Economic Modelling 28 (2011) 2054–2070

Contents lists available at ScienceDirect

Economic Modelling

journal homepage: www.elsevier.com/locate/ecmod

A structural dynamic microsimulation model of household savings and labour supply

Justin van de Ven

NIESR, 2 Dean Trench Street, London, SW1P 3HE, UK

ARTICLE INFO

Article history:
Accepted 18 April 2011

JEL classification:
C51
C61
C63
H31

Keywords:
Dynamic programming
Savings
Labour supply
Taxes and benefits

ABSTRACT

A structural model of the household is described that represents current best-practice in the analysis of savings and labour supply responses to the policy environment. Care has been taken in specifying the model so that it represents an appropriate basis for the analysis of incentive effects to policy change, and for exploring the empirical support for alternative structural assumptions. Matching the model to survey data for the UK reveals some interesting puzzles in relation to the timing of retirement.

1. Introduction

Structural models of behaviour are the clock-work to which economic theorising is devoted, and upon which analyses of our changing environment depend. Current best practice in the economic analysis of intertemporal decision making is conducted within the life-cycle framework of Modigliani and Brumberg (1955) and Friedman (1957). Despite more than half a century of intense analysis, however, the life-cycle framework has had, and continues to have, a very limited bearing upon practical issues of policy design and reform. One of the reasons for this disappointing yield is the high computational burden that is generally involved when using the life-cycle model to identify decisions in contexts where future circumstances are uncertain. It is reasonable to suppose, therefore, that this aspect of economic analysis should have benefitted immensely from the rapid advance of computing technology, which has persisted for as long as the life-cycle framework has existed. In this paper I describe a model that represents the current state of the art in the economic analysis of savings and labour supply decisions, thereby revealing how much progress has been made, and – just as important – how far we still have to go before it can reasonably be said that this field of study has truly “come of age”.

Policy makers are faced with a multitude of difficult problems when considering the relative merits of policy alternatives. One of the most important of these is the evaluation of a policy counterfactual's incentive effects, which are often highly opaque and context dependent. Consider, for example, a reduction in the rate at which a means-tested retirement benefit is withdrawn in respect of private income. The most commonly expressed view is that this type of policy change reduces effective tax rates on saving, thereby encouraging an increase in private retirement provisions. But, as is well understood in the economic literature, there is no reason to suppose that this ‘intuitive’ view will hold in practice. This is because, for net savers who would have been eligible to receive the means-tested benefit in the absence of the policy change, the substitution and wealth effects of reducing the withdrawal rate act in opposite directions. Furthermore, for those who would only be made eligible to receive the means-tested retirement benefit following the policy change (as the benefit is extended up the income distribution), the income and substitution effects work in the same direction, to *reduce* incentives to save. And this stylised discussion ignores any interactions that might exist between the means-tested benefit and the wider tax and benefits system, which are often sufficiently complex in advanced economies to add substantially to the difficulty of inferring incentive effects.

In response to these complexities, microsimulation models are now one of the tools commonly applied in the analysis of the implications of policy reform. The feature that distinguishes microsimulation models from their macro based counterparts is that each agent from a population is individually represented. As such, microsimulation models are particularly useful for policy analyses where the effects depend upon individual specific circumstances (as in the preceding example), or where the distributional implications are a focus of interest.

Two types of microsimulation model can be distinguished. Static microsimulation models, as their name suggests, determine the

E-mail address: jvandeven@niesr.ac.uk.

0264-9993/$ – see front matter © 2011 Elsevier B.V. All rights reserved.
doi:10.1016/j.econmod.2011.04.009

impact of counterfactual conditions upon a population of agents at a point in time. In contrast, dynamic microsimulation models age each agent in response to stochastic variation and an accumulated history, eventually generating panel data over the entire life course. Most dynamic microsimulation models are consequently designed to consider the intertemporal and long term effects of counterfactual conditions, rather than the impact effects with which static models are concerned.

The vast majority of microsimulation models that are currently in use are based either upon simple assumptions regarding behavioural responses (the simplest being that there are none), or on regression equations that are selected for their empirical appeal more than their structural content. As the behavioural assumptions upon which these models are based are likely to be sensitive to changes in the policy environment, they are usually ill-suited to an evaluation of the incentive effects of policy reform. Addressing this short-coming is now a vibrant issue of academic research, and is the fundamental motivation for the modelling approach discussed in this paper.

The model that is described here is best referred to as a structural dynamic microsimulation model of the household. It departs from 'classical' dynamic microsimulation models (e.g. Curry (1996), Caldwell (1997)) because it is founded upon "deep parameters" that are defined to be structurally stable in context of environmental change. This is what motivates the emphasis placed on the 'structural' nature of the model, and makes it robust to the Lucas critique when exploring the implications of changes to the decision making environment.

The theoretical basis of the model is specified in response to two sets of empirical observations. First, are the empirical regularities that motivated formulation of the life-cycle framework in the 1950s as an alternative to the "fundamental psychological laws" of Keynes (1936). And second, are observations reported in the extensive literature that has sought to test specific assumptions within the life-cycle framework. Before discussing these, note that an explicit distinction is drawn here between the life-cycle 'framework' and a life-cycle 'model'. Following Browning and Lusardi (1996), I refer throughout to the life-cycle framework as the hypothesis that decisions are made to maximise expected lifetime utility, subject to the decision maker's circumstances, constraints, and expectations about the future. A life-cycle model obtains when additional assumptions are added to the framework that allow behaviour to be explored.

In a recent review, Attanasio and Weber (2010) cite the following stylised empirical facts as providing the underlying motivation for the life-cycle framework: the marginal propensity to consume out of disposable income is lower in the short-run than the long run; there is an inverse relationship between saving rates and population sub-group average income after controlling for income level; and there is a positive correlation between saving rates and income changes. These stylised observations, which continue to hold in contemporary survey data, can all be coherently explained within the life-cycle framework (e.g. Attanasio and Weber (2010)).

One of the central issues addressed by the literature that has tested alternative assumptions within the life-cycle framework is whether consumption growth is orthogonal to predictable changes in income, as is implied by a highly stylised life-cycle model. This implication has been rejected over a wide range of dimensions considered in the literature: age specific measures of consumption and income are both hump shaped, with consumption appearing to track income across the life-course over a wide range of alternative population subgroups (e.g. Carroll and Summers (1991)); consumption growth appears to respond significantly to instrumented income growth (e.g. Campbell and Mankiw (1989, 1991), and Flavin (1981)); and consumption is significantly influenced by predictable changes in both behaviour (e.g. retirement, Hamermesh (1984), Banks et al. (1998), and Berneim et al. (2001)) and the policy environment. These deviations between the data and the stylised life-cycle model first considered for analysis are referred to in the literature as 'excess sensitivity'.

Furthermore, observed aggregate consumption growth is more smooth than aggregate income growth, despite the positive serial correlation that is exhibited by income growth (e.g. Campbell and Deaton (1989)). This is a discrepancy with the implications of common life-cycle model specifications that is referred to in the literature as 'excess smoothness'. The literature has also documented cases where individuals appear to exhibit a preference for commitment mechanisms that are absent from most life-cycle models (e.g. Thaler and Benartzi (2004)), where decisions show a significant bias in favour of the choice involving the least effort (where such a distinction exists; e.g. Choi et al. (2002), Madrian and Shea (2001), and Beshears et al. (2008) on the influence of default options for 401 (k) pension plans in the US), and where investment decisions depart radically from model implications (e.g. Gross and Souleles (2002) on coincident holding of assets and debt, and Mehra and Prescott (1985) on the equity premium puzzle).

Excess sensitivity violations of life-cycle models have been resolved in the contemporary literature by relaxing a number of specific assumptions made in the models that predominated in the early empirical literature. The hump shaped profile of consumption is now commonly understood to be the product of population demographics; changing household needs due to marriage and child rearing during the working lifetime, and uncertainty regarding the time of death late in life (see, e.g., Tobin (1967), Attanasio and Weber (1995), and Browning and Ejrnæs (2002)). Precautionary saving in response to income uncertainty has been identified as an important factor underlying the timing of the peak in household consumption, which occurs later in life than the peak in income (e.g. Carroll (1992), Browning and Ejrnæs (2002), and Gourinchas and Parker (2002)). Heckman (1974) recognised that we should not expect consumption to be independent of retirement decisions if consumption and leisure are not additively separable in utility, and Smith (2006) identifies the role of negative shocks to life-cycle income associated with involuntary retirement. Furthermore, a number of authors have reported evidence that liquidity constraints are important in understanding the responsiveness of consumption growth to income growth described by survey data more generally (e.g. Zeldes (1989), and Agarwal et al. (2007)). The model discussed in this paper is designed to reflect all of these considerations.

The desire for commitment mechanisms that has been identified in some contexts is commonly attributed in the literature to time-inconsistency of preferences, with the majority of attention paid to quasi-hyperbolic discounting in the aggregation of intertemporal utility (e.g. Laibson et al. (2007) and van de Ven (2010) for empirical investigations). Less substantive progress has, however, been made in relation to the other deviations between common life-cycle models and the statistical record that are set out above, with the most promising candidates for model amendments being the introduction of information asymmetries, decision costs, and habit formation, which are all closely related concepts. A common explanation for excess smoothness, for example, is that intertemporal consumption is influenced by prevailing habits (e.g. Fuhrer (2000), Sommer (2001)).[1] Alternatively, Ludvigson and Michaelides (2001) emphasise the potential of imperfect information about the state of the macroeconomy to explain excess smoothness within the life-cycle framework. Where biases are identified in favour of decision alternatives that impose the minimum effort on the decision maker, then decision costs represent an intuitive explanation; most of the related literature in this regard has focussed on stock accumulation at the firm level and the demand for factors of production (see Khan and Thomas (2008), for a review). Similarly, the role of decision costs and habit formation in explaining apparently errant investment decisions has been explored at length (see e.g. O'Donoghue and Rabin (1999) on decision

[1] There is, however, weak evidence of habits in panel microdata; see Dynan (2000).

costs and procrastination, and Campbell and Cochrane (1999) on the role of habits in explaining the equity premium puzzle). With regard to these observations, the current model is designed to capture quasi-hyperbolic discounting, and the addition of decision costs/habit formation remains an issue for further research.

A singular complication in formulating a model that captures the various considerations that are discussed above is that closed form solutions in general do not exist when uncertainty of the decision making environment is taken into account. Where the analytical focus is on the identification of certain preference parameters, or empirical testing of selected model assumptions, then the non-existence of closed form solutions can be addressed by specifying the empirical analysis in terms of associated Euler conditions. But such conditions do not, in general, permit behaviour in levels – such as a consumption function – to be retrieved. Indeed, this is the most likely candidate to explain why the life-cycle framework has had such a limited bearing upon the practical issues of policy design and reform. To bring a life-cycle model that captures the considerations discussed above to bear upon practical policy issues, it is usually necessary to solve the decision making problem using numerical dynamic programming methods. That is the approach that is adopted for the model described here, with the dual objective to obtain an analytical framework that can help to better inform policy makers, and provide a suitable basis to test empirically alternative behavioural assumptions.

A crucial stage in setting out a model for practical policy analysis is defining parameters to matching the model against survey data. This process, is commonly undertaken either manually via a trial and error process (commonly referred to as calibration), or by an automated econometric procedure.[2] Here, I report results of the former approach. The advantage of this approach is that the process of calibration throws up a number of puzzles in relation to the model's ability to reflect observations drawn from survey data; puzzles that can be easily missed, or difficult to pin down if the parameter adjustment is undertaken by an automated black-box. I show that the unobserved model parameters are tied down tightly by the data considered for calibration. Although the calibration does a fine job of capturing much of the detail that is described by the underlying survey data, the preferred parameter combination fails to capture some aspects of observed behaviour. This is particularly true around the discrete change in tax and benefits policy that occurs at state pension age, which results in an abrupt change in behaviour generated by the simulation model that is not evident in survey data. This influences the calibrated parameters in a way that is comparable to a form of omitted variable bias.

The paper is divided into seven sections. First, a full description of the characteristics that are reflected by the model, and the behavioural framework upon which it is based, are provided in Section 2. Model solution procedures are described in Section 3, and the calibration strategy used to adjust model parameters to reflect survey data is described in Section 4. Section 5 reports estimates for observable model parameters, and Section 6 discusses the fit obtained by the model calibration. A summary and directions for further research are provided in the conclusion.

2. Detailed description of the model

The decision unit in the model is the household, defined as a single adult or partner couple and their dependant children. The model divides the life course into annual increments, and can be used to consider household decisions regarding consumption, labour supply, the portfolio allocation of liquid wealth between safe and risky assets, and private pension contributions. These decisions are made to maximise expected lifetime utility, given a household's prevailing circumstances, their preferences, and beliefs regarding the future. A household's circumstances are described by their age, number of adults, number of children, wage rate, liquid wealth, private pension rights, occupational pension rights, and time of death. Importantly, the belief structure is rational in the sense that expectations are consistent with the intertemporal decision making environment.

Incorporating an appreciation of uncertainty into individual expectations increases the complexity of the utility maximisation problem. Of the eight characteristics that define the circumstances of a household, seven can be considered stochastic (relationship status, number of children, private pension rights, occupational pension rights, wage rates, liquid wealth, and time of death), and only one is forced to be deterministic (age).

In the terminology of the dynamic programming literature, consumption, labour supply, private pension contributions, and the portfolio allocation of liquid wealth are control variables, that are selected to maximise the value function described by a time separable utility function, subject to eight state variables, seven of which can be stochastic. This section begins by defining the assumed preference relation, before describing the wealth constraint, the simulation of pensions, and the processes assumed for the evolution of income and household demographics.

2.1. The utility function

Expected lifetime utility of household i at age t is described by the time separable iso-elastic function:

$$U_{i,t} = \frac{1}{1-1/\gamma}\Bigg\{u\left(\frac{c_{i,t}}{\theta_{i,t}}, l_{i,t}\right)^{1-1/\gamma} + E_t\Bigg[\beta_1\delta\left(\phi_{1,t}u\left(\frac{c_{i,t+1}}{\theta_{i,t+1}}, l_{i,t+1}\right)^{1-1/\gamma} + \left(1-\phi_{1,t}\right)\left(\zeta_a + \zeta_b w^{+}_{i,t+1}\right)^{1-1/\gamma}\right) + \beta_1\beta_2\sum_{j=t+2}^{T}\delta^{j-t}\left(\phi_{j-t,t}u\left(\frac{c_{i,j}}{\theta_{i,j}}, l_{i,j}\right)^{1-1/\gamma} + \left(1-\phi_{j-t,t}\right)\left(\zeta_a + \zeta_b w^{+}_{i,j}\right)^{1-1/\gamma}\right)\Bigg\}\Bigg] \quad (1)$$

where $\gamma > 0$ is the iso-elastic parameter; E_t is the expectations operator; T is the maximum potential age; β_1, β_2, and δ are discount factors (assumed to be the same for all households); $\phi_{j-t,\ t}$ is the probability of living to age j, given survival to age t; $c_{i,t} \in R^+$ is discretionary composite consumption; $l_{i,t} \in [0, 1]$ is the proportion of household time spent in leisure; $\theta_{i,t} \in R^+$ is adult equivalent size based on the "revised" or "modified" OECD scale; the parameters ζ_a and ζ_b reflect the "warm-glow" model of bequests; and $w^{+}_{i,t} \in R^+$ is net liquid wealth when this is positive and zero otherwise.

The labour supply decision (if it is included in the model) is made between discrete alternatives, which reflects the view that this provides a closer approximation to reality than if it is defined as a continuous decision variable for given wage rates (e.g. Fagan (2003)). When adults are modelled explicitly, then households with one adult can choose from up to three labour options; full-time ($l^{FT}_{i,t}$), part-time ($l^{PT}_{i,t}$), and not employed ($l_{i,t} = 1$). Similarly, couples can choose from up to five labour options; both full-time employed ($l^{2FT}_{i,t}$), one full-time and one part-time employed ($l^{FtPt}_{i,t}$), one full-time and the other not employed ($l^{FtNe}_{i,t}$), one part-time and the other not employed ($l^{PtNe}_{i,t}$), and both not employed ($l_{i,t} = 1$); the option for both adults of a couple to be part-time employed is omitted because it is observed only very rarely in survey data. When adults are not modelled explicitly, then labour supply is restricted to one of two options: employed or not employed.

To the extent that the focus on discrete labour options limits employment decisions relative to the practical reality, it will dampen the responsiveness of labour supply behaviour implied by the

[2] A range of alternative methods is also available — see, for example, Lee and Ingram (1991), Stern (1997), Gourieroux et al. (1993), and Gallant and Tauchen (1996). All of these methods are consistent under similar conditions.

simulation model, and dampen variation in employment incomes. The former of these effects implies that the parametrisation of the model may require a labour elasticity that overstates the practical reality, while the latter suggests that excessive variation in labour incomes may be required to reflect the wage dispersion described by survey data.

The modified OECD scale assigns a value of 1.0 to the household reference person, 0.5 to each additional adult member and 0.3 to each child, and is currently the standard scale for adjusting before housing costs incomes in European Union countries.

The model incorporates an allowance for behavioural myopia, through its assumption of quasi-hyperbolic preferences, which are interesting because they are time inconsistent, giving rise to the potential for "conflict between the preferences of different intertemporal selves" (Diamond and Köszegi (2003), p., 1840). The current version of the model focuses exclusively on rational expectations, and consequently does not permit consideration of decisions by so-called "naïve" consumers, who are unaware of their self-control problems in the context of quasi-hyperbolic discounting. The model assumes that all discount parameters are the same for all individuals, and time invariant. This is in contrast to the approach that is adopted by Gustman and Steinmeier (2005), who allow variation in the rate of time preference to be an important factor in reflecting heterogeneity in household retirement behaviour. I have chosen not to do this to ensure that heterogeneity of household behaviour generated by the model is driven by heterogeneity of observable household characteristics.

The warm-glow model of bequests simplifies the associated analytical problem, relative to alternatives that have been considered in the literature.[3] Including a bequest motive in the model raises the natural counter-party question of who receives the legacies that are left. The most accurate approximation to reality would involve including the possibility that households receive a bequest at any age, and then to growth adjust the value of bequests received to the value of bequests made. This would add to the uncertainty associated with the decision problem, and so is omitted from the current version of the model. Rather, it is assumed that households leave their legacies to the state (potentially in the form of a 100% inheritance tax), which is a common simplifying assumption.

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

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

where $\varepsilon>0$ is the (period specific) elasticity of substitution between equivalised consumption ($c_{i,t}/\theta_{i,t}$) and leisure ($l_{i,t}$). The constant $\alpha>0$ is referred to as the utility price of leisure. The specification of intertemporal preferences described by Eqs. (1) and (2) is standard in the literature, despite the contention that is associated with the assumption of time separability (see Deaton and Muellbauer (1980), pp. 124–125, or Hicks (1939), p. 261). This specification of preferences implicitly assumes that characteristics which affect utility, but are not explicitly stated, enter the utility function in an additive way.

2.2. The wealth constraint and simulation of disposable income

Eq. (1) is maximised subject to an age specific credit constraint imposed on liquid net worth, $w_{i,t} \geq D_t$ for household i at age t.[4] D_t can be either exogenously set, or be relaxed subject to the constraint that all households must have repaid their debts by an exogenously defined age, $t_D \leq T$. Liquid net worth is defined as the sum of safe liquid assets, $w_{i,t}^s \in [D_t, \infty)$, and risky liquid assets, $w_{i,t}^r \in [0, \infty)$. Intertemporal variation of $w_{i,t}$ is described by:

$$w_{i,t} = \begin{cases} \hat{w}_{i,t} & t \neq t_{SPA} \\ \left(1-\pi_a^l\right)\hat{w}_{i,t} + (1-\pi_a^p)w_{i,t}^p + (1-\pi_a^o)w_{i,t}^o & t = t_{SPA} \end{cases} \tag{3a}$$

$$\hat{w}_{i,t} = \begin{cases} \pi_{div}\left(w_{i,t-1} - c_{i,t-1} + \tau_{i,t-1}\right) & n_t^a < n_{t-1}^a, t < t_{SPA} \\ w_{i,t-1} - c_{i,t-1} + \tau_{i,t-1} & \text{otherwise} \end{cases} \tag{3b}$$

$$\tau_{i,t} = \tau\left(l_{i,t}, x_{i,t}, n_{i,t}^a, n_{i,t}^c, r_{i,t}^s w_{i,t}^s, r_t^r w_{i,t}^r, pc_{i,t}^o, pc_{i,t}^p, t\right) \tag{3c}$$

$$ln(1 + r_t^r) \sim N\left(\mu_r - \frac{\sigma_r^2}{2}, \sigma_r^2\right) \tag{3d}$$

where $w_{i,t}^p$ denotes wealth held in personal pensions; $w_{i,t}^o$ is wealth held in occupational pensions; π_a^l, π_a^p, and π_a^o are, respectively, the proportions of liquid wealth, private pension wealth, and occupational pension wealth that are used to purchase a life annuity at state pension age, t_{SPA}; π_{div} is the proportion of liquid wealth that is assumed to be lost upon marital dissolution prior to t_{SPA} (to capture the impact of divorce); and $\tau(.)$ denotes disposable income net of non-discretionary expenditure.

As the model has been designed explicitly to undertake public policy analysis, particular care was taken in formulating the effects of taxes and benefits on household disposable incomes. Eq. (3c) indicates that taxes and benefits are calculated with respect to labour supply, $l_{i,t}$; private non-property income, $x_{i,t}$; the numbers of adults, $n_{i,t}^a$, and children, $n_{i,t}^c$; the return to safe liquid assets, $r_{i,t}^s w_{i,t}^s$ (which is negative when $w_{i,t}^s<0$); the return realised on risky liquid assets, $r_t^r w_{i,t}^r$ (possibly negative); private contributions to private and occupational pensions, $pc_{i,t}^p$ and $pc_{i,t}^o$; and age, t.

The lifetime is divided into three periods for the purpose of calculating disposable income: the working lifetime $t<t_{ER}$, early retirement $t_{ER} \leq t < t_{SPA}$, and pension receipt $t_{SPA} \leq t$. In each of these periods of life, household disposable income is calculated by a series of linear splines in taxable income. This aspect of the model is designed to obtain a close match to prevailing tax and benefits policy, while at the same time providing a high degree of flexibility over the policy environment.

Calculation of taxable income for each adult in a household depends on the household's age, with property and non-property income being treated separately. Prior to state pension age, $t<t_{SPA}$, household non-property income $x_{i,t}$ for tax purposes is equal to labour income $g_{i,t}$ less the proportion of pension contributions that is tax exempt, π^{pe}; from state pension age it is equal to labour income plus the proportion of pension annuity income that is taxable, π^{pt}:

$$x_{i,t} = \begin{cases} g_{i,t} - \pi^{pe}\left(pc_{i,t}^o + pc_{i,t}^p\right) & t < t_{SPA} \\ g_{i,t} + \pi^{pt} p_{i,t} & t \geq t_{SPA} \end{cases} \tag{4}$$

$$\text{where: } p_{i,t} = \begin{cases} \chi\left(\pi_a^p w_{i,t}^p + \pi_a^o w_{i,t}^o + \pi_a^l \hat{w}_{i,t}\right) & t = t_{SPA} \\ \left(\dfrac{\pi^s + (1-\pi^s)\cdot\left(n_{i,t}^a - 1\right)}{\pi^s + (1-\pi^s)\cdot\left(n_{i,t-1}^a - 1\right)}\right) p_{i,t-1} & t > t_{SPA} \end{cases} \tag{5}$$

$p_{i,t}$ denotes pension annuity income, and χ is the annuity rate. The annuity purchased at age t_{SPA} is assumed to be inflation linked, and to reduce to a fraction π^s of its (real) value in the preceding year if one member of a couple departs the household in response to the mortality of a spouse.

Where the household is identified as supplying labour, and is younger than state pension age, then non-property (employment) income is split between spouses (in the case of married couples) on the basis of their respective labour supplies. A household that is identified with a single wage earner has all of its non-property income

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

[4] Note that $w_{i,t}^+$ referred to above is related to $w_{i,t}$, with $w_{i,t}^+ = 0$ if $w_{i,t}<0$, and $w_{i,t}^+ = w_{i,t}$ otherwise.

allocated to that one earner; a household with one full-time and one part-time earner has non-property income allocated on the basis of a user defined ratio; and a separate ratio is used to divide non-property income when both spouses of a household are full-time employed. A household without an employed adult has all of its non-property (pension) income allocated to a single spouse.

Similarly, property income is only allocated between spouses for households below state pension age, and who supply some labour. In this case, property income is allocated on the basis of an exogenous ratio that defines the proportion of wealth that is assumed to be held in the name of the lowest earning spouse. Property income, $y_{i,t}$, is equal to the sum of returns from the safe and risky liquid assets:

$$y_{i,t} = \begin{cases} r_t^r w_{i,t}^r + r_{i,t}^s w_{i,t}^s & \text{if } w_{i,t}^s > 0; r_t^r > 0 \\ r_{i,t}^s w_{i,t}^s & \text{if } w_{i,t}^s > 0; r_t^r \leq 0 \\ r_t^r w_{i,t}^r & \text{if } w_{i,t}^s \leq 0; r_t^r > 0 \\ 0 & \text{if } w_{i,t}^s \leq 0; r_t^r \leq 0 \end{cases}. \tag{6}$$

Hence, the model assumes that the interest cost on loans, and losses due to negative risky asset returns cannot be written off against labour income for tax purposes.

The interest rate on safe liquid assets depends upon whether $w_{i,t}^s$ indicates net investment assets, or net debts:

$$r_{i,t}^s = \begin{cases} r^I & \text{if } w_{i,t}^s > 0 \\ r_l^D + \left(r_u^D - r_l^D\right) min\left\{\dfrac{-w_{i,t}^s}{max\left[g_{i,t}, 0.7g\left(h_{i,t}, l_{i,t}^{ft}\right)\right]}, 1\right\}, r_l^D < r_u^D & \text{if } w_{i,t}^s \leq 0 \end{cases}$$

where $l_{i,t}^{ft}$ is household leisure when one adult in household i at age t is full-time employed. This specification for the interest rate implies that the interest charge on debt increases from a minimum of r_l^D when the debt to income ratio is low, up to a maximum rate of r_u^D, when the ratio is high. The specification also means that households that are in debt are treated less punitively if they have at least one adult earning a full-time wage than if they do not. The assumption that the maximum rate of interest is charged when net debt is equal to or greater than the household full-time employment wage reflects the observation that less than 1% of households recorded by the 2000/01 BHPS with some labour income had unsecured debt that exceeded their annual gross labour income.[5]

The model is specified on the assumption that r_t^r is distributed such that $\mu_r < r_l^D$, in which case no rational (and risk averse) household will choose to borrow to fund investment in the risky liquid asset ($w_{i,t}^r > 0$ only if $w_{i,t}^s \geq 0$). Disposable income is consequently given by:

$$\tau_{i,t} = \begin{cases} \hat{\tau}_{i,t} & \text{if } r_t^r \geq 0; w_{i,t}^s \geq 0 \\ \hat{\tau}_{i,t} + r_t^r w_{i,t}^r & \text{if } r_t^r < 0; w_{i,t}^s \geq 0 \\ \hat{\tau}_{i,t} + r_t^s w_{i,t}^s & \text{if } w_{i,t}^s < 0 \end{cases} \tag{7}$$

$$\hat{\tau}_{i,t} = \begin{cases} x_{i,t} + y_{i,t} - tax_{i,t} + benefits_{i,t} - cc_{i,t} - hsg_{i,t} - (1-\pi^{pe})\left(pc_{i,t}^o + pc_{i,t}^p\right) & \text{if } t < t_{SPA} \\ x_{i,t} + y_{i,t} - tax_{i,t} + benefits_{i,t} - hsg_{i,t} + \left(1-\pi^{pt}\right)p_{i,t} & \text{if } t \geq t_{SPA} \end{cases} \tag{8}$$

where $tax_{i,t}$ denotes the simulated tax burden, $benefits_{i,t}$ welfare benefits received, $cc_{i,t}$ non-discretional child care costs net of associated benefits, and $hsg_{i,t}$ non-discretional housing costs net of associated benefits.

The form of the budget constraint described by Eq. (3a) has been selected to minimise the computational burden of the utility maximisation problem. For the purposes of taxation, and in a discrete time model such as this, investment returns can be calculated on the basis of wealth held at the beginning of a given period, or at the end of the period. Calculating taxes with respect to wealth held at the beginning of a period (as it is here) implies that disposable income is made independent of consumption. This is advantageous when consumption is a choice variable, as it implies that the numerical routines that search for utility maximising values of consumption do not require repeated evaluations of disposable income for each consumption alternative that is tested.

[5] See Juster and Shay (1964) and Pissarides (1978) for early examples of the literature concerned with behavioural responses to the nature of liquidity constraints.

2.3. Occupational and private pensions

Up to two pension schemes can be run in parallel in the model. One of these schemes is designed to reflect occupational pensions, and the other is designed to reflect private pensions. The two schemes are identical with the exception that membership and contributions to occupational pensions are exogenously defined, whereas they can be specified as endogenous decisions for private pensions.

Both pension schemes are modelled at the household level, and are defined contribution in the sense that every household is assigned an account into which their respective pension contributions are notionally deposited. The account associated with each scheme, j, accrues a (post-tax) rate of return, r_t^j, which can be specified as uncertain. At the state pension age, t_{SPA}, a proportion of the balance of each pension, π_a^j, is converted into an inflation adjusted life annuity, and the remainder of the pension pot is received as a tax free lump-sum. Hence, a separate account is kept of occupational and personal pensions only until state pension age. The remainder of this section describes how pension rights are simulated to state pension age.

For occupational pensions, $j = o$, membership is exogenously identified with respect to a lower threshold on labour income, π_l^o. Employee contributions to an occupational pension are calculated as a fixed percentage, π^{oc}, of (total) labour income, and can be tax-shielded (as discussed in Section 2.2). Employer (and government) contributions can also be considered for analysis, paid at the rate π_{ec}^o. The balance in the notional account of household i's occupational pension at any age, $t \leq t_{SPA}$, is given by:

$$w_{i,t}^o = \begin{cases} (1 + r_{t-1}^o)w_{i,t-1}^o + (\pi^{oc} + \pi_{ec}^o)g_{i,t} & g_{i,t} \geq \pi_l^o \\ (1 + r_{t-1}^o)w_{i,t-1}^o & o.w. \end{cases} \tag{9}$$

$$ln(1 + r_t^o) \sim N\left(\mu_o - \frac{\sigma_o^2}{2}, \sigma_o^2\right) \tag{10}$$

where $g_{i,t}$ defines aggregate household labour income.

Two alternatives are available for analysis of private pensions, $j = p$. The first is to consider private pensions as exogenous, in a similar fashion to that described for occupational pensions. Alternatively, membership and contribution rates to private pensions can be considered as endogenous decisions. I discus the former case before moving on to the latter.

When membership and contribution rates to private pensions are exogenously imposed, then their simulation progresses in a similar way to that described for occupational pensions, with two notable exceptions. First, membership of private pensions can be specified with respect to an upper limit on labour income, π_u^p. And second, income below the lower bound, π_l^p, is exempt from contributions. Hence, in this case, accrued rights to a private pension are described by:

$$w_{i,t}^p = \begin{cases} (1 + r_{t-1}^p)w_{i,t-1}^p + \left(\pi_{i,t}^{pc} + \pi_{ec}^p\right)\left(g_{i,t} - \pi_l^p\right) & \pi_l^p \leq g_{i,t} \leq \pi_u^p \\ (1 + r_{t-1}^p)w_{i,t-1}^p & \text{otherwise} \end{cases} \tag{11a}$$

$$ln(1+r_t^p) \sim N\left(\mu_p - \frac{\sigma_p^2}{2}, \sigma_p^2\right). \tag{11b}$$

This structure, in conjunction with that assumed for occupational pensions, consequently permits different pension arrangements to be considered for households distinguished by their respective labour incomes.

When private pensions are simulated endogenously, then households choose the proportion of their labour incomes to contribute. The focus on labour incomes implies that households which choose not to supply labour in a given year are excluded from making private pension contributions during that year. As above, contributions to private pensions can be subject to a series of lower (π_l^p) and upper (π_u^p) bounds on eligible incomes, lower (π_l^{pc}) and upper (π_u^{pc}) bounds on contribution rates, and a ceiling on the value of the aggregate pension pot, π_{max}^p. In this case, rights to private pensions evolve as described by Eqs. (11a) and (11b), but with π^{pc} chosen by each household rather than exogenously imposed.

2.4. Labour income dynamics

Up to three household characteristics influence labour income: the household's labour supply decision, the household's latent wage, $h_{i,t}$, and whether the household receives a wage offer $wo_{i,t}$. Households can be exposed to an exogenous, age and relationship specific probability of receiving a wage offer, $p^{wo}(n_{i,t}^a, t)$. This facility is designed to capture the incidence of (involuntary) unemployment/retirement. If a household receives a wage offer, then its labour income is equal to a fraction of its latent wage, with the fraction defined as an increasing function of its labour supply. A household that receives a wage offer and chooses to supply the maximum amount of labour receives its full latent wage, in which case $g_{i,t} = h_{i,t}$. A household that does not receive a wage offer, in contrast, is assumed to receive $g_{i,t} = 0$ regardless of its labour supply decision (implying no labour supply where employment incurs a leisure penalty).

The decision to measure wage potential at the household level rather than at the level of the individual significantly simplifies the analytical problem. Separately accounting for the wages of each adult in a household is properly addressed only by the addition of a state variable to the model where households are comprised of an adult couple. Furthermore, there is significant empirical evidence to suggest that men and women have quite different labour market opportunities, with those of women exhibiting a relatively high degree of heterogeneity.[6] Hence, accounting for the wage potential of individuals could not ignore the sex of adult household members, thereby introducing an additional state variable. These issues are further complicated by the difficulties involved in characterising sex-specific wage generating processes, imperfect correlation of temporal innovations experienced by spouses, and so on. The model side-steps these issues, due to the limitations of existing computing technology.

In the first period of the simulated lifetime, t_0, each household is allocated a latent full-time wage, h_{i,t_0}, via a random draw from a log-normal distribution, $log(h_{i,t_0}) \sim N(\mu_{n^a,t_0}, \sigma^2_{n^a,t_0})$, where the parameters of the distribution depend upon the number of adults in the household, n^a. Thereafter, latent wages follow the stochastic process described by the equation:

$$log\left(\frac{h_{i,t}}{m\left(n_{i,t}^a, t\right)}\right) = \psi\left(n_{i,t-1}^a\right) log\left(\frac{h_{i,t-1}}{m\left(n_{i,t-1}^a, t-1\right)}\right) + \kappa\left(n_{i,t-1}^a, t-1\right)\frac{\left(1-l_{i,t-1}\right)}{\left(1-l_W\right)} + \omega_{i,t} \tag{12}$$

where the parameters $m(.)$ account for wage growth (and depend on age, t, and the number of adults in the household, $n_{i,t}^a$), $\psi(.)$ accounts for time persistence in earnings, $\kappa(.)$ is the return to another year of experience, and $\omega_{i,t} \sim N(0, \sigma_{\omega,n_{i,t-1}^a}{}^2)$ is a household specific disturbance term.

A change in the number of adults in a household affects wages through the persistence term, ψ, the experience effect, κ, and the wage growth parameters m. This model is closely related to alternatives that have been developed in the literature (see Sefton and van de Ven (2004) for discussion), and has the practical advantage that it depends only upon variables from the current and immediately preceding periods ($t-1$, $n_{i,t-1}^a$, $n_{i,t}^a$, $h_{i,t-1}$, $l_{i,t-1}$), which limits the number of characteristics that describe the circumstances of a household (and thereby the number of state variables in the optimisation problem). Furthermore, although the concept of an experience term in a wage regression is not new, its inclusion is an innovation for the related literature (e.g. Low (2005), and French (2005)).[7] Most related studies omit an experience term because it complicates the utility maximisation problem by invalidating two-stage budgeting. We have, however, found that its inclusion enables us to better capture the profile of labour supply during the life-course.

2.4.1. Complicating the standard decision making problem

The preferences defined by Eqs. (1) and (2) are homothetic. Hence, if consumption and leisure were each defined over a continuous domain, and if the price of leisure was exogenous, then the preferred consumption to leisure ratio would be independent of an agent's wealth endowment. In this case, within period utility – Eq. (2) – at the decision making optimum can be expressed in terms of the period specific measure of total expenditure (on goods and leisure), and the maximisation problem can be resolved by two-stage budgeting. This decision making structure is fully consistent with the original analysis of Arrow, so that interpretation of $1/\gamma$ as a measure of relative risk aversion carries over.

However, the focus on discrete labour options, and the inclusion of an experience effect on wages, complicate the intertemporal decision making problem. The discrete nature of labour supply implies that it is not possible to restate intratemporal utility at the decision making optimum as a function of within period total expenditure. Nevertheless, optimised intratemporal utility remains a continuous function of total within-period expenditure (albeit one that is subject to kinks at labour transitions) so that it remains sensible to interpret $1/\gamma$ as a measure of relative risk aversion. Meanwhile, the experience effect on wages implies that the price of leisure is endogenous to the decision making problem, thereby invalidating two stage budgeting.

2.5. Allowing for adults and children

The model allows for two alternatives in the way that the numbers of adults and children in each household evolve. When household demographics are not an important issue of concern, the associated computational burden can be minimised by imposing the same demographic structure by age on all households. In this case, the numbers of adults and children in all households evolve following an exogenously defined age profile. Alternatively, the numbers of adults and children in each household can be defined to evolve stochastically, following a (reduced form) nested logit model. The logit model is comprised of two levels, where the first (highest) determines the evolution of the number of adults in a household, and the second determines the number of children, given the evolution considered for the number of adults and age.

If the number of adults is selected to be uncertain, then a household can be comprised of either a single adult or adult couple, subject to

[6] On recent evidence regarding the labour market experience of women see, for example, Connolly and Gregory (2008).

[7] See, Shaw (1989) for an early study that considers the influence of learning-by-doing within a life-cycle modelling context.

stochastic variation between adjacent years. The fact that children typically remain dependants in a household for a limited number of years implies that it is necessary to record both their numbers and ages when including them explicitly in the model. This substantially increases the computational burden. If, for example, a household was able to have children at any age between 20 and 45, with no more than one birth in any year, and no more than six dependent children at any one time, then this would add an additional 334,622 discrete states to the computation problem (with a proportional increase in the associated computation time). In view of this, the model is currently specified to permit households to have up to three children at each of two discrete ages, so that the maximum number of dependent children in a household at any one time is limited to six.

The logit model that describes the evolution of adults in a household is given by Eq. (13)[8]:

$$s_{i,t+1} = \alpha_0^A + \alpha_1^A t + \alpha_2^A t^2 + \alpha_3^A t^3 + \alpha_4^A dk_{i,t} + \alpha_5^A s_{i,t} \quad (13)$$

where $s_{i,t}$ is a dummy variable, that takes the value 1 if household i is comprised of a single adult at age t and zero otherwise, and $dk_{i,t}$ is a dummy variable that equals 1 if household i at age t has at least one child. With regard to the simulation of births, four separate ordered logit equations are applied; one for each of single and couple households, at each of the specified child-birth ages. The ordered logit equations assumed for the first child birth age, for both singles and couples, do not include any additional household characteristics. The ordered logit equations for the second child birth age include the number of children born at the first child birth age as an additional descriptive characteristic.

3. Model solution procedures

The allowance for stochastic labour income and asset returns implies that an analytical solution to the utility maximisation problem does not exist, so that numerical solution routines need to be employed. Furthermore, the flexible nature of the decision making problem that can be selected for analysis implies that some model configurations are amenable to stylisations that reduce computational burden, while others are not. This section begins with a broad overview of the solution procedure, before describing the alternative numerical routines that are employed.

3.1. Overview of modelling procedure

The solution procedures use backward induction to solve the required inter-temporal Bellman equation. Starting in the last possible period of the household's life, T, it is assumed that households have no children, and are unable to invest in a risky asset or to supply labour. These assumptions imply a deterministic optimisation problem over consumption that is relatively simple to solve for given numbers of adults n_t^a, wealth w_T, and annuity income p_T, omitting the household index i for brevity. Given this level of consumption we can denote the maximum achievable utility, the value function, by $V_T(n_T^a, w_T, p_T)$. Similarly, we can calculate the intermediate measures of welfare:

$$\hat{u}(n_T^a, w_T, p_T) = u\left(\frac{\hat{c}_T(n_T^a, w_T, p_T)}{\theta_T}, 1\right) \quad (14)$$

$$\hat{X}(n_T^a, w_T, p_T) = E_t\left(\frac{1}{(1-1/\gamma)}\left(\zeta_a + \zeta_b \hat{w}_{T+1}^{+}(n_T^a, w_T, p_T)\right)^{1-1/\gamma}\right) \quad (15)$$

where $\hat{c}_T$ and $\hat{w}_{T+1}$ denote the optimised measures of consumption and next period wealth. These functions are calculated at all nodes of a three dimensional grid in the number of adults, wealth, and retirement annuity.

At time $T-1$, households can be permitted to invest in risky assets and to supply labour. Here, the problem reduces to solving the Bellman equation:

$$\begin{aligned} V_{T-1}&(n_{T-1}^a, w_{T-1}, h_{T-1}, wo_{T-1}, p_{T-1}) \\ &= \max_{c_{T-1}, \nu_{T-1}, l_{T-1}} \left\{ \frac{1}{1-1/\gamma} u\left(\frac{c_{T-1}}{\theta_{T-1}}, l_{T-1}\right)^{1-1/\gamma} \right. \\ &\quad + E_{T-1}\left[\frac{\beta_1\delta}{1-1/\gamma}\left(\phi_{1,T-1}\hat{u}(n_T^a, w_T, p_T)^{1-1/\gamma}\right.\right. \\ &\quad + \left.\left(1-\phi_{1,T-1}\right)\left(\zeta_a + \zeta_b w_T^{+}\right)^{1-1/\gamma}\right) \\ &\quad \left.\left. + \beta_1\beta_2\delta^2\phi_{1,T-1}\hat{X}(n_T^a, w_T, p_T)\right]\right\} \end{aligned} \quad (16)$$

subject to the intertemporal dynamics that are described above, where wo_{T-1} is a wage offer identifier taking the value 1 if a wage offer is received and zero otherwise, and ν_{T-1} is the proportion of liquid wealth invested in the risky asset. This optimisation problem is solved for the $T-1$ value function, at each node of the five dimensional grid over the permissable state-space. The expectations operator is evaluated in the context of the log-normal distributions assumed for wages and risky asset returns using the Gauss–Hermite quadrature procedure, which permits evaluation at a set of discrete abscissae. Interpolation methods are used to evaluate the value function at points between the assumed grid nodes throughout the simulated lifetime.

Solutions for earlier ages then proceed via backward induction, based upon the solutions obtained for later ages. Prior to t_{SPA}, solutions may also be required for pension contributions, and the state space may be expanded to include children, and the two pension assets permitted in the model.

Having solved for utility maximising behavioural responses at grid nodes as described above, the life-courses of individual households are simulated by running households forward through the grids. This is done by first populating a simulated sample by taking random draws from a joint distribution of all potential state variables at the youngest age considered for analysis. The behaviour of each simulated household, i, at the youngest age is then identified by reading the decisions stored at their respective grid co-ordinates. Given household i's characteristics (state variables) and behaviour, its characteristics are aged one year following the processes that govern their intertemporal variation. Where these processes depend upon stochastic terms, random draws are taken from their defined distributions (commonly referred to as Monte Carlo simulation). This process is repeated for the entire simulated life of each household. Analyses are then based upon the data generated for the simulated cohort.

3.2. Details of solution routines

The value function in this problem is neither smooth, nor concave (though it is designed to be non-decreasing and continuous). Non-smoothness arises because of our focus on a discrete labour supply decision, the allowance for quasi-hyperbolic discounting, and the consideration of a flexible budget set that may be non-convex (due, for example, to means testing of welfare benefits). Non-concavities of the value function imply that the optimisation problem (16) can have

[8] When children are not modelled explicitly, then the cubic term in age and the dummy variable for children is omitted from the logit equation.

local maxima. This observation, combined with the idiosyncratic nature of the model and its level of complexity, emphasises the importance of checks to determine the validity of model solutions, and for methods to determine the degree of numerical accuracy obtained.

The model includes three tools for assessing the accuracy of the numerical solutions obtained: increasing the solution detail, varying the interpolation methods, or varying the numerical search routines. The first of these is the most simple, and often the most powerful of the three. When *increasing the solution detail*, it is possible to choose the size and number of grid points adopted for each of the continuous state variables that are included in the decision making problem, as well as to increase the number of abscissae used in the Gaussian quadrature. Increasing the grid points provides a more detailed solution of the utility maximising problem, though it can also imply a rapid increase in computational burden. Increasing the grid points in multiple dimensions increases the computational burden geometrically rather than arithmetically; commonly referred to as the curse of dimensionality.

The model allows for either linear or cubic *interpolation methods* to evaluate behaviour between discrete grid points, where cubic interpolation is preferable when the value function displays a reasonable degree of smoothness. If cubic interpolation is selected, then the model uses cubic interpolation wherever the surface is identified as displaying an acceptable degree of smoothness, and applies linear interpolation otherwise.[9] Cubic interpolation does, however, imply a larger computational burden, requiring evaluations at 4^n grid points over n dimensions, compared with 2^n grid points required for linear interpolation.

Finally, the model includes three alternative *numerical search routines* to identify utility maximising decisions. The most simple and computationally burdensome is a 'brute force' approach that conducts a comprehensive grid search over the permissable decision space. A second search routine uses Brent's method to search for an optimum by repeatedly evaluating the value function associated with alternative decision options. This approach, which combines parabolic interpolation with a golden section search, is efficient where the surface over which the search is conducted is reasonably well behaved. The third search routine uses the Bus and Dekker (1975) bisection algorithm to search for a zero in relation to Euler conditions, and is only appropriate where time inconsistent preferences are suppressed (so that $\beta_1=\beta_2=1$ in Eq. (1)). As the last two of these procedures are not designed to distinguish between global and local optima, a supplementary search routine is also provided, which tests over a localised grid above and below an identified optimum for a preferred decision set. If a preferred decision set is identified, then the supplementary routine searches recursively for any further solutions. This process is repeated until no further solutions are found. Of all feasible solutions, the one that maximises the value function is selected.

4. Calibration approach

The calibration that is reported here was undertaken by a two stage process. In the first stage, estimates for observable model parameters were calculated from available survey data. Given the estimates obtained in the first stage, values for the unobserved parameters of the model were adjusted in the second stage. This second stage involved generating a simulated population cohort based upon assumed starting values for the unobserved model parameters, and the model parameters estimated in the first stage (following the methods described in Section 3). 'Simulated moments' for various population characteristics were calculated from the data generated for the simulated population, and these were compared against associated 'sample moments' estimated from survey data. The unobserved model parameters were then adjusted to match the simulated moments to the sample moments.

4.1. The model specification

A very large number of trial parameter combinations are usually required to conduct the second stage of the model's calibration. This gives rise to a trade-off between including sufficient detail in the model to provide a close reflection of the practical reality, and limiting the computation time to permit sufficient parameter combinations to be explored. For the calibrations that are reported here, households were assumed to have time consistent preferences ($\beta_1=\beta_2=1$ in Eq. (1)), and were distinguished by the following seven characteristics:

— Age	— Number of adults	— Net liquid assets	— Wage offers
—	— Private pension rights	—	—
— Wage rates	— Time of death		

This restricted model focuses on decisions over labour supply (including the possibility of part-time employment), consumption, pension saving, and bequests, given a household's age, its number of adults, liquid assets, pension wealth, wage offer, wage rate, and survival. Household decisions were considered at annual intervals between ages $t_0=20$ and $T=110$, with labour supply possible to age 69. State pension age was set to $t_{SPA}=65$, the pension age that prevailed for men in 2005. Uncertainty was taken into consideration for the intertemporal development of the number of adults in a household, wage offers, wage rates, and the time of death — age, liquid wealth, and pension wealth were all considered to evolve deterministically.

As noted above, the model solves decision making problems by dividing the permissable state space into a series of grids. The domains of wages and wealth between ages 20 to 69 were each divided into 35 points using a log scale. The domain of pension wealth between ages 20 to 64 was divided into 16 points using a log scale. It was assumed that 25% of pension wealth at $t=t_{SPA}$ is taken as a tax free lump, with the remainder taken as a retirement annuity. The domain of retirement annuity was divided into 16 points using a log scale between ages 65 and 69. From age 70 to age 110, the wealth and retirement annuity domains were each divided into 151 points using a log scale.

Two additional dimensions – one to reflect the number of adults in a household and one to reflect a wage offer – complete the grids that were considered for the calibration. These grid dimensions differ from those described above in that they refer to characteristics that take discrete values. From age 20 to 95 (inclusive), solutions were required for single adults and couples; from age 96 all households were assumed to be comprised of a single adult. Between ages 20 and 69, solutions were required for households with and without a wage offer.

This specification of the model required utility maximising decisions to be numerically evaluated for 5,447,667 different combinations of household characteristics, for each alternative parameter combination tested as part of the calibration process.[10] For reference, this specification of the model takes 44 min to run on a computer with an Intel Core2 Duo T9300 (2.50 GHz) processor and 2 Gb of RAM, and 15 min to run on a workstation with dual Intel Xeon E5410 (2.33 GHz) processors and 4 Gb of RAM.

[9] The criterion for identifying smoothness involves distinguishing the "inner" 2^n points in closest proximity to the co-ordinate to be interpolated, from the "outer" 4^n points considered in evaluating the cubic interpolation. If the smallest difference between any of the outer points and any of the inner points is more than 5 times the maximum difference between the inner points, then the model reverts to linear interpolation.

[10] $=(69-19).35.35.16.2.2+(95-69).151.151.2+(110-95).151$.

4.2. The specific calibration strategy adopted

Three basic sets of model parameters required adjustment in the second stage of the calibration process[11]:

- wage parameters, which are described in Section 2.4: $p^{wo}(n_{i,t}^{a}, t)$, $m(n_{i,t}^{a}, t)$, $\kappa(n_{i,t}^{a}, t)$, μ_{n^a,t_0}, σ_{n^a,t_0}^2, $\sigma_{\omega,n_{i,t}^a}^2$.
- preference parameters, which are described in Section 2.1: γ, ε, δ, α.
- and a subset of tax parameters, which are described below.

Although the monolithic nature of the model complicates piecewise parameter adjustment, extensive experimentation revealed that some aspects of the model calibration are more sensitive to alternative parameter specifications than are others. This observation motivated the strategy that was adopted for the current study, which is structured round an ordered series of calibration sub-problems. The calibration strategy focussed upon matching the simulated to sample moments in the following order:

1. employment income
2. disposable income
3. the time profile of employment decisions
4. the time profile of consumption
5. the relationship between wealth and employment late in the working lifetime
6. the relationship between employment and consumption late in the working lifetime.

The calibration procedure began by assuming a set of starting values for the parameters to be calibrated. These starting values were selected to broadly match the model to survey data, with a tendency to understate the preference for leisure (to ensure sufficient data to construct simulated moments for employment income). The parameters of the model were then adjusted to match the simulated to sample statistics in a cascading process, focussing first upon the moments described in 1 above, and then moving up through the list, re-adjusting parameters associated with lower levels each time parameters associated with higher levels were revisited. Each of the six stages of the calibration is now described in turn.

4.2.1. Matching employment income

The parameters μ_{n^a,t_0} and σ_{n^a,t_0}^2 were adjusted to match the distribution of household employment incomes at age 20. At the same time, the wage parameters $m(n_{i,t}^{a}, t)$ were adjusted to match simulated moments to sample estimates for geometric mean employment income. As these aspects of the calibration involve the same number of model parameters as the number of matched moments, this part of the calibration is exactly identified, which facilitates a very close matching between simulated and sample moments. The same is not true, however, for $\sigma_{\omega,\, n_{i,t}^a}^2$, which is adjusted to capture the age trend described by the variance of (log) employment income. These model parameters were adjusted by an iterative search routine designed to minimise the mean absolute difference between age specific geometric mean and log variances of employment income generated by the model, and moments estimated from survey data.

4.2.2. Matching disposable income

The stylisations assumed for the model produced disparities between the simulated and sample moments of employment income on the one hand, and of disposable income on the other. One of the most important considerations in this respect is the role played by accommodation costs in determining benefits eligibility in the UK. This represents a problem for the simulation model, as the model does not take explicit consideration of owner-occupation, nor does it reflect the incidence of rental costs or council taxes.

A wedge between the simulated and sample moments of employment income and disposable income poses a problem for model calibration because each measure of income tends to influence different aspects of behaviour. Employment incomes – and their relation to child care costs, housing costs, and welfare benefits – tend to have an important bearing on labour supply decisions, whereas the age profile of disposable income has an important bearing on savings decisions. To address these concerns, the current calibration adjusts the scale of housing costs to match age specific geometric means of disposable income generated by the model against associated sample moments, given the parameters assumed for employment income. Like the first stage of the model calibration, this stage was found to exhibit sufficient stability to be implemented by an automated adjustment procedure, where the procedure is designed to minimise the mean absolute difference between simulated and sample moments (of disposable income).

4.2.3. Matching the employment profile

The experience parameters $\kappa(n_{i,t}^{a}, t)$ were adjusted to match age profiles of full-time and part-time employment. The utility price of leisure, α, was adjusted to match average rates of full-time employment. To avoid over-identifying the model, the age dependency that can be permitted for the probabilities of a low wage offer was suppressed, so that the parameters $p^{wo}(n_{i,t}^{a}, t) = p^{wo}(n_{i,t}^{a})$ for all t: these two parameters (one for each of singles and couples) were adjusted to match rates of non-employment. These parameters were adjusted manually.

4.2.4. Matching the consumption profile

Age profiles for the geometric mean of consumption were matched by manually adjusting the discount rate, δ. Lower values of δ imply less patience in the preference relation, and consequently produce consumption profiles that embody less saving. This parameter was adjusted manually.

4.2.5. Matching the relationship between wealth and retirement

Increasing the intra-temporal elasticity, ε, tends to decreases the demand for leisure relative to consumption (equivalent to later retirement) for high income households, and *vice versa* for low income households, Sefton et al. (2008). ε, was consequently adjusted to match the simulation model to rates of employment participation by wealth quintile observed in survey data between ages 60 and 64. The associated calibration was undertaken manually.

4.2.6. Matching the relationship between consumption and retirement

Given ε, increasing $1/\gamma$ tends to make consumption and leisure direct substitutes, thereby increasing the correlation between labour income and consumption, Sefton et al. (2008). γ was consequently adjusted to reflect the influence of employment on consumption, as reflected by separate estimates for single adults and couples aged between 60 and 64 for the following linear equation:

$$log\left(c_{i,t}\right) = \alpha_0 + \alpha_1 emp_{i,t} + \varepsilon_{i,t} \tag{17}$$

where $emp_{i,t}$ is a dummy variable taking the value 1 if household i at age t chose some employment and 0 otherwise.

4.3. The survey data

There is no internally consistent basis upon which to calibrate the model. This is because the model does not capture real-world

[11] Conspicuous exceptions from the list of parameters provided here are (ζ_a, ζ_b), which govern preferences over bequests, and $\psi(n_{i,t}^{a})$ which describes the intertemporal persistence of the wage generating process. These parameters are not "tied down" by the calibration strategy that is described below, and were consequently exogenously set.

uncertainty in the prevailing economic environment, including variation over a range of parameters that reflect the evolving tax and benefits system, conditions of the macro-economy, household demographics, and so on.[12] The related literature has consequently focussed on three alternative data sources, which are all imperfect in their own ways. The first is (pseudo) panel data for an actual cohort. These data are affected by time and cohort effects that make them unrepresentative for the population in general.[13] Alternatively, it is possible to control for time and cohort effects by econometric estimation. In this case, collinearity between age, cohort and time effects implies the need to introduce an additional restriction to permit identification. One popular 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 (pseudo) panel data set used for estimation. Calibrating the model against this sort of age profile implicitly assumes – as do calibrations based upon data for a single population cohort – that behaviour is invariant to changes in the policy environment that occurred during the period of estimation. This assumption is difficult to maintain when one of the objectives of the analysis is to consider behavioural responses to alternative policy experiments.[14]

The third approach, which was adopted for the current paper, involves matching the model against suitably adjusted age profiles described by cross-sectional survey data. The adjustments are necessary to reflect the fact that cross-sectional data do not usually include a description of agent expectations regarding their future circumstances; expectations that are crucial in determining the decisions projected by the model. It is consequently assumed that households behave as though they will be subject to the policy environment that prevailed in 2005 for the remainder of their lives; that they expect their labour incomes to grow at a constant rate; that the anticipated probabilities governing cohabitation reflect 2005 cross-sectional data; and that the probabilities governing mortality reflect official projections for the cohort aged 20 in 2005.

5. Estimates for observable model parameters

The model parameters for which exogenous estimates were obtained are concerned with five key issues: taxation, life expectancy, intertemporal indexing, the distinction between alternative labour classes, and household demographics. A conspicuous omission from this list is the treatment of wages, the parameters for which were addressed as part of the second stage calibration to ensure the approach taken to account for sample selection is consistent with the wider analytical framework. The specification of these five aspects of the model is described in turn below.

5.1. Taxation

Taxes and benefits considered for the calibration were specified to closely reflect the rates and thresholds of schemes that were applied in the UK in 2005. Recall from Section 2.2 that the model divides the lifetime into three periods for the purpose of calculating disposable income: the working lifetime $t<t_{ER}$, early retirement $t_{ER}\leq t<t_{SPA}$, and pension receipt $t_{SPA}\leq t$. Taxes and benefits during the working lifetime, $t<t_{ER}=60$, were specified to closely reflect schedules reported in the April 2005 edition of the *Tax Benefit Model Tables* (TBMT), issued by the Department for Work and Pensions.[15] The period of early retirement in the model was defined to capture the effects of income taxes and the Pension Guarantee. The Pension Guarantee was a benefit payable to individuals from age 60 in 2005, and was subject to a 100% withdrawal (claw-back) rate on private income until the benefit was exhausted. Finally, the period of pension receipt, $65=t_{SPA}\leq t$, was specified to reflect income taxes, the basic state pension, and the Pension Credit. The basic state pension was represented by a flat rate (universal) benefit, and the Pension Credit by a means tested benefit subject to a 40% withdrawal rate on private income.

5.2. Life expectancy

The survival probabilities assumed for calibrating the model were based upon the period life tables published by the Government Actuary's Department (GAD). These data were used to calculate the age specific probabilities of survival for a couple, where both members of the couple were aged 20 in 2005. The projections embody official estimates for the trend improvement of future survival rates, and their use in calibrating the model is consistent with our assumption of rationally behaved households.

The official data permit survival rates to be calculated to age 94. Age specific survival probabilities between 95 and 110 were exogenously specified to obtain a smooth progression from the official estimate at age 94 to a 0% survival probability at age 110. These probabilities are reported in Table 3 from age 61.

5.3. Intertemporal indexing

A number of aspects of the economic environment were subject to indexing as part of the calibrations that are reported here. Wages and consumption statistics estimated from cross-sectional data were adjusted to reflect growth of 1.3% p.a., based upon ONS data which indicate that the consumer price deflator grew 3.0% p.a., and that the rate of nominal wage growth was 4.3% p.a. on average between 1990 and 2007. Tax thresholds were indexed to grow at the rate of 0.1% p.a. This is the real rate of growth of the basic rate tax threshold between 2000/01 and 2006/07. A growth rate of 1.6% p.a. was assumed for welfare benefits, based upon a weighted average of the real rates of growth observed for the Child Benefit, unemployment benefits, and the Basic State Pension. Both housing and child care costs were assumed to grow at the same rate of 2.1% p.a. This rate is based upon a nominal rate of growth of 5.7% p.a. in rental costs between 1987 and 2006, and an associated rate of inflation of 3.6% p.a. The rate of return to positive holdings of net liquid wealth (r^l) was assumed to be 4.0% p.a. This is double the real rate of return on cash held with banks, building societies and ISAs during the period between April 1999 and December 2006, and was selected because simulated net liquid wealth is a composite asset that includes investment classes which usually provide a higher yield than cash balances. The average real interest charges applied between January 1995 and January 2006 to credit card loans, and to unsecured personal loans in excess of £10,000, were, respectively, 18.0% and 9.7% per annum. The lower limit cost of debt (r_l^D) was consequently set to 8% per annum, and the upper limit (r_u^D) to 18%.

5.4. Distinguishing the implications of alternative labour supply decisions

The labour supply and wage effects of alternative labour supply decisions were estimated using data from the 2005/06 Family Resources Survey. Weighted averages were calculated for the number of hours worked, and for gross labour income, distinguishing population sub-samples by the number of adults in a household and labour market status. The proportions considered for the calibration are reported in Table 1.

[12] See Nelissen (1998) for an early simulation study that takes into consideration the influence of the evolving tax structure. It is of note that this critique may be applied more broadly to much of the literature that is concerned with behavioural responses in an evolving economic environment.

[13] See, for example, Attanasio et al. (2005).

[14] See, for example, Sefton et al. (2008).

[15] See http://www.dwp.gov.uk/asd/tbmt.asp.

Table 1
Model parameters to distinguish the effects on leisure and labour income of alternative labour supply decisions.
Source: author calculations from 2005/06 Family Resource Survey data.

Number of adults	Employment decision	Leisure cost[a]	Propn of full-time wage
1	Not employed	0.000	0.000
1	Part-time employed	0.480	0.357
1	Full-time employed	1.000	1.000
2	Not employed	0.000	0.000
2	1 part-time employed	0.220	0.188
2	1 full-time employed	0.514	0.614
2	1 part-time and 1 full-time	0.774	0.808
2	2 full-time employed	1.000	1.000

[a] Leisure cost specified relative to full-time employment of all adult members.

Table 2
Regression statistics for logit model of relationship status.

Variable	Coefficient	std. error
Constant	−3.37007	0.14255
Age	−3.16E−02	5.64E−03
Age^2	6.00E−04	5.24E−05
Single	5.61881	0.03347
Correct predictions	0.94663	
Sample	79,551	
Proportion single	0.33078	

Author calculations based on pooled data from waves 1 to 15 of the BHPS.

5.5. Household demographics

The logit equation used to model the number of adults in a household was estimated using pooled data derived from waves 1 (1991/92) to 15 (2005/06) of the BHPS. Regression statistics are reported in Table 2.

As is implicit in the description of the calibration strategy that is stated above, dependant children were modelled deterministically in the simulations used to calibrate the model. In this case, the number of children in a household was considered to be a function of the household's age and number of adults. The number of children assumed for a simulated household was specified to reflect age specific averages (by relationship status) observed in the same sample population that was used to obtain estimates for the processes used to generate the number of adults in a household, as discussed above.[16] These age specific averages are reported in Table 3.

6. Calibrated model parameters

Following an extensive search, the calibrated parameter values are reported in Tables 4–6. This section begins by interpreting the calibrated parameters in context of the exogenously assumed decision making environment. Sensitivity of the match obtained between simulated and sample moments is then discussed, before drawing conclusions regarding the apparent inconsistencies between the behavioural model as it is described here and associated observations drawn from survey data.

6.1. Calibrated parameters and match to sample moments

Although no consensus exists regarding a generally representative value for relative risk aversion, it is common for comparable studies to assume a value between 1.0 and 4.0. This range is consistent with the calibrated value of 2.9 reported in Table 4. The relative values of the intra-temporal elasticity (ε) and relative risk aversion ($1/\gamma$) imply that consumption and leisure are direct substitutes, and simulations under the calibrated parameter specification imply an average intertemporal elasticity of substitution for consumption of 0.41, which lies within the range of values identified in the related literature.[17] The discount factor indicates slightly less impatience than the assumed real rate of return (3.6% c.f. 4.0% per annum), and the utility price of leisure is in the region of 1.0 by construction.[18] Finally, the calibrations assume that full-time employment reduces leisure time by 50% prior to age 65 (state pension age), and by 67% from age 65. I discuss the reasons for this disparity in the leisure price of employment at further length below.

The calibrated age specific wage experience effects are stronger for singles than for couples throughout the working lifetime. The age profiles assumed for the experience effect are non-decreasing for both singles and couples, rising from a minimum at the beginning of the simulated lifetime of 2.5% per annum for singles and 2.1% per annum for couples, to peak at 15% and 13% respectively at the end of the working lifetime. Hence, if a single person aged 20 chooses to work full-time, then they can expect to have a latent wage at age 21 that is 2.5% higher than it would have been had they chosen not to work at all.

The age specific weights on housing costs assumed for analysis describe similar profiles for both singles and couples, increasing from zero at the beginning of the simulated lifetime, to peak in the late 50s, before falling away at higher ages. The scale of the calibrated housing costs are, however, quite different for singles and couples. For singles, the weight on housing costs peaks at around 60% of the statistics assumed in the DWP Tax Benefit Model Tables. For couples, in contrast, the housing costs weight peaks at just under 100%. This disparity in scale is perhaps not very surprising, given that the maximum housing costs assumed for couples are less than £10 per week higher than those for singles after adjusting for the number of children in a household. The calibrations consequently suggest that the housing costs assumed for couples in the simulated tax and benefits framework may understate the practical reality. Nevertheless, it is important to bear in mind that the calibrated housing costs adjustments reflect the aggregate disparity between the simulated and actual tax and benefits functions. The age profiles reported for the housing costs weights in Table 5 consequently aggregate a range of characteristics that extend beyond housing costs *per sé*.

It is useful to discuss this issue in a little more depth. Although the practical reality regarding the tax and benefit treatment of households in the UK is very complex, six principal factors can be identified as determinants of disposable income. They are, in approximate order of importance, the private income, labour force status, and age of each household member, the number of dependant children, housing costs, and the health status of household members. To recap, the calibration that is considered here takes explicit account of the first two of these factors, subject to the implicit assumption that children do not earn income. A single age is attached to each household for the purposes of taxation, which implicitly focuses analysis on same-aged spouses. The numbers of dependant children and housing costs are specified as deterministic and non-stochastic functions of the numbers of adults and household age. Finally, the model ignores variation in health, thereby focussing upon the healthy population subgroup. It is not immediately clear what the stylisations assumed by the model imply for mean incomes by age, which is the determining factor for the housing costs adjustments that are reported in Table 5. It is, however, likely that the omission of household heterogeneity that is relevant for determining taxes and benefits in practice is likely to dampen the simulated dispersion of disposable incomes, relative to the practical reality.

[16] The age profiles considered for the number of children in a household by relationship status were smoothed using a five point moving average (between annual age groups).

[17] Calculated at population averages for consumption (£592 per week), leisure (0.659), and the equivalence scale (1.702*470) between ages 20 and 50, weighting each age equally.

[18] This is achieved by multiplying the equivalence scale by 470, to normalise equivalised consumption (c/θ).

Table 3
Estimates for observed age specific model parameters.

Number of children						Mortality rates			
Age	Singles	Couples	Age	Singles	Couples	Age	Mortality	Age	Mortality
20	0.28	0.70	45	0.43	1.29	61	0.001	86	0.016
21	0.33	0.69	46	0.39	1.17	62	0.001	87	0.017
22	0.40	0.66	47	0.35	1.05	63	0.001	88	0.020
23	0.46	0.65	48	0.30	0.93	64	0.001	89	0.034
24	0.52	0.65	49	0.25	0.79	65	0.001	90	0.043
25	0.58	0.67	50	0.21	0.67	66	0.001	91	0.046
26	0.62	0.71	51	0.16	0.55	67	0.001	92	0.053
27	0.65	0.79	52	0.13	0.45	68	0.002	93	0.044
28	0.69	0.88	53	0.10	0.36	69	0.001	94	0.072
29	0.74	0.99	54	0.08	0.29	70	0.002	95	0.102
30	0.80	1.09	55	0.06	0.24	71	0.002	96	0.135
31	0.83	1.20	56	0.05	0.20	72	0.003	97	0.170
32	0.89	1.32	57	0.04	0.16	73	0.003	98	0.209
33	0.92	1.42	58	0.03	0.13	74	0.003	99	0.250
34	0.94	1.51	59	0.02	0.11	75	0.005	100	0.295
35	0.91	1.59	60	0.02	0.10	76	0.003	101	0.343
36	0.90	1.65	61	0.02	0.09	77	0.005	102	0.396
37	0.85	1.68	62	0.02	0.08	78	0.007	103	0.452
38	0.81	1.70	63	0.01	0.07	79	0.007	104	0.513
39	0.75	1.70	64	0.01	0.06	80	0.007	105	0.580
40	0.70	1.68	65	0.01	0.05	81	0.010	106	0.651
41	0.65	1.64	66	0.01	0.05	82	0.011	107	0.728
42	0.60	1.59	67	0.00	0.04	83	0.010	108	0.812
43	0.53	1.50	68	0.00	0.04	84	0.015	109	0.902
44	0.49	1.40	69+	0.00	0.00	85	0.019	110	1.000

Notes: average numbers of children by age and relationship status estimated from waves 1 to 15 of the BHPS morality rates calculated on cohort life expectancies for couple aged 20 in 2005 published by the ONS.

Table 4
Calibrated values for age independent model parameters.

	Value	Parameter[a]
Preference parameters		
Relative risk aversion	2.900	$1/\gamma$
Intra-temporal elasticity	0.560	δ
Discount factor	0.964	δ
Utility price of leisure	1.221	α
Bequest motive: constant	1.000	ζ_a
Bequest motive: slope	0.100	ζ_b
	Singles	Couples
Relationship specific parameters		
Prob of low wage offer	0.267	0.039
Housing weight (aged 79+)	0.002	0.434
Per period income shocks[b]	0.201	0.126
Age 20 income distribution[b]	0.383	0.381
Age 20 mean of log income	5.435	6.062
Leisure cost of full-time work prior to age 65 (SPA)	0.500	0.500
From age 65	0.670	0.670

[a] Preference parameters, as defined in Eqs. (1) and (2).
[b] Specified as standard errors (of logs).

We turn now to the relation between the simulated statistics obtained from the model using the calibrated parameter values reported in Tables 4 and 5, and the associated sample statistics calculated from survey data, displayed in Figs. 1–4. All monetary values are reported relative to average gross annual income for all full-time employees in the UK during 2005/06, equal to £474 per week.[19]

Figs. 1–4 reveal a close match between the simulated and sample moments. Starting with Fig. 1, the top two panels indicate that the model does a very good job of capturing the incidence of full-time, part-time and non-employment described by survey data throughout the simulated lifetime. Panel C of the figure also indicates that the model does a good job of capturing the influence of wealth on the timing of retirement, with poorer households choosing to retire sooner than richer households.

Fig. 2 indicates that the match obtained between the simulated and sample moments for employment income is less close than the match obtained for rates of employment participation. In particular, the top panel of Fig. 2 indicates that the model closely reflects the geometric mean of employment incomes described by survey data up to age 55, but tends to overstate the geometric mean thereafter. The largest disparities between the simulated and sample moments are observed at age 64, just prior to state pension age in the simulations. In this regard, it is important to note that the disparities between the model and survey data that are revealed by Panel A of Fig. 2 are affected by sample selection in the timing of retirement: by age 60 more than half of all adults in a couple and 70% of single adults no longer choose to work under the simulations. The bottom panel of Fig. 2, by contrast, indicates a relatively close match between simulated and sample moments for the dispersion of employment incomes throughout the simulated lifetime.

That a poorer match should be obtained between simulated and sample moments for employment income, relative to rates of employment participation, may come as something of a surprise, given the relative flexibility of the model parameters that govern the simulation of employment incomes. If employment profiles were calibrated by adjusting 14 model parameters[20] – most of which have an opaque influence on age specific employment incentives – whereas moments of employment income were calibrated by adjusting 106 model parameters[21] – all of which have a clear correspondence to age specific moments of employment income – then why was a better fit obtained to the former than the latter? The inter-related nature of the simulated moments upon which the calibration is based complicates piece-wise explanation of the incentives that underlie this aspect of the calibration. The current subsection consequently proceeds by discussing the remaining moments against which the model was calibrated, and I return to discuss the disparities that are identified here at length in the following subsection.

The top panel of Fig. 3 indicates that, in contrast to the statistics reported for employment income, the model does a very good job of matching the geometric mean of disposable incomes throughout the simulated lifetime. This result is not particularly surprising, as the age specific weights on housing costs were adjusted to match the model to sample moments for geometric mean disposable income. Although not included as a formal part of the calibration, associated statistics of dispersion are reported in the bottom panel of Fig. 3. These indicate that the model tends to understate the dispersion of disposable incomes by age for both singles and couples, which is best interpreted as providing a measure of the extent to which the simulated tax and benefits system departs from the practical reality (as noted above).

At age 65, the simulated geometric mean of disposable income jumps up for both singles and couples in a way that is not evident in the associated sample moments. This is in contrast to the associated moments for employment income displayed in Fig. 2, which jump down at age 65. Note that the jump in the geometric mean of

[19] Average gross annual income for all full-time employees reported for winter 2005/06, reported in *Labour Force Survey Historical Quarterly Supplement*, Table 37, available from the National Statistics website: http://www.statistics.gov.uk/.

[20] The intra-temporal elasticity, ε, the utility price of leisure, α, four measures of the leisure price of full-time employment (singles and couples, pre age 65 and post 65), six parameters of the experience effect on wages, κ, (in bold in Table 5), and two age independent probabilities of receiving a low wage offer, p^{wo}, one for singles and one for couples.

[21] 50 age specific parameters of wage growth, $m(t)$, two parameters for the distribution of employment incomes at the beginning of simulated life, μ_{t_0}, $\sigma_{t_0}^2$, and one parameter for the uncertainty of wages from one year to the next, σ_ω^2, for each of singles and couples.

Table 5
Calibrated values for age specific model parameters of single adults.

Age	Trend income[a]	Experience effect[a]	Housing weight[b]	Age	Trend income[a]	Experience effect[a]	Housing weight[b]
20	78.57	**0.02500**	0.01477	45	223.15	0.13077	0.48664
21	85.91	0.02923	0.03494	46	225.97	**0.13500**	0.50246
22	96.83	0.03346	0.05485	47	228.41	0.13607	0.51733
23	108.77	0.03769	0.07462	48	228.70	0.13714	0.53121
24	122.71	0.04192	0.09407	49	226.35	0.13821	0.54407
25	137.93	0.04615	0.11313	50	222.58	0.13929	0.55598
26	152.44	0.05038	0.13171	51	220.42	0.14036	0.56701
27	165.64	0.05462	0.14976	52	219.15	0.14143	0.57703
28	178.07	0.05885	0.16732	53	216.53	0.14250	0.58554
29	191.18	0.06308	0.18452	54	211.88	0.14357	0.59158
30	204.88	0.06731	0.20157	55	207.03	0.14464	0.59389
31	215.12	0.07154	0.21874	56	204.46	0.14571	0.59130
32	222.17	0.07577	0.23632	57	201.88	0.14679	0.58289
33	229.30	0.08000	0.25449	58	199.30	0.14786	0.56804
34	235.99	0.08423	0.27335	59	196.72	0.14893	0.54631
35	242.81	0.08846	0.29289	60	194.14	**0.15000**	0.51752
36	251.86	0.09269	0.31299	61	191.56	0.15000	0.48185
37	258.09	0.09692	0.33346	62	188.98	0.15000	0.43983
38	258.57	0.10115	0.35411	63	186.40	0.15000	0.39224
39	256.26	0.10538	0.37468	64	183.83	0.15000	0.33999
40	253.08	0.10962	0.39497	65	181.25	0.15000	0.28392
41	249.73	0.11385	0.41477	66	178.67	0.15000	0.22507
42	242.83	0.11808	0.43392	67	176.09	0.15000	0.16400
43	232.23	0.12231	0.45233	68	173.51	0.15000	0.10429
44	224.41	0.12654	0.46991	69	170.93	0.15000	0.03948

[a] Tend income = m(t), and experience effect = kappa in Eq. (12), values in bold adjusted during calibration.
[b] Housing weight refers to age specific factor used to adjust housing costs.

Table 6
Calibrated values for age specific model parameters of adult couples.

Age	Trend income[a]	Experience effect[a]	Housing weight[b]	Age	Trend income[a]	Experience effect[a]	Housing weight[b]
20	371.47	**0.02100**	0.00002	45	711.68	0.04888	0.73464
21	372.47	0.02212	0.03294	46	709.74	**0.05000**	0.76513
22	384.18	0.02323	0.06634	47	696.16	0.05571	0.79684
23	397.51	0.02435	0.09986	48	675.52	0.06143	0.82891
24	414.61	0.02546	0.13364	49	651.56	0.06714	0.86001
25	430.49	0.02658	0.16757	50	628.74	0.07286	0.88875
26	442.83	0.02769	0.20155	51	607.07	0.07857	0.91401
27	453.83	0.02881	0.23538	52	580.58	0.08429	0.93521
28	469.38	0.02992	0.26886	53	554.91	0.09000	0.95228
29	488.68	0.03104	0.30179	54	533.91	0.09571	0.96556
30	506.18	0.03215	0.33399	55	511.36	0.10143	0.97553
31	524.08	0.03327	0.36532	56	498.19	0.10714	0.98250
32	547.24	0.03438	0.39567	57	485.02	0.11286	0.98616
33	572.64	0.03550	0.42493	58	471.85	0.11857	0.98532
34	594.28	0.03662	0.45298	59	458.68	0.12429	0.97807
35	609.97	0.03773	0.47980	60	445.51	**0.13000**	0.96248
36	621.46	0.03885	0.50545	61	432.35	0.13000	0.93741
37	633.18	0.03996	0.53017	62	419.18	0.13000	0.90294
38	646.88	0.04108	0.55428	63	406.01	0.13000	0.86017
39	659.15	0.04219	0.57816	64	392.84	0.13000	0.81074
40	671.71	0.04331	0.60220	65	379.67	0.13000	0.75627
41	685.54	0.04442	0.62672	66	366.50	0.13000	0.69837
42	695.84	0.04554	0.65198	67	353.34	0.13000	0.63806
43	700.60	0.04665	0.67823	68	340.17	0.13000	0.57896
44	705.51	0.04777	0.70570	69	327.00	0.13000	0.51493

[a] Tend income = m(t), and experience effect = kappa in Eq. (12), values in bold adjusted during calibration.
[b] Housing weight refers to age specific factor used to adjust housing costs.

disposable income at age 65 is not due to a coincident jump in the housing weights that are applied, as the age profiles for the housing weights were subject to kernel smoothing. The jumps that are evident in the geometric means reported in Figs. 2 and 3 are attributable to discontinuities in labour supply responses, an issue that is discussed at length in the following subsection.

Finally, the top panel of Fig. 4 indicates that the simulation model does a good job of capturing the age profiles described for the geometric mean of household consumption for both singles and couples. Consistent with the statistics reported for the dispersion of disposable incomes, the bottom panel of Fig. 4 indicates that the dispersion of household consumption by age and relationship status that is described by survey data is understated by the simulation model.

6.2. Sensitivity to principal preference parameters

6.2.1. Adjusting the intra-temporal elasticity, ε

As noted in Section 4.2, ε was adjusted to match the simulation model to rates of employment participation by wealth quintile observed in survey data between ages 60 and 64. Sensitivity of simulated employment participation to alternative values for ε, holding all other model parameters fixed, is reported in Table 7. Table 7 confirms the influence of ε on the timing of retirement: a higher value for the intra-temporal elasticity, ε, tends to reduce employment participation among households in the lowest wealth quintile, and *vice versa* for households in the top wealth quintile.

6.2.2. Adjusting the parameter of relative risk aversion, $1/\gamma$

Section 4.2 reports that increasing $1/\gamma$ for a given value of ε tends to make leisure and consumption direct substitutes, and that γ was consequently adjusted to reflect the impact of employment on consumption about retirement. Table 8 reports the sensitivity of simulated behaviour to variations in $1/\gamma$ alongside the associated sample moments upon which the calibration was based.

6.3. Factors complicating the model calibration

There are two aspects of the calibration that is reported in Section 6.1 that stand out as anomalies. The first, is the higher leisure costs that are imposed on employment from age 65, and the second is the substantial disparities reported between the simulated and sample geometric means of employment income. The terms of the preferred calibration that are reported in Section 6 were selected after rejecting an alternative calibration that imposed the same leisure cost of employment throughout the simulated lifetime, and in which the parameters of the wage generating processes were set exclusively to match simulated to sample moments for employment income. This 'restricted calibration', which is available from the author upon request, was found to obtain a very close match to moments of employment income, but to also generate insufficient savings among single adult households, and to obtain a poor fit to statistics for labour supply.

Regarding the first of these anomalies identified under the restricted calibration, the discount rate δ was the principal parameter used to match the model to consumption/saving profiles. As the same discount rate is assumed for both single adults and couples, and as the consumption profile generated for couples under the restricted calibration was found to obtain a decent match to survey data, no compensating adjustment could be found to improve the associated match for single adult households. Hence, the restricted calibration provides a poor reflection of the consumption/savings decisions with which the model is intimately concerned.

Regarding the second of the anomalies identified for the restricted calibration, rates of labour market participation generated by the model were found to respond abruptly to age specific changes in benefits policy that applied in 2005, which are not evident in the associated survey data. In 2005, any man aged 60 or over was eligible to a means-tested retirement benefit (the Guarantee Credit) that was subject to a 100% withdrawal rate — meaning that £1 was withdrawn for every £1 of private income. From age 65 in 2005, men in the UK became eligible to the basic State Pension and to a means-tested benefit (the Savings Credit) that reduced the effective withdrawal rate

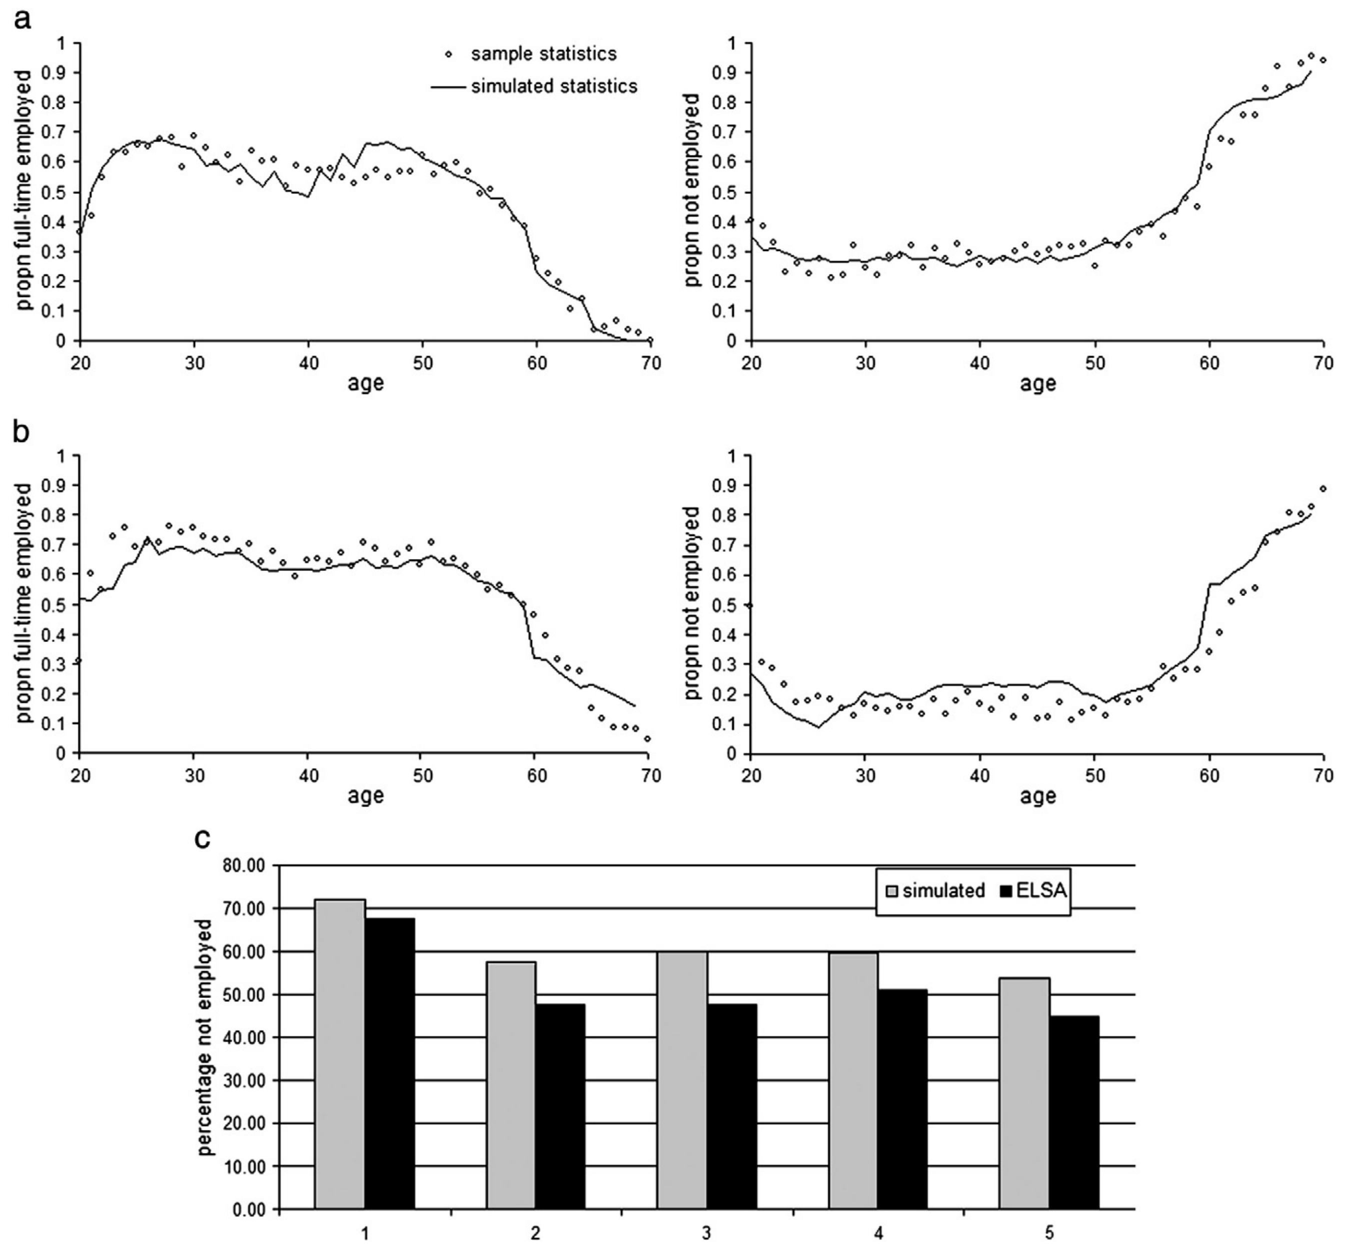

Fig. 1. Employment statistics — simulated versus survey data. (a) Single adults; (b) adult couples; and (c) percentage of wealth quintile not economically active, 60-64 age band. Notes: sample statistics — age profiles calculated from 2005/06 FRS data. ELSA statistics — percentage of males not economically active, reported in Marmot et al. (2003), p. 156. Simulated statistics — age profiles generated from model, using calibrated parameters reported in Tables 4 and 5.

on income over the basic State Pension from 100% to 40%.[22] The simulations suggest that the benefits structure available to men between ages 60 and 64 in 2005 represents a substantial disincentive to work for low wealth households, disincentives that are partially offset under the benefits structure that applied for men aged 65 and over. As a consequence of these incentives, the restricted model calibration generated a substantial fall in rates of employment participation from age 60, and a return to employment from age 65.

The anomalies that are identified above for the unrestricted calibration (a higher utility cost of leisure from state pension age, and poor fit to the geometric mean of employment income) are the consequence of an attempt to improve the match between simulated and sample moments of employment participation. On the one hand, higher labour incomes from age 55 help support higher rates of employment participation late in the working lifetime, and on the other a higher leisure cost of employment from age 65 suppresses labour market re-entry. These adjustments were made necessary by the form of the tax and benefits system that was assumed for the calibration.

Extensive experimentation failed to identify model parameters that could unambiguously reconcile the step-wise transitions in employment incentives described by the economic environment assumed for the model, and the smooth age profiles described by survey data. An obvious question, which remains unanswered at the time of writing, is what important aspects of the economic environment are omitted that are likely to be responsible for the mis-match that is identified here?

[22] Women were eligible to the Savings Credit from age 60. The model makes no distinction on sex, and assumes the ages of benefits eligibility that applied for men. It also assumes eligibility to the full basic State Pension, and ignores the (earnings related) State Second Pension.

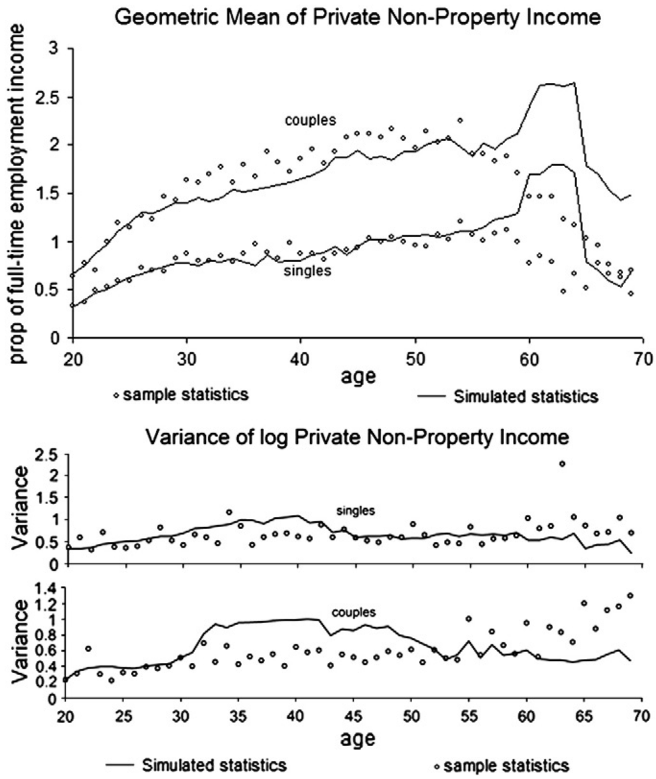

Fig. 2. . Private non-property income profiles by age — simulated versus sample moments. Notes: sample statistics — age profiles calculated from 2005/06 FRS data. Simulated statistics — age profiles generated from model, using calibrated parameters reported in Tables 4 and 5.

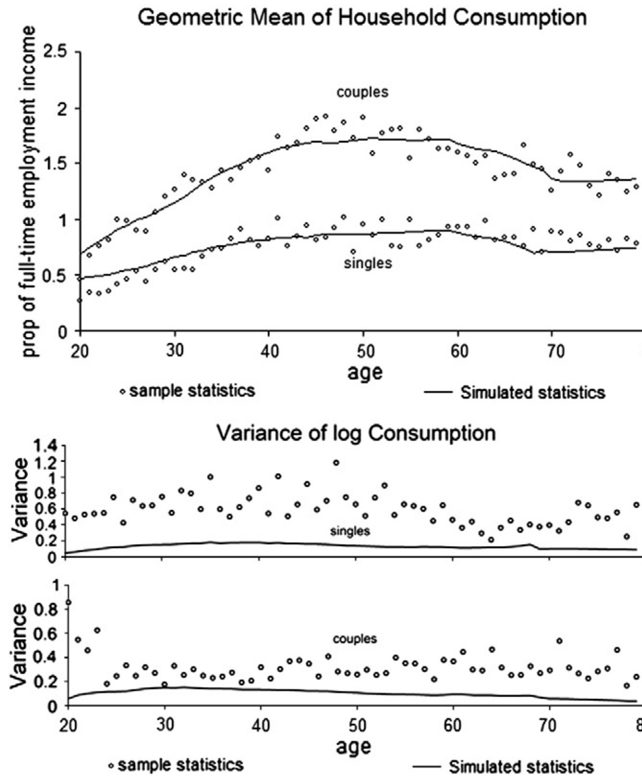

Fig. 4. Consumption profiles by age — simulated versus survey data. Notes: sample statistics — age profiles calculated from 2005 EFS data, and adjusted to meet budget balance over the simulated lifetime. Simulated statistics — age profiles generated from model, using calibrated parameters reported in Tables 4 and 5.

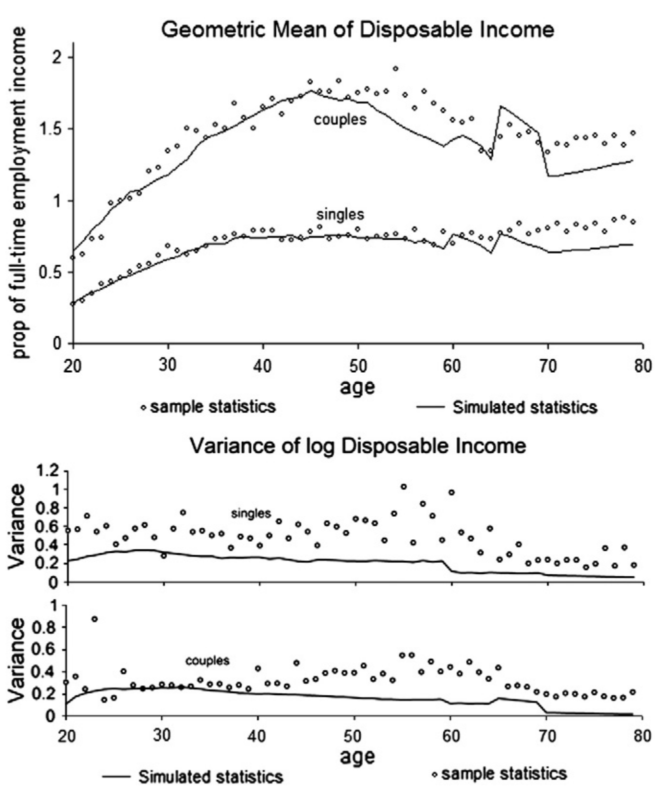

Fig. 3. Disposable income profiles by age — simulated versus sample moments. Notes: sample statistics — age profiles calculated from 2005/06 FRS data. Simulated statistics — age profiles generated from model, using calibrated parameters reported in Tables 4 and 5.

7. Conclusion

Policy makers are faced with a complex problem when evaluating the relative merits of reform alternatives, suggesting that substantial welfare gains may be achieved by improving the evidence base upon which such decisions are made. It is consequently of note that the life-cycle framework, which represents current best practice in the economic analysis of intertemporal decision making, has had a very limited bearing on the policy debate in contexts where the affected behavioural margins display a clear intertemporal dimension. Valid questions are then, what progress has been made towards adapting the life-cycle framework so that it provides a suitable basis to inform the policy debate, and when will such a framework be made accessible to policy makers? In answer to the first of these questions, the last fifty years has seen a great deal of preparatory work to identify suitable assumptions upon which to found an analysis of policy within the life-cycle framework — we are now able to explain much of the behaviour that is described by the available survey data, even if some conspicuous gaps remain. Furthermore, the rapid pace of progress in computing technology means that we are able to implement increasingly detailed model specifications, which is crucial in capturing the complexity that characterises real-world circumstances. On the second question, life-cycle models that incorporate the advances described above are now being adapted for policy makers, and will hopefully start to have an important bearing upon the policy debate within a decade or two. This paper describes one such model.

The model described here is designed to explore savings and labour supply responses to tax and benefits policy. The focus upon savings provides the fundamental motivation for founding the model upon the life-cycle framework, and the allowance that is made for an experience effect on wages makes a forward-looking framework of decision making important for labour supply decisions too. The

Table 7
Percentage of population aged 60–64 not economically active, by wealth quintile.

Intra-temporal elasticity	Population wealth quintile				
	Lowest	2	3	4	Highest
0.56 (calibrated value)	72.2	57.4	60.0	59.6	53.7
0.45	66.2	60.3	64.1	69.1	67.3
0.65	79.6	58.9	52.2	51.7	43.6
ELSA[a]	67.6	47.6	47.7	51.1	44.9

[a] Statistics for men reported by Marmot et al. (2003), p. 156. All other statistics reported for full simulated adult population.

Table 8
Least squares regression statistics for effects of employment on consumption between ages 60 and 64, simulated and sample data.

Relative risk aversion	Singles	Couples
2.9 (calibrated value)	0.4171	0.3763
1.5	0.2360	0.2684
4.5	0.5157	0.4375
EFS[a]	0.3816	0.3547
(std error)	(0.0867)	(0.0666)

[a] Author's calculations using 2005/06 wave of EFS all other statistics reported for simulated populations.

central concern with taxes and benefits makes the household, defined as a single adult or partner couple and their dependant children, the natural unit to adopt for analysis. It also motivates an explicit consideration of relationship status within the modelling framework, which is in contrast to many comparable models in the contemporary literature. The emphasis on saving and labour supply decisions makes pensions policy a central focus of concern, and particular care has been taken in endowing the model with sufficient flexibility to capture the terms of commonly observed pension schemes. As a consequence of this structure, the model permits a very close reflection to be obtained to contemporary tax and benefit systems that are applied in practice.

The form of the preference structure that motivates decisions within the model is standard in the related literature, and has much intuitive appeal (e.g. Browning and Lusardi (1996)). Consumption enters the preference relation in equivalised form, where the equivalence scale reflects the number of adults and children in a household. Households are considered to be uncertain about the timing of mortality but aware of age specific mortality rates. And wages are uncertain, subject to a preference structure that is sufficiently flexible to reflect a precautionary savings motive. These three considerations have been identified as important factors in reflecting the age profiles of consumption that are commonly displayed by survey data.

Furthermore, the preference relation assumed for the model does not impose additive separability between consumption and leisure, and the model has been structured to allow for age specific probabilities of involuntary unemployment. The literature suggests that these factors can be important in capturing the change in consumption on retirement. The model also allows for time-inconsistent preferences in the form of quasi-hyperbolic discounting, so that it is sufficiently flexible to capture demand for commitment mechanisms, where these are observed. And the model allows for both hard and soft liquidity constraints in the intertemporal budget set, in response to empirical observations that these can have an important bearing on behaviour in practice.[23]

This model structure is shown to be sufficiently flexible to obtain a close match to a wide range of sample moments estimated from contemporary survey data. Nevertheless, the calibration that is reported here does throw up an important puzzle: why do rates of employment participation described by contemporary survey data reveal reasonably smooth trends by age, when the underlying tax and benefits system implies substantial discontinuities in the (post-tax) returns to labour? This type of issue is one that macro-modellers are well acquainted with. One solution to this issue that has been explored in the macro-literature, and which we hope to address in the near future, is to allow for decision making rigidities such as habit formation. Indeed, it is hoped that by allowing for habit formation in the preference relation, we might obtain a model structure that provides an appropriate basis for exploring behavioural responses to policy shocks in the short to medium term. This is in contrast to the current model specification, which is best adapted to the identification of long-run behavioural responses, which provide qualitative detail of medium term incentive effects.

The rapid advance of computing technology shows no sign of abating, and the economics profession continues to search for, and adapt to discrepancies between our theoretical understanding of the decisions that people make and the associated statistical record. In view of these trends, and given how much progress has been made to date, it is reasonable to hope that the next decade or two will see analyses based formally upon the life-cycle framework appear as a routine feature of the policy reform process, in much the same way that classical micro-simulation analyses do today. To the extent that this improves the information base used to evaluate alternatives for policy reform, it is likely to deliver substantial welfare gains — a milestone that will indicate the real coming of age of the micro-economic analysis of intertemporal behaviour.

[23] 'Hard' constraints are those that impose a fixed upper limit to the available credit, and 'soft' constraints impose higher interest rates as the (negative) wealth position deteriorates.

Acknowledgement

The model described here is based upon programming architecture initially devised by James Sefton. Routines for the Gauss–Hermite Quadrature, numerical solutions of (local) optima, and polynomial extrapolation were obtained from Press et al. (1986).

The model also makes use of the Bus and Dekker (1975) zero finding algorithm. Interpolation methods are based upon Keys (1981). The model is written in Fortran, with parallelisations implemented using OpenMP. I should also like to thank the following institutions for valuable funding: the Economic and Social Research Council, the Department for Work and Pensions, Her Majesty's Revenue and Customs, the Department for Business, Enterprise and Regulatory Reform, the National Association of Pension Funds, the Leverhulme Trust, and the European Commission. Many thanks to Martin Weale and an anonymous referee for comments on an earlier draft. The usual disclaimer applies.

References

Agarwal, S., Liu, C., Souleles, N.S., 2007. The reaction of consumer spending and debt to tax rebates — evidence from consumer credit data. Journal of Political Economy 115, 986–1019.

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

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

Attanasio, O.P., Weber, G., 2010. Consumption and saving: models of intertemporal allocation and their implications for public policy. Journal of Economic Literature 48, 693–751.

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.

Banks, J., Blundell, R., Tanner, S., 1998. Is there a retirement-savings puzzle? The American Economic Review 769–788.

Berneim, B.D., Skinner, J., Weinberg, S., 2001. What accounts for the variation in retirement wealth among U.S. households? The American Economic Review 91, 832–857.

Beshears, J., Choi, J.J., Laibson, D., Madrian, B.C., 2008. The importance of default options for retirement saving outcomes: evidence from the United States. In: Kay, S.J., Sinha, T. (Eds.), Lessons from Pension Reform in the Americas. Oxford University Press, Oxford.

Browning, M., Ejrnæs, M., 2002. Consumption and children. Centre for Applied Microeconometrics Discussion Paper 2002-06.
Browning, M., Lusardi, A., 1996. Household saving: micro theories and macro facts. Journal of Economic Literature 34, 1797–1855.
Bus, J.C.P., Dekker, T., 1975. Two efficient algorithms with guaranteed convergence for finding a zero of a function. ACM Transactions on Mathematical Software 1, 330–345.
Caldwell, S., 1997. Corsim 3.0 User and Technical Documentation. Ithaca, New York.
Campbell, J.Y., Cochrane, J.H., 1999. By force of habit: a consumption-based explanation of aggregate stock market behavior. Journal of Political Economy 107, 205–251.
Campbell, J.Y., Deaton, A., 1989. Why is consumption so smooth? Review of Economic Studies 56, 357–373.
Campbell, J.Y., Mankiw, N.G., 1989. Consumption, income, and interest rates: reinterpreting the time series evidence. In: Blanchard, O.J., Fischer, S. (Eds.), NBER Macroeconomics Annual 1989. MIT Press, Cambridge, Mass.
Campbell, J.Y., Mankiw, N.G., 1991. The response of consumption to income: a cross country investigation. European Economic Review 35, 723–756.
Carroll, C., 1992. The buffer-stock theory of saving: some macroeconomic evidence. Brookings Papers on Economic Activity 23, 61–156.
Carroll, C., Summers, L., 1991. Consumption growth parallels income growth: some new evidence. In: Bernheim, B.D., Shoven, J. (Eds.), National Saving and Economic Performance. University of Chicago Press, Chicago.
Choi, J.J., Laibson, D., Madrian, B.C., Metrick, A., 2002. Defined contribution pensions: plan rules, participation choices, and the path of least resistance. Tax Policy and the Economy 16, 67–113.
Connolly, S., Gregory, M., 2008. Moving down: women's part-time work and occupational change in Britain 1991–2001. Economic Journal 118, F52–F76.
Curry, C., 1996. PENSIM: a dynamic simulation model of pensioners' income. Government Economic Service Working Paper 129.
Deaton, A., 1997. The Analysis of Household Surveys. John Hopkins University Press, Baltimore.
Deaton, A., Muellbauer, J., 1980. Economics and Consumer Behavior. Cambridge University Press, Cambridge.
Diamond, P., Köszegi, B., 2003. Quasi-hyperbolic discounting and retirement. Journal of Public Economics 87, 1839–1872.
Dynan, K., 2000. Habit formation in consumer preferences: evidence from panel data. The American Economic Review 90, 391–406.
Fagan, C., 2003. Working-time Preferences and Work-life Balance in the EU: Some Policy Considerations for Enhancing the Quality of Life. European Foundation for the Improvement of Living and Working Conditions, Dublin.
Flavin, M., 1981. The adjustment of consumption to changing expectations about future income. Journal of Political Economy 89, 974–1009.
French, E., 2005. The effects of health, wealth, and wages on labour supply and retirement behaviour. Review of Economic Studies 72, 395–427.
Friedman, M., 1957. A Theory of the Consumption Function. Princeton University Press, Princeton.
Fuhrer, J.C., 2000. An optimizing model for monetary policy: can habit formation help? The American Economic Review 90, 367–390.
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.
Gross, D.B., Souleles, N.S., 2002. Do liquidity constraints and interest rates matter for consumer behaviour? Evidence from credit card data. Quarterly Journal of Economics 16, 149–185.
Gustman, A.L., Steinmeier, T.L., 2005. The social security early entitlement age in a structural model of retirement and wealth. Journal of Public Economics 89, 441–464.
Hamermesh, D.S., 1984. Consumption during retirement: the missing link in the life cycle. The Review of Economics and Statistics 66, 1–7.
Heckman, J., 1974. Life cycle consumption and labor supply: an explanation of the relationship between income and consumption over the life cycle. The American Economic Review 64, 188–194.
Hicks, J.R., 1939. Value and Capital. Clarendon Press, Oxford.
Juster, F.T., Shay, R.P., 1964. Consumer sensitivity to finance rates: an empirical and analytical investigation. NBER Occasional Paper 88.
Keynes, J.M., 1936. The General Theory of Employment, Interest, and Money. MacMillan, London.
Keys, R.G., 1981. Cubic convolution interpolation for digital image processing. IEEE Transactions on Acoustics, Speech, and Signal Processing 29, 1153–1160.
Khan, A., Thomas, J.K., 2008. Adjustment costs. In: Durlauf, S.N., Blume, L.E. (Eds.), The New Palgrave Dictionary of Economics. Palgrave Macmillan, Basingstoke.
Laibson, D., Repetto, A., Tobacman, J., 2007. Estimating discount functions with consumption choices over the lifecycle. Oxford University Department of Economics Discussion Paper 341.
Lee, B.S., Ingram, B., 1991. Simulation estimation of time-series models. Journal of Econometrics 47, 197–205.
Low, H., 2005. Self-insurance in a life-cycle model of labour supply and savings. Review of Economic Dynamics 8, 945–975.
Ludvigson, S.C., Michaelides, A., 2001. Does buffer-stock saving explain the smoothness and excess sensitivity of consumption? The American Economic Review 91, 631–647.
Madrian, B., Shea, D., 2001. The power of suggestion: inertia in 401(k) participation and savings behaviour. Quarterly Journal of Economics 141, 1149–1187.
Marmot, M., Banks, J., Blundell, R., Lessof, C., Nazroo, J., 2003. Health, Wealth and Lifestyles of the Older Population in England: The 2002 English Longitudinal Study of Ageing. Institute for Fiscal Studies, London.
Mehra, R., Prescott, E.C., 1985. The equity premium: a puzzle. Journal of Monetary Economics 15, 145–161.
Modigliani, F., Brumberg, R., 1955. Utility analysis and the consumption function: an interpretation of cross-section data. In: Kurihara, K.K. (Ed.), Post Kerynesian Economics. Rutgers University Press, New Brunswick, NJ.
Nelissen, J.H.M., 1998. Annual versus lifetime income redistribution by social security. Journal of Public Economics 68, 223–249.
O'Donoghue, T., Rabin, M., 1999. Procrastination in preparing for retirement. In: Aaron, H.J. (Ed.), Behavioral Dimensions of Retirement Economics. Brookings Institute, Washington.
Pissarides, C.A., 1978. Liquidity considerations in the theory of consumption. Quarterly Journal of Economics 92, 279–296.
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.
Shaw, K.L., 1989. Life-cycle labor supply with human capital accumulation. International Economic Review 30, 431–456.
Smith, S., 2006. The retirement-consumption puzzle and involuntary early retirement: evidence from the British Household Panel Survey. Economic Journal 116, C130–C148.
Sommer, M., 2001. Habits, sentiment, and predictable income in the dynamics of aggregate consumption. Johns Hopkins University Department of Economics Working Paper 458.
Stern, S., 1997. Approximate solutions to stochastic dynamic programs. Econometric Theory 13, 392–405.
Thaler, R.H., Benartzi, S., 2004. Save more tomorrow: using behavioural economics to increase employee saving. Journal of Political Economy 112, S164–S187.
Tobin, J., 1967. Life cycle saving and balanced growth. In: Fellner, W., Hall, C.A., Koopmans, T.C., Miller, J.P., Nerlove, M., Ruggles, R., Samuelson, P.A., Scarf, H., Tobin, J., Wallich, H.C. (Eds.), Ten Economic Studies in the Tradition of Irving Fisher. Wiley, New York.
van de Ven, J., 2010. An empirical investigation of quasi-hyperbolic discounting. National Institute of Economic and Social Research Discussion Paper 355.
Zeldes, S.P., 1989. Consumption and liquidity constraints: an empirical investigation. Journal of Political Economy 97, 305–346.