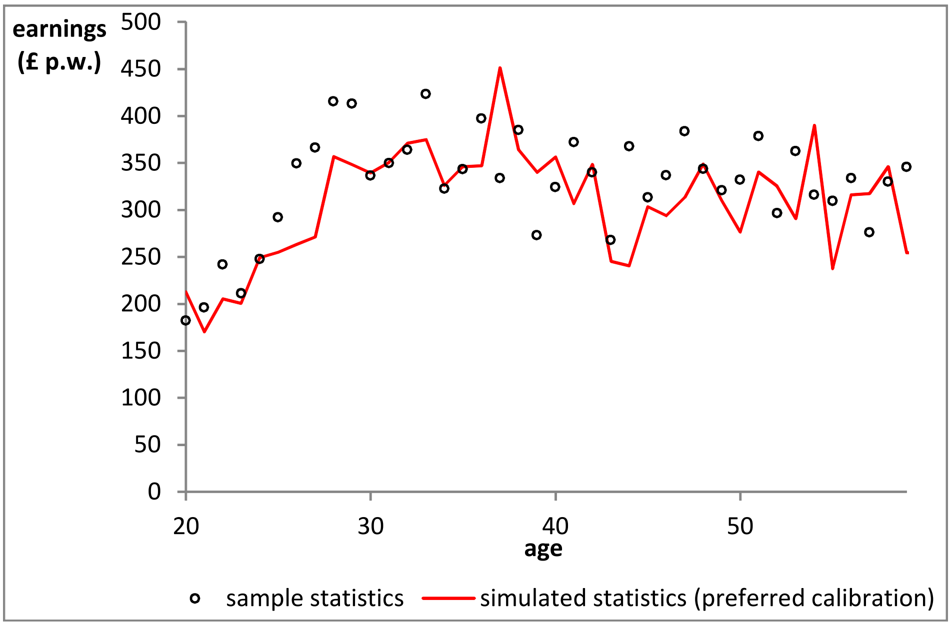

Panel A: geometric mean of employment income

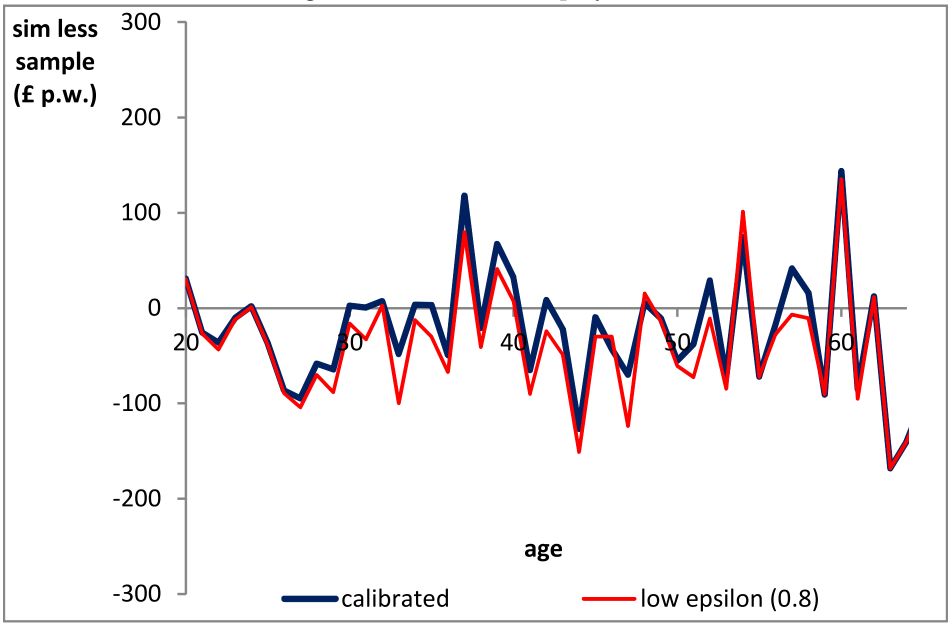

Panel B: simulated moments less sample moments

Figure 7: The geometric mean of employment income of single adults

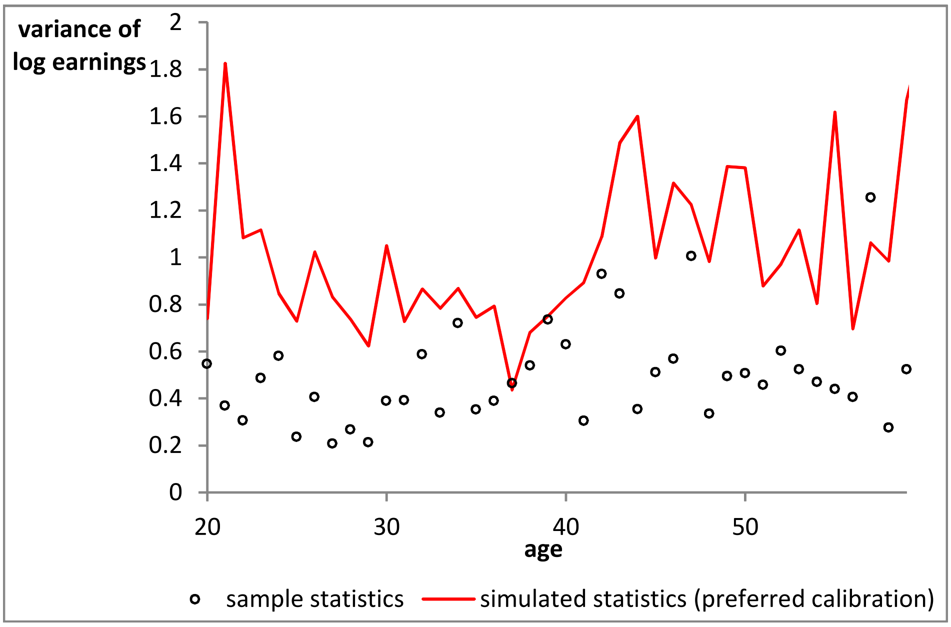

Figure 8: The variance of log employment income of single adults

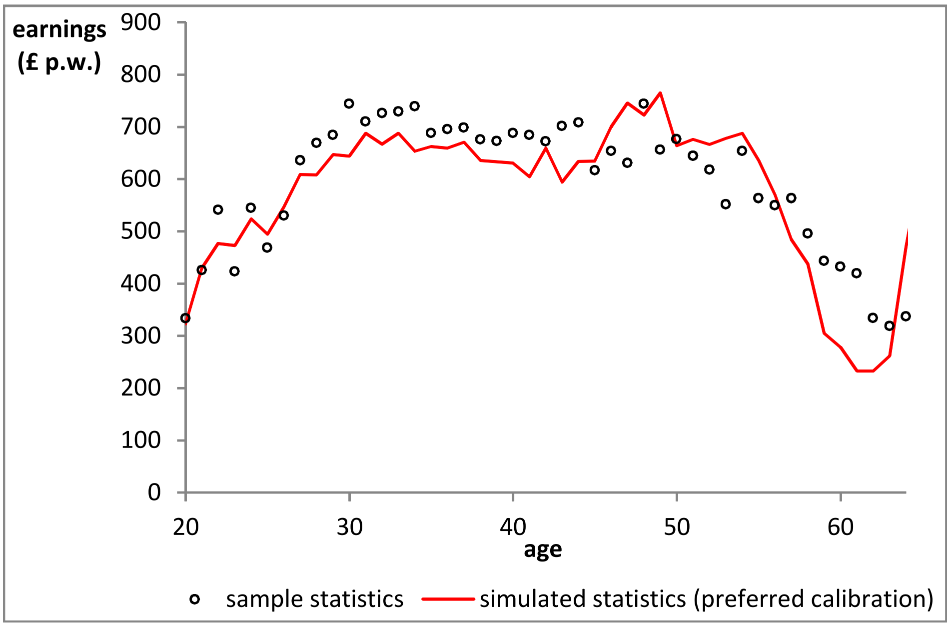

Panel A: geometric mean of employment income

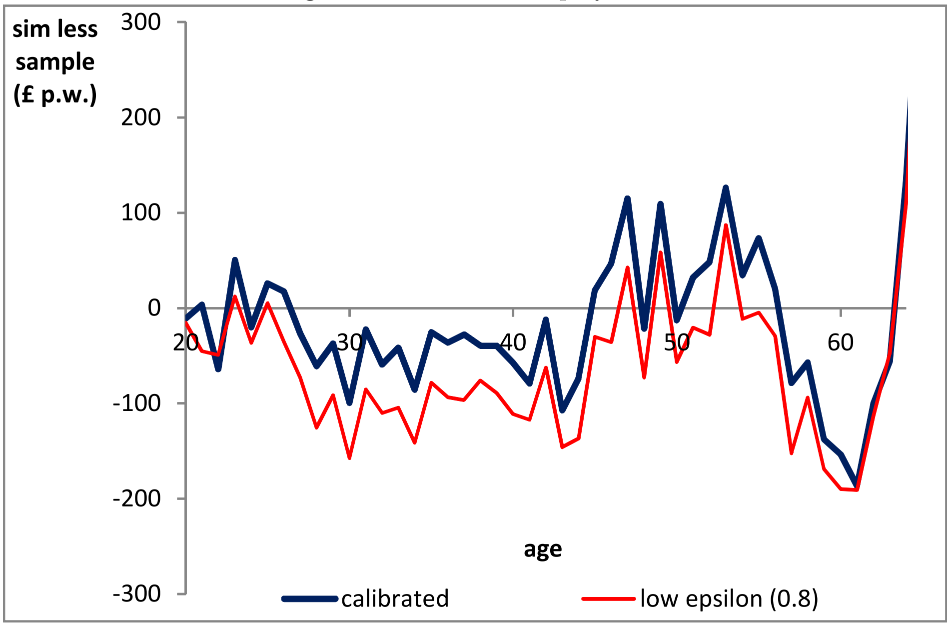

Panel B: simulated moments less sample moments

Figure 9: The geometric mean of employment income for cohabitating couples

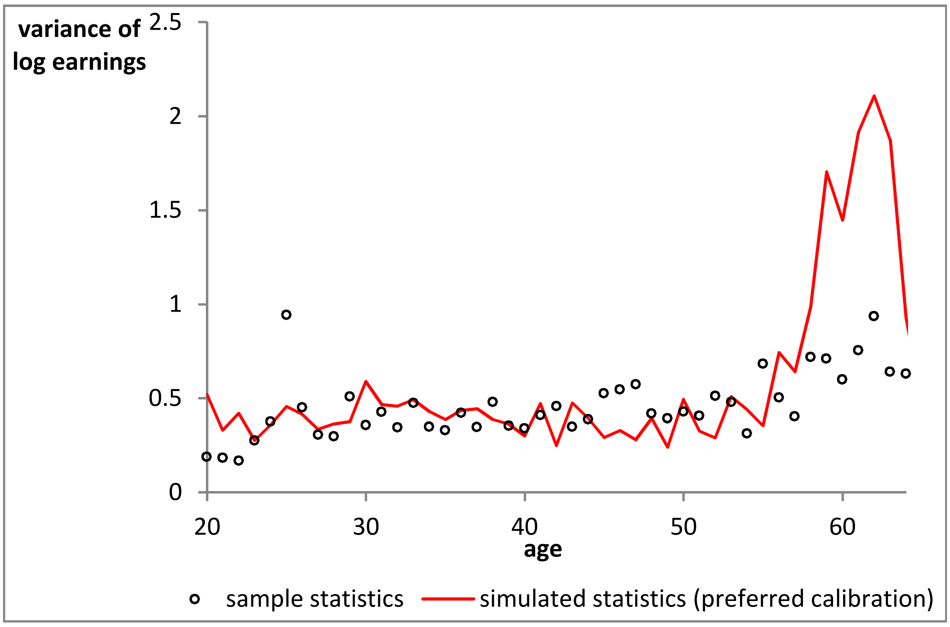

Figure 10: The variance of log employment income for cohabitating couples

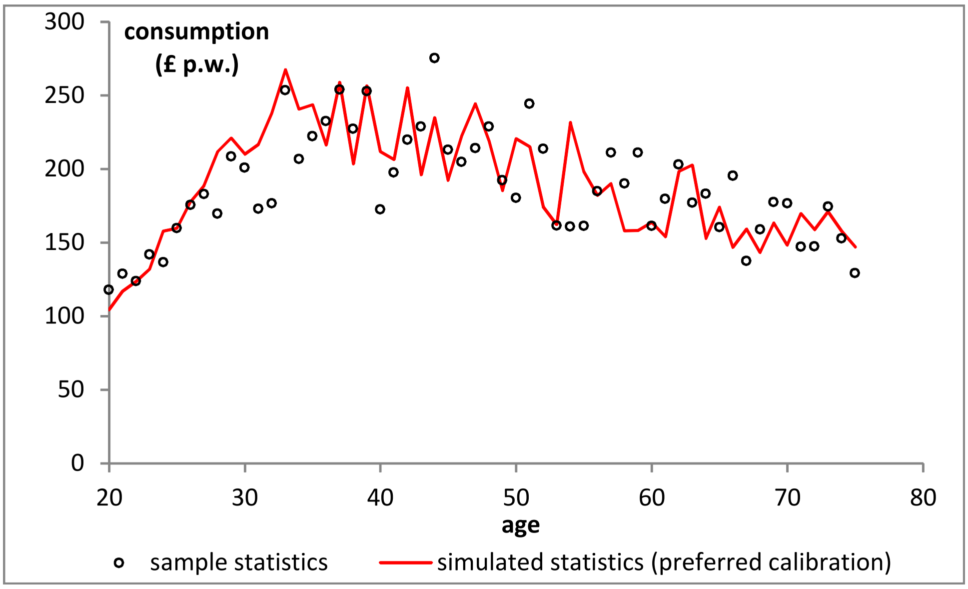

Panel A: geometric mean of employment income

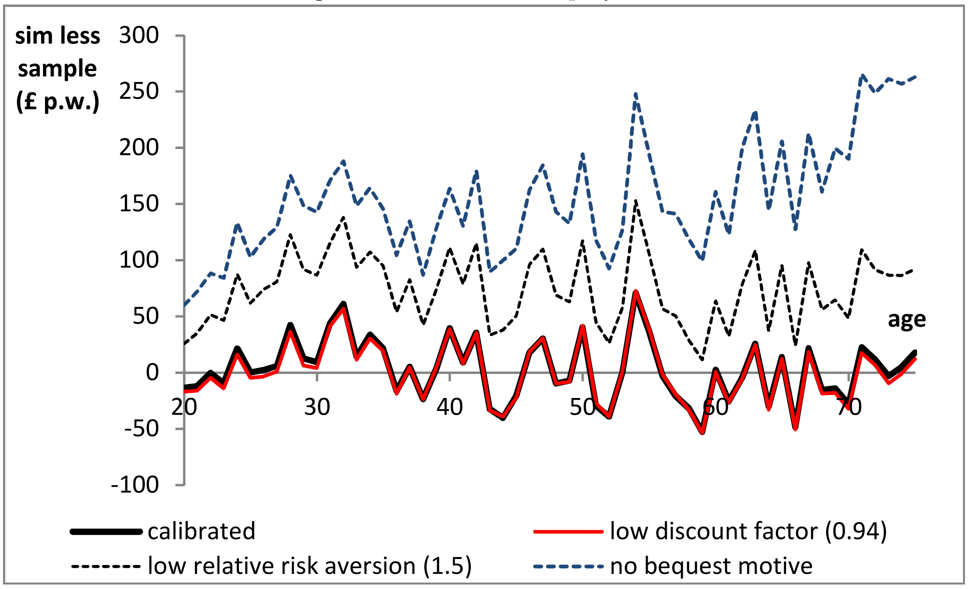

Panel B: simulated moments less sample moments

Figure 11: The geometric mean of consumption of single adults

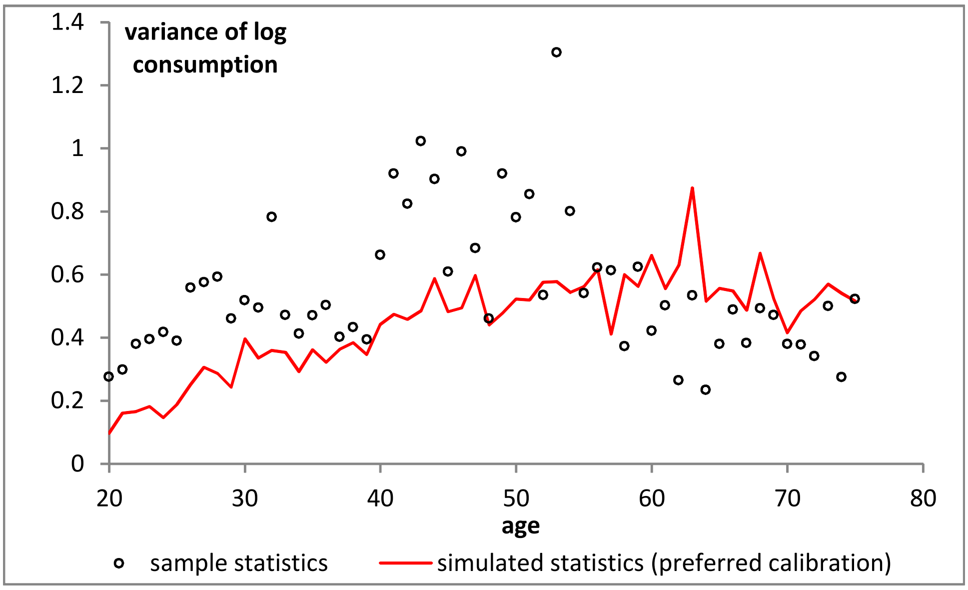

Figure 12: The variance of log consumption of single adults

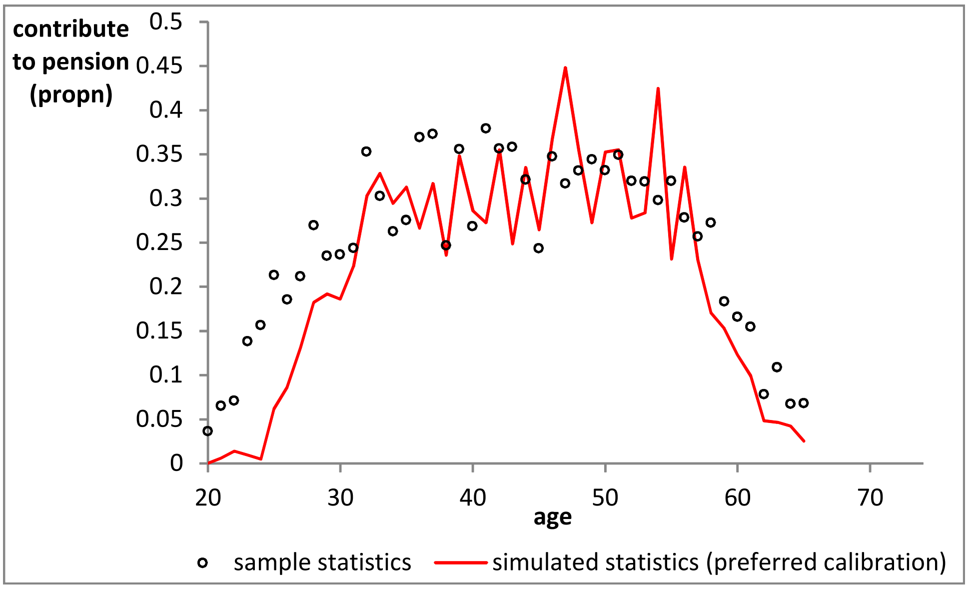

Panel A: rates of pension scheme participation

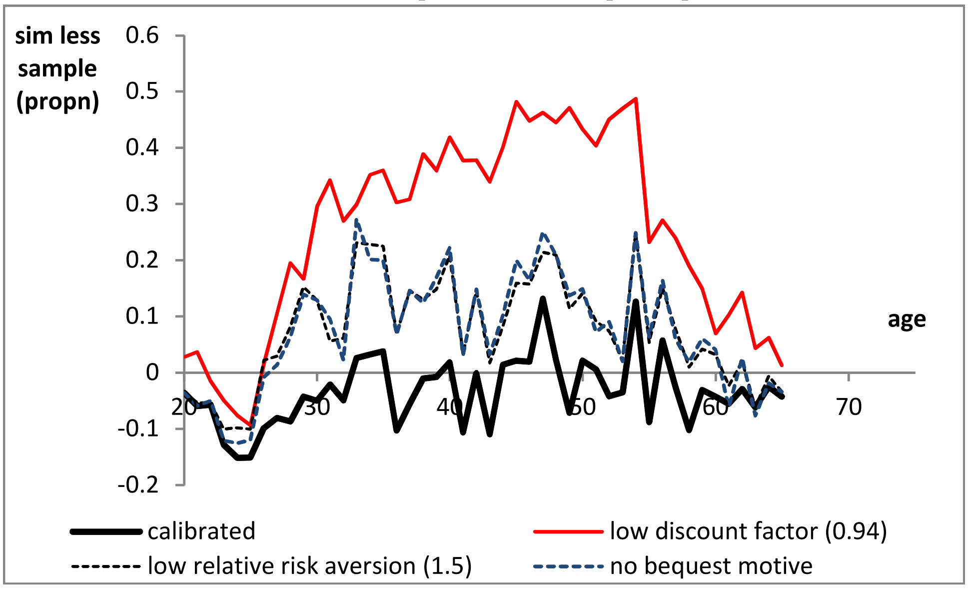

Panel B: simulated moments less sample moments

Figure 13: Rates of pension scheme participation of single adults

Panel A: geometric mean of employment income

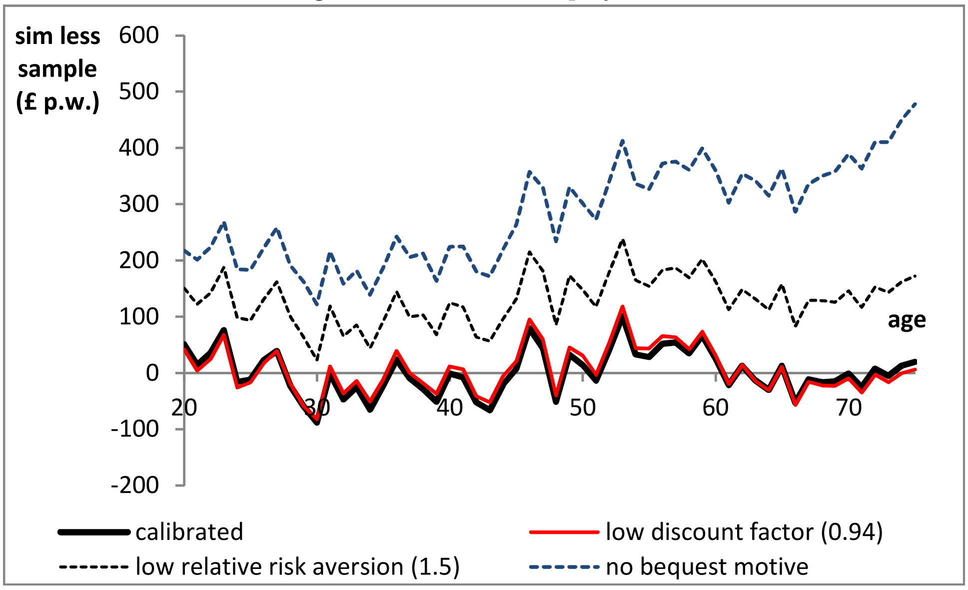

Panel B: simulated moments less sample moments

Figure 14: The geometric mean of consumption for cohabitating couples

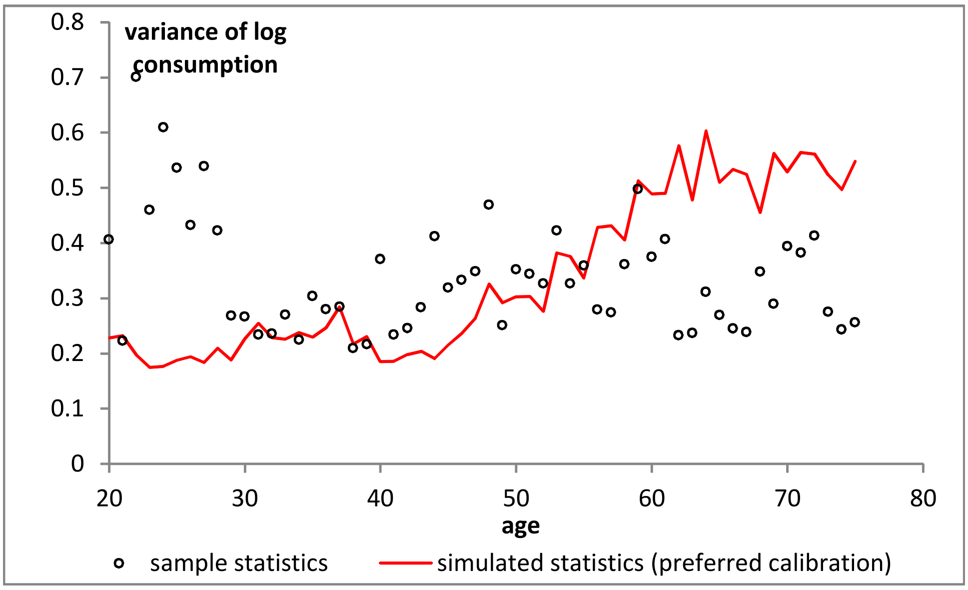

Figure 15: The variance of log consumption for cohabitating couples

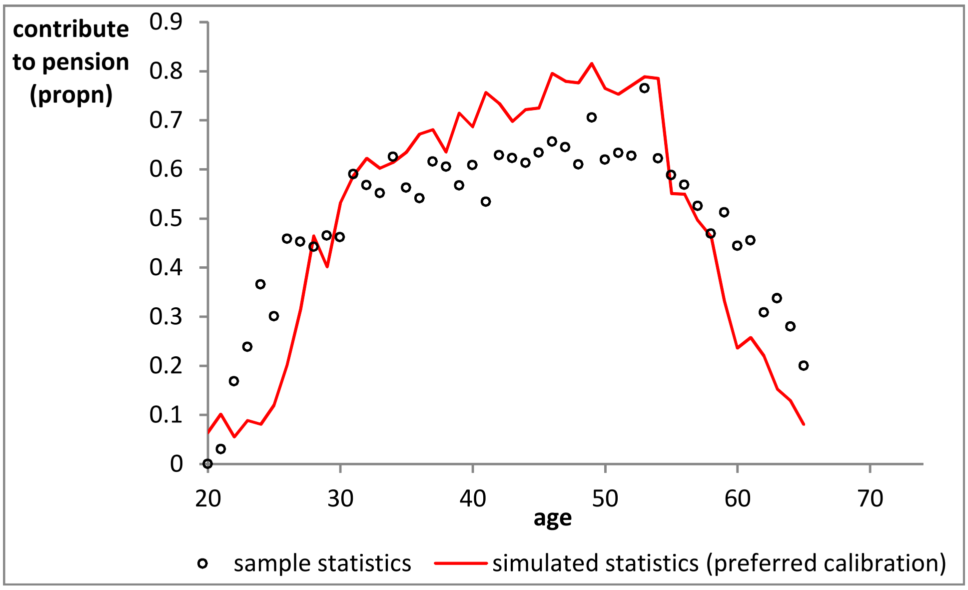

Panel A: rates of pension scheme participation

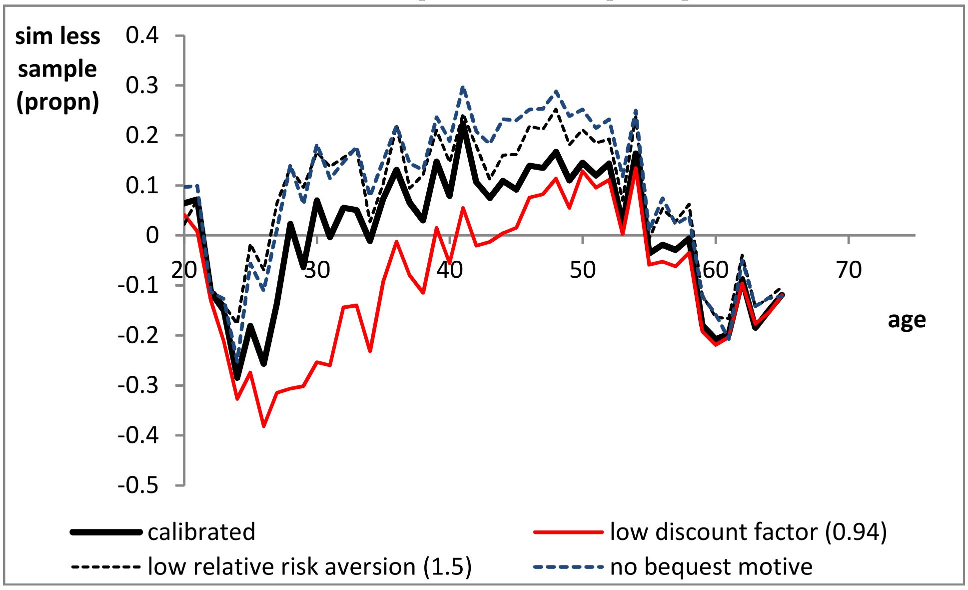

Panel B: simulated moments less sample moments

Figure 16: Rates of pension scheme participation of adults in cohabitating relationships