

A Framework for Pension Policy Analysis in Ireland: PENMOD, a Dynamic Simulation Model*

Tim Callan[†], Justin van de Ven[‡] & Claire Keane[§]

1. INTRODUCTION

Public policy towards both private pensions and state-provided pensions must be framed in a long-term context. Decisions regarding participation in private schemes, and the extent of contributions thereto, have implications which unfold over time. In defined contribution (DC) schemes, an individual's pension fund is built up over the working lifetime, and then drawn down in retirement. Government's budget constraint also leads to trade-offs between the level of the state pension, the age at which it becomes payable, and the taxes required to finance it. Because of the essential dynamic elements in pension contributions and payments, the impact of policy changes is not well captured by static models, which take a "snapshot" of the impact at a point in time. While such models (including SWITCH, the ESRI tax-benefit model) can provide some insights into the impact of pension-related policies, a fuller analysis must take account of the complex interplay of forces over time.

The approach taken here is well established in the economic literature on pensions. Essentially, our model (PENMOD) takes a representative cohort of individuals and simulates key elements of their lifetime experience. This includes both economic elements such as labour market participation and wages as they age as well as demographic elements (marriage, divorce, children, death). Crucially, decisions

* The model described here is based upon the NIBAX model architecture, as described in van de Ven (2011). We are grateful to Gerry Hughes and to Pete Lunn for helping to plug gaps in our knowledge. We are also indebted to several members of CSO staff for assistance: Tom McMahon, Pamela Lafferty, Marion McCann, Deirdre Cullen, and Shaun McLaughlin. Responsibility for any errors or obscurities rests with the authors.

[†] Economic and Social Research Institute, Dublin. Email: tim.callan@esri.ie

[‡] National Institute of Economic and Social Research, London. Email: jvanden@niesr.ac.uk

[§] Economic and Social Research Institute, Dublin. Email: claire.keane@esri.ie

regarding savings and pensions are also taken into account. Policy instruments in terms of income tax and social welfare are also included.

The need to take into account a sequence of decisions over each individual's full lifetime (up to the age of 120) imposes a very strict discipline on the degree of detail that can be incorporated into the model. Static models (such as SWITCH) can include a very high degree of detail in their description of the tax and welfare system. Dynamic models (such as PENMOD) must use a broader brush, in order to be able to provide greater depth in terms of the analysis over time. Thus, it is not a case of one class of model being "better" than another; rather it is a question of different classes of model being more suitable for different purposes.

A number of strategic simplifications are needed to ensure that the dynamic microsimulation model captures key features of the tax/welfare and pension systems while remaining tractable. One major simplification is that the model does not attempt to deal with public sector pensions, where the issues which arise are of a different type. We focus instead on the private sector, where decisions regarding the balance between contributions towards pension savings and the income in retirement are more subject to the influence of economic and policy variables. Second, we focus on private sector employees rather than the self-employed. This is because the terms of retirement for the self-employed often depend upon the envisaged income arising from ownership of a family business, or revenues arising from its sale that are distinct from the pension system with which we are immediately concerned. Third, we do not attempt to deal with issues of illness and longer-term incapacity to work. There are both state schemes (Illness Benefit, Invalidity Pension, Disability Allowance) and private schemes (permanent health insurance) which are geared towards dealing with income support for those unable to work. The issues arising are, however, too complex to include when modelling the long-term evolution of incomes and pensions and are, therefore, outside the scope of the present model. Simplifications of this type are common in the international literature in this area.

As regards the pension regime itself, this is characterised by up to five different types of pension scheme running in parallel. Each scheme takes a defined contribution form, where the approach adopted is to allow for schemes of differing "quality", with the probability of obtaining higher quality pensions rising with income – details of the approach are set out in Section 2.3.

The remainder of the paper is set out in 7 sections as follows. 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. Section 3 provides technical details of how the model generates behavioural responses to policy change. The approach

taken to calibrate the model against Irish survey data is described in Section 4, and calibrated model parameters are reported in Sections 5 and 6. A brief example of the type of policy analysis that can be conducted using the model is reported in Section 7, and a summary and directions for further research are provided in the conclusion.

2. MODEL SPECIFICS

The decision unit in the model is the nuclear family unit, defined as a single adult or partner couple and their dependent children.¹ 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 simulated on the assumption that households maximise expected lifetime utility, given their prevailing circumstances, 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, pension opportunities, private sector pension rights, and time of death. The belief structure is rational, in the sense that expectations are calculated on probability distributions that are consistent with the intertemporal decision making environment.

Of the eight characteristics that define the circumstances of a household, seven can be considered stochastic (relationship status, number of children, private sector pension scheme eligibility, private sector pension rights, wage rates, liquid wealth, and time of death), and only age is forced to be deterministic.

As a brief overview, the model permits:

- the adjustment of preferences over consumption, leisure, and bequests
- adjustment of the imposed liquidity constraints, which are defined both in terms of hard credit limits and variable interest charges that depend on the debt to income ratios
- inclusion of uncertainty over relationship status (single or couple)
 - provided that relationship status is considered to be uncertain, the number of children in a household can also be modelled stochastically
- alternative options in regard to the nature of uncertainty associated with labour incomes, including the possibility of receiving a low (zero) wage offer
- households to invest some of their liquid wealth in a risky asset

¹ For convenience, we use the term “household” interchangeably with “narrow nuclear family unit”, which has the advantage of brevity at the cost of a slight abuse of language. It should nevertheless be understood that adult children are treated as independent units in our analysis.

- the nature of the uncertainty associated with returns to the risky asset can also be altered
- households to choose their labour supply between discrete alternatives
- adjustment of a detailed tax and benefits structure
- private sector pensions
 - contribution rates (and ultimately membership) can also be made a decision variable
 - contribution rates (employee and employer) can be made stochastic
 - the stochastic nature of the return to private pension wealth can be adjusted

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 size.

2.1. The utility function

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

$$\begin{aligned}
 U_{i,t} = & \frac{1}{1 - 1/\gamma} \left\{ u \left(\frac{c_{i,t}}{\theta_{i,t}}, l_{i,t} \right)^{1-1/\gamma} + \right. \\
 & + E_t \left[\beta_1 \delta \left(\varphi_{1,t} u \left(\frac{c_{i,t+1}}{\theta_{i,t+1}}, l_{i,t+1} \right)^{1-1/\gamma} + (1 - \varphi_{1,t}) \left(\zeta_a + \zeta_b w_{i,t+1}^+ \right)^{1-1/\gamma} \right) + \right. \\
 & \left. \left. + \beta_1 \beta_2 \sum_{j=t+2}^T \delta^{j-t} \left(\varphi_{j-t,t} u \left(\frac{c_{i,j}}{\theta_{i,j}}, l_{i,j} \right)^{1-1/\gamma} + (1 - \varphi_{j-t,t}) \left(\zeta_a + \zeta_b w_{i,j}^+ \right)^{1-1/\gamma} \right) \right] \right\} \quad (1)
 \end{aligned}$$

where $1/\gamma > 0$ is the (constant) coefficient of relative risk aversion; 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); $\varphi_{j-t,t}$ 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 considered to be 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. When adults are modelled explicitly, then households with one adult can choose from up to three labour options; full-time $(l_{i,t}^{FT})$, part-time $(l_{i,t}^{PT})$,

and not employed ($I_{i,t} = 1$). Similarly, couples can choose from up to five labour options; both full-time employed ($I_{i,t}^{2FT}$), one full-time and one part-time employed ($I_{i,t}^{FtPt}$), one full-time and the other not employed ($I_{i,t}^{FtNe}$), one part-time and the other not employed ($I_{i,t}^{PtNe}$), and both not employed ($I_{i,t} = 1$). 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 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 incomes before housing costs in European countries. Its inclusion in the preference relation reflects the fact that household size has been found to have an important influence on the timing of consumption (e.g. Attanasio & Weber (1995) and Blundell et al. (1994)).²

The model incorporates an allowance for behavioural myopia, through its assumption of quasi-hyperbolic preferences following Laibson (1997). Such preferences are interesting because they are time inconsistent, giving rise to the potential for “conflict between the preferences of different intertemporal selves” (Diamond & 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 & Steinmeier (2005), who allow variation in the rate of time preference to be an important factor in reflecting heterogeneity in household retirement behaviour. We have chosen not to do this to ensure that heterogeneity of household behaviour generated by the model is driven by heterogeneity in observable household characteristics.

² An empirical study by Fernandez-Villaverde & Krueger (2006) of US data from the Consumer Expenditure Survey suggests that roughly half of the variation observed for lifetime household consumption can be explained by changes in household size, as described by equivalence scales. See Balcer & Sadka (1986) and Muellbauer & van de Ven (2004) on the use of this form of adjustment for household size in the utility function.

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,

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

where $\epsilon > 0$ 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 equations (1) and (2) is standard in the literature, despite the contention that is associated with the assumption of time separability (see Deaton & 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

Equation (1) is considered to be 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 The age profile of D_t can either be exogenously defined in the model, or be relaxed subject to the constraint that all households must have repaid their debts by an exogenously defined age, $t_D \leq T$ (the maximum terminal age assumed for the model).⁵ 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:

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.

⁵ Note that the structure of the decision problem considered here implies that relaxing the upper limit on debt does not permit households to consume an infinite amount prior to the age by which all debts are forced to be repaid. In the context of uncertainty, and when marginal utility approaches infinity as (discretionary) consumption tends toward zero, relaxing the constraint on debt implies an upper bound on consumption that is defined in terms of the minimum potential income stream that a household may receive in all future periods up to the date by which all debts must be repaid.

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

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

$$\tau_{i,t} = \tau(l_{i,t}, x_{i,t}, n_{i,t}^a, n_{i,t}^c, r_{i,t}^s, w_{i,t}^s, r_{i,t}^r, w_{i,t}^r, pc_{i,t}, t) \quad (3c)$$

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

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 pensionable 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 $\tau(\cdot)$ 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 module that simulates the effects of taxes and benefits on household disposable incomes. Equation (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, $r_{i,t}^r w_{i,t}^r$ (possibly negative); contributions to private sector pensions, $pc_{i,t}$; and age, t .

The form of the budget constraint described by equation (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 wealth held 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.

We now describe details of the function that is used to evaluate disposable income. The lifetime is divided into two periods for the purpose of calculating disposable

income: the working lifetime $t < t_{SPA}$, and pension receipt $t_{SPA} \leq t$. In each of these periods of life, household disposable income is calculated by:

1. evaluating aggregate *take-home pay* from the taxable incomes of each adult member of a household – this reflects the taxation of individual incomes in the Ireland
2. simulating receipt of benefits from aggregate household take-home pay – this reflects the fact that benefits tend to be provided at the level of the family unit
3. household *disposable income* is then equal to aggregate take-home pay, plus benefits.

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 pensionable age, $t < t_{SPA}$, household non-property income $x_{i,t}$ considered for tax purposes is equal to labour income $g_{i,t}$ less the proportion of pension contributions that is considered tax exempt, π^{pe} ; from state pensionable age it is equal to labour income plus the proportion of pension annuity income that is considered taxable, π^{pt} :

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

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

$p_{i,t}$ denotes pension annuity income, and χ is the annuity rate considered for analysis. 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 pensionable 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 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 an exogenously 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 pensionable 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:

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

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 is assumed to depend upon whether $w_{i,t}^s$ indicates net investment assets, or net debts:

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

where $l_{i,t}^t$ 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 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:

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

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

where $tax_{i,t}$ denotes the simulated tax burden, and $benefits_{i,t}$ welfare benefits received.

2.2.1. Intertemporal indexing

It is likely that individuals take some account of wage growth when planning for the future: a 20 year old today can reasonably expect that labour incomes will be higher when they reach age 45 than are currently paid to today's 45 year olds. If this is true, then it is important that the rational agent model be calibrated against data that take wage growth into account (discussed at further length in Section 4). This gives rise to a host of complications regarding the appropriate intertemporal development to assume for the tax and benefits system: holding taxes and benefits fixed in the context of rising wages, for example, will result in widespread tax bracket creep and marginalisation of the welfare state, with important implications for simulated behaviour.

Two parameters of the model control the way in which the tax system evolve with time in the model. The first controls the rate at which tax thresholds grow with time, thereby offsetting bracket creep, and the second controls the rate of growth of welfare benefits. These parameters adjust the tax and benefits schedules in a way that is designed to omit the creation of poverty traps. Nevertheless, rapid temporal adjustment of the tax system can give rise to analytical problems, and the the model is programmed in a way that is designed to indicate when excessive variation has been called for. Separate routines have been developed that allow the disposable income schedules that are generated by the model to be viewed directly, and these are reviewed to verify that a model simulation is sensible.

2.3. Private Sector Pensions

Private sector pensions in the model 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. Although DC pensions account for less than half of all pensions that currently attract contributions in Ireland, there has been a strong temporal trend toward DC schemes since the 1990s (in common with countries throughout the OECD), which motivates our modelling in this regard. Up to five private sector pension schemes can be considered in parallel in the model, where schemes are distinguished by their respective rates of (exogenously defined) employer contributions. Pension contribution rates are defined as percentages of (total) labour income, implying that pension membership requires employment participation. Households are considered to be eligible to participate in only one pension scheme in any year, where eligibility to each scheme is identified stochastically with reference to a set of income dependent probabilities, and uncertainty between adjacent years can be suppressed in cases of continuous pension participation. Membership of a pension to which a household is eligible can either be exogenously imposed, or modelled as an endogenous decision. Similarly, the rate of private contributions to a pension scheme, $\pi_{i,t}^p$, can either

be exogenously imposed, or considered endogenous in the model. Where private pension contributions are considered endogenous in the model, then these 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 .

Accrued rights to a private pension are described by:

$$w_{i,t}^p = \begin{cases} (1 + r_{t-1}^p) w_{i,t-1}^p + (\pi_{i,t}^p + \pi_{ec,j}^p) (g_{i,t} - \pi_l^p) & \text{member of scheme } j \\ (1 + r_{t-1}^p) w_{i,t-1}^p & \text{otherwise} \end{cases} \quad (9a)$$

$$\ln(1 + r_t^p) \sim N\left(\mu_p - \frac{\sigma_p^2}{2}, \sigma_p^2\right) \quad (9b)$$

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 $w_{i,t}$. 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. 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

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

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, as the current state of computing technology makes it impractical to address them, and to analyse endogenous decisions over pension contributions.

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, $\log(h_{i,t_0}) \sim N(\mu_{n^a,t_0}, \sigma_{n^a,t_0}^2)$, where the parameters of the distribution depend upon the number of adults in the household, n^a . Thereafter, latent wages follow a random walk with drift described by the equation:

$$\log\left(\frac{h_{i,t}}{m(n_{i,t}^a, t)}\right) = \log\left(\frac{h_{i,t-1}}{m(n_{i,t-1}^a, t-1)}\right) + \kappa(n_{i,t-1}^a, t-1) \frac{(1-l_{i,t-1})}{(1-l_w)} + \omega_{i,t} \quad (10)$$

where the parameters $m(\cdot)$ account for wage growth (and depend on age, t , and the number of adults in the household, $n_{i,t}^a$), $\kappa(\cdot)$ 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 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 well established⁷, its inclusion is an innovation for the related literature (e.g. Low, 2005, and French, 2005). 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 lifecycle.

⁷ With regard to statistical evidence of the effect of experience on income, Mincer & Ofek (1982) report that in the short run, every year out of the labour market can result in a 3.3%-7% fall in wages relative to those who remain employed. This study also finds, however, that the restoration of human capital tends to be faster than the original accumulation, so that the impact of early labour breaks reduce to 1.3%-1.8% in the long run. Eckstein & Wolpin (1989) do not make a distinction between the long run and short run impact of actual experience, but find that the first year out of the labour market reduces wages by around 2.5%, with subsequent years having a marginally diminishing effect. See also, Waldfogel (1998) and Myck & Paull (2004) for the role of experience in explaining the gender wage gap.

2.4.1. Complicating the standard decision making problem

The preferences defined by equations (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 – equation (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 relative risk aversion (and, similarly, of γ as a measure of the intertemporal elasticity of substitution of total expenditure) carries over.⁸

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 relative risk aversion (and, similarly, γ as a measure of the intertemporal elasticity of substitution of total expenditure). Meanwhile, the experience effect on wages implies that the price of leisure is endogenous to the decision making problem, thereby invalidating two stage budgeting. Furthermore, a positive experience effect on wages tends to depress savings rates as wealth rises.⁹

2.5. Household composition

The model allows for households to form and to split, for the arrival of children, and for the risk of death at different ages. The technical approach in terms of numbers of adults and children in a household is to allow these to evolve stochastically, following a “reduced form” nested logit model. The first (highest) level determines the number of adults in a household, and the second (“nested” within that) determines the number of children, given the age and number of adults in the household.

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 stochastic variation

⁸ There is the separate issue of disentangling the intertemporal elasticity of substitution from relative risk aversion, which is not addressed here. See Epstein & Zin (1989).

⁹ To see this, note that an experience effect on wages tends to increase the present-discounted cost of reduced labour supply, to the extent that an individual expects to want to work in the future. As wealth rises, labour attachment is weakened, which also weakens the experience effect on incentives to work in the short run. Including an experience effect on wages consequently tends to exaggerate the negative relationship between wealth and labour supply, thereby depressing the savings rate as wealth rises.

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 rational agent model. This substantially increases the computational burden. If, for example, a household was considered to be 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 state variables 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.

This may seem somewhat artificial (it is as if larger families must involve multiple births, and births only occur at two specific ages). The precise timing of births is not a central focus of interest, however, and the approach taken here means that the presence and number of children can be taken into account, while abstracting somewhat from the associated detail.

The logit model that is considered to describe the evolution of adults in a household is given by equation (11):¹⁰

$$s_{i,t+1} = \alpha_0^A + \alpha_1^A t + \alpha_2^A t^2 + \alpha_3^A t^3 + \alpha_4^A dk_{i,t} + \alpha_5^A s_{i,t} \quad (11)$$

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 childbirth ages. The ordered logit equations assumed for the first childbirth age, for both singles and couples, do not include any additional household characteristics. The ordered logit equations for the second childbirth age includes the number of children born at the first childbirth age as an additional descriptive characteristic.

3. SOLVING THE LIFE-TIME DECISION PROBLEM

This section begins by discussing the conceptual approach adopted to solve the lifetime decision making problem, before describing details of the analytical routines used to implement the numerical solution.

¹⁰ When children are not modelled explicitly, then the cubic term in age and the dummy variable for children is omitted from the logit equation.

3.1. Conceptual approach

The procedures that we adopt use backward induction to solve for decisions that maximise expected lifetime utility. A terminal age T is assumed, following which death occurs with certainty. Utility maximising decisions at this terminal age are free of temporal dynamics, and are consequently straightforward 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. We refer to the utility associated with this solution as the value function, $V_T(n_T^a, w_T, p_T)$. Furthermore, we can calculate the intermediate measures of welfare:

$$\widehat{u}(n_T^a, w_T, p_T) = u\left(\frac{\widehat{c}_T(n_T^a, w_T, p_T)}{\theta_T}, 1\right) \quad (12)$$

$$\widehat{X}(n_T^a, w_T, p_T) = E_t\left(\frac{1}{(1-1/\gamma)} (\zeta_a + \zeta_b \widehat{w}_{T+1}^+(n_T^a, w_T, p_T))^{1-1/\gamma}\right) \quad (13)$$

where \widehat{c}_T and \widehat{w}_{T+1} denote the optimised measures of consumption and next period wealth, on the assumption that labour supply at the terminal age is not possible. We calculate these functions at all nodes of a three dimensional grid in the number of adults, wealth, and retirement annuity.

At age $T-1$, suppose that households are permitted to invest in risky assets and to supply labour. Here, the problem reduces to solving the Bellman equation:

$$\begin{aligned} V_{T-1}(n_{T-1}^a, w_{T-1}, h_{T-1}, wo_{T-1}, p_{T-1}) = \\ = \max_{c_{T-1}, v_{T-1}, l_{T-1}} \left\{ \frac{1}{1-1/\gamma} u\left(\frac{c_{T-1}}{\theta_{T-1}}, l_{T-1}\right)^{1-1/\gamma} + \right. \\ \left. + E_{T-1} \left[\frac{\beta_1 \delta}{1-1/\gamma} \left(\varphi_{1,T-1} \widehat{u}(n_T^a, w_T, p_T)^{1-1/\gamma} + (1 - \varphi_{1,T-1}) (\zeta_a + \zeta_b w_T^+)^{1-1/\gamma} \right) + \right. \right. \\ \left. \left. + \beta_1 \beta_2 \delta^2 \varphi_{1,T-1} \widehat{X}(n_T^a, w_T, p_T) \right] \right\} \quad (14) \end{aligned}$$

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 v_{T-1} is the proportion of liquid wealth invested in the risky asset. We solve this optimisation problem for the $T-1$ value function, at each node of the five dimensional grid over the permissible 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, 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 recursive repetition of the procedure outlined for age $T - 1$, given the solutions (previously) 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 pension assets permitted in the model. A more complete description of the analytical problem, including the treatment of boundary conditions, is reported in the technical appendix.

The above procedure generates a grid that spans all possible combinations of characteristics that the model considers a household might have (the state space). The utility maximising decisions identified by the numerical procedure are stored at each grid intersection, alongside the numerical approximation of expected lifetime utility (the value function). Although this set of information can be informative in its own right, most analyses are based upon panel data for the life-course of a cohort of households that are generated using the grid defined above. The life course of a birth cohort is generated 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 are considered to 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 to produce data for the entire life-course.

3.2. Details of solution routines

The model described here is complex and generates behaviour where no analytical solution exists. As such, it is reasonable to describe it as a 'black-box' routine, which raises concerns over the accuracy of the behavioural responses that it generates. These concerns are exaggerated by the fact that the value function may be both non-smooth and/or non-concave (although it is designed to be increasing and continuous), which can complicate the solution due to the existence of multiple local maxima.

It is important to recognise from the outset that *any* numerical solution is likely to be associated with a degree of error – the problem is to assess whether the scale of the inaccuracies generated by the model are qualitatively important for the purpose to which it is applied. The model includes *three* principal tools for assessing the accuracy of the numerical solutions that it derives: variation of solution detail, variation of interpolation methods, and variation of the numerical search routines that are used. The first is the most simple, and often the most powerful of the three.

When *varying the solution detail*, the size and number of grid points adopted for each of the continuous state variables can be altered, as can the number of abscissae used in the Gaussian quadrature.¹¹ 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; a problem that is commonly referred to as the curse of dimensionality.

The model includes both linear and cubic *interpolation methods*, for evaluating behaviour between discrete grid points.¹² Relative to linear interpolation, cubic interpolation produces a smoother functional form, and ensures continuous differentiability. Cubic interpolation also requires evaluations at 4^n grid points, rather than 2^n points, where n is the number of dimensions over which the interpolation is being taken. If the user indicates that cubic interpolation is to be used, then the model performs an internal check to determine whether the surface over which an interpolation is being conducted is reasonably smooth, before selecting the cubic interpolation for analysis; otherwise, it selects the linear interpolation.¹³ It is of note that the cubic interpolation, and linear interpolation routines have been programmed separately, and so can be used to validate against one another.

Finally, the model includes three alternative *numerical search routines*, which are used to find utility maximising values of consumption. A ‘brute force’ procedure uses grid search methods to identify the local optimum associated with the highest numerical approximation of the value function. The advantage of this approach is that it makes no assumptions regarding the form that the value function takes. This advantage is, however, purchased at a very substantial increase in the computational burden associated with the search routine. Alternatively, Brent’s method can be used to search over the consumption domain, based upon parabolic interpolation with a golden section search of repeated evaluations of the value function. This approach has been found to be efficient, particularly where the surface over which the search is conducted is reasonably well behaved, but is not designed to take account of multiple local optima. The third search alternative is based upon the Bus & Dekker (1975) bisection algorithm, which can be used to identify the consumption that evaluates the Euler condition to zero. Like Brent’s method, the Bus & Dekker (1975) algorithm is recognised as efficient, and is not designed to account for

¹¹ Evaluation of weights and abscissae of the Gauss-Hermite quadrature are based upon a routine reported in Chapter 4 of Press et al. (1986).

¹² The interpolation routines that are used are based on Keys (1981).

¹³ This 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.

multiple local optima. Relative to Brent's method, optimisation of the Euler condition can – in some circumstances – result in improved accuracy, but at the cost of increased computational burden (as repeated calls to the value function do not require the additional computational burden involved in evaluating first derivatives). Furthermore, some analytical contexts may argue against the use of Euler conditions, as in the case where non-exponential discounting is assumed.

A *supplementary search routine* is included in the model to mitigate concerns regarding identification of multiple local optima where Brent's method or the Bus & Dekker algorithm are applied. Here the model can be directed to explore a localised grid above and below an identified optimum for a preferred level of consumption, based upon value function calls. If an alternative value of consumption is identified by this supplementary routine as strictly preferred to the original local maximum, then the routine will search recursively for any further solutions above and below. This process is repeated until no further solutions are found. Of all feasible solutions, the one that maximises the value function is selected.

4. DATA AND CALIBRATION METHODOLOGY

4.1. Data considered for calibrating the model

Cross-sectional data for Ireland observed in 2005 were primarily considered for calibrating the model. This focus on cross-sectional data was adopted after careful consideration, taking into account the limitations of the structural model and the primary purpose for which the model has been devised. The model is limited in the sense that it does not capture real-world uncertainty over a range of characteristics, including the evolving tax and benefits system, conditions of the macro-economy, household demographics, and so on. As such, calibrating the model to survey data reported for a population birth cohort requires the implicit assumption that either changes in the policy environment have an incidental impact on behaviour, or are perfectly foreseen. The former of these assumptions is difficult to maintain when the primary purpose of the model is to explore behavioural responses to policy reform, and the latter is patently inaccurate. Cross-sectional data avoid these problems because they describe behaviour observed under a single policy environment. The assumptions implicit in the calibration are then that: a) individuals base their decisions on the belief that the existing policy environment will be maintained into the indefinite future; and¹⁴ b) that expectations regarding the future evolution of individual specific characteristics – including demographics, wages, employment opportunities, and so on – can be based upon age profiles exhibited by contemporary

¹⁴ This assumption is not uncommon in the associated literature.

survey data. The former of these assumptions appears to us to be plausible (if not necessarily accurate), as does the latter after an allowance is made for trend improvements in wages and survival probabilities. These underlying assumptions should be borne in mind when interpreting the discussion that follows.¹⁵

4.2. Calibration approach

Models of the type referred to above are parameterised against observed data via a two-stage process that adapts to the large number of parameters involved and the computation times required to determine the implications of a given parameter combination. In the first stage, observable model parameters – including those governing inter-temporal wage dynamics and transitions in relationship status – are each estimated separately on available survey data. Given the parameter estimates obtained in the first stage, the second stage involves adjusting the (unobserved) model parameters to match simulated moments implied by the structural model to sample moments estimated from survey data.

The second stage of the model parameterisation is usually conducted either by manual calibration or optimisation of a loss function using an econometric criterion.¹⁶ Various methods exist, each with their own advantages and disadvantages. We chose to manually adjust unobserved model parameters in the second stage of the parameterisation, based upon the sum of squared errors for each age specific model characteristic and graphical representations of the respective characteristics, following the approach by Sefton et al. (2008). Although this approach sacrifices some objectivity in the method by which the parameter estimates are obtained, it also facilitates a detailed understanding of the behavioural implications of alternative parameter combinations, relative to an automated econometric “black-box”.

4.2.1. Specification of the model considered for calibration

As the second stage of the calibration requires testing over a very large number of parameter combinations, the model was limited to the following eight characteristics:

¹⁵ A third possibility, which has been considered in the associated literature, is to calibrate the model to population characteristics that control for time and cohort effects (e.g. (Sefton et al. 2008)). This option has the problem that the details of the policy environment implicit in such profiles represent an average of the policy environments that applied during the period considered for estimation, and as a consequence are not well defined.

¹⁶ Econometric methods include Simulated Minimum Distance (Lee and Ingram, 1991), Method of Simulated Moments (Stern, 1997), Indirect Estimation (Gourieroux et al., 1993) and Efficient Method of Moments (Gallant and Tauchen, 1996).

- age
- number of adults
- wage offers
- wage rates
- net liquid assets
- pension eligibility
- pension rights
- time of death

This restricted model focuses on decisions over labour supply (including the possibility of part-time employment), consumption, and pension participation, given a household's age, its number of adults, liquid assets, wage offer, wage rate, pension scheme eligibility, pension wealth, and survival. Household decisions were considered at annual intervals between ages $t_0 = 20$ and $T = 120$, with labour supply possible to age 75. State Pensionable Age was set to $t_{SPA} = 65$, the pensionable age that prevailed in 2005. Uncertainty was taken into consideration for the intertemporal development of the number of adults in a household, wage offers, wage rates, private pension eligibility, 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 permissible state space (the range of characteristics that any household might conceivably have) into a series of grids. The domains of wages and wealth between ages 20 to 69 were each divided into 34 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 age $t = t_{SPA}$ is taken as a tax free lump, with the remainder taken as a retirement annuity. The domain of the retirement annuity was divided into 16 points using a log scale between ages 65 and 75. From age 76 to age 120, the wealth and retirement annuity domains were each divided into 151 points using a log scale.

Three additional dimensions – reflecting the number of adults in a household, wage offers, and pension scheme eligibility – 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 considered to be comprised of a single adult. Between ages 20 and 75, solutions were required for households with and without a wage offer. Furthermore, 3 private sector pension schemes were considered for analysis.

This specification of the model required utility maximising decisions to be numerically evaluated for 12,283,729 different combinations of household characteristics, for each alternative parameter combination tested as part of the calibration process.¹⁷ For reference, this specification of the model takes 25 minutes to run on a computer with an Dell T5500 workstation with dual Xeon X5650 processors and 6Gb of RAM.

¹⁷ $= (64-19).34.34.16.3.2.2 + (75-64).34.34.16.2.2 + (95-75).151.151.2 + (120-95).151.151$

4.2.2. Calibration strategy

The parameters of the model that were not estimated on observable data (or otherwise exogenously assumed) were calibrated by comparing age profiles at the household level for both singles and couples of:

1. the geometric mean of household employment income
2. the variance of household log employment income
3. the proportion of adult household members employed full-time, part-time and not at all
4. the geometric mean of household consumption
5. the variance of household log consumption

Statistics for calibration were drawn mainly from 3 surveys by the Irish Central Statistics Office. Statistics on employment were derived from SILC 2005. Proportions employed full-time, part-time or not in work were derived from the Quarterly National Household Survey (April 2005), while statistics on consumption expenditure were derived from the Household Budget Survey, again for 2004/2005.

Age specific geometric means of household employment income were matched by altering the distribution mean of the simulated cohort at entry to the model (age 20), μ_{n^a, t_0} , and by adjusting the age and relationship specific trend parameters of human capital described by $m(n_{i, t}^a, t)$ in equation (10). The variance of log employment income by age and relationship status was matched by adjusting the variance of the distribution at entry to the model, σ_{n^a, t_0}^2 , and the variance of age specific innovations, $\sigma_{\omega, n_{i, t-1}^a}^2$. Age and relationship specific rates of employment participation were matched by adjusting the utility price of leisure, α , the learning by doing effects, $\kappa(n_{i, t-1}^a, t-1)$. Learning by doing effects, $\kappa(n_{i, t-1}^a, t-1)$ were also adjusted to match the model to the split between full-time and part-time employment described by survey data, as were the ratios of part-time to full-time wages. Finally, the timing of consumption was adjusted by altering the exponential discount rate δ , and the parameter of relative risk aversion $1/\gamma$. The variance of consumption by age was a residual that depends heavily upon the associated income parameters $\left\{ \sigma_{n^a, t_0}^2, \sigma_{\omega, n_{i, t-1}^a}^2 \right\}$.

It was necessary to select a set of starting values for the model from which to commence the calibration process. Starting with the wage parameters, we began with a flat wage profile over the life course, assuming zero experience effects, $\kappa = 0$, and no risk of a low wage offer. The leisure cost of full-time and part-time employment were defined as non-stochastic and age invariant proportions of the total time available to an adult, assuming 18 'viable' hours per day. Similarly, the ratio of the part-time to full-time wage was assumed to be independent of

age and relationship status. The initial ratios considered for the calibration were calculated using data from the 2005 wave of the SILC; associated statistics are reported in Table 1. Finally, the preference parameters of the model were taken from UK econometric regressions (see (van de Ven 2010)), but the search routine meant that parameters were free to vary in response to characteristics in the Irish data. The exception is the utility price of leisure, which was set deliberately low to ensure an adequate sample for calculating moments of employment income.

The model calibration was conducted using a cascading procedure designed to subject the most flexible aspects of the model to the most frequent instances of re-adjustment. From the list of moments referred to above, the model exhibits the greatest degree of flexibility in relation to the geometric means of household employment income, where the number of associated model parameters $\left\{ \mu_{n^a, t_0}, m \left(n_{i, t}^a, t \right) \right\}$ is identical to the number of moments considered for the calibration. The calibration consequently focused in the first instance upon adjusting the parameters $\left\{ \mu_{n^a, t_0}, m \left(n_{i, t}^a, t \right) \right\}$ until a close match was obtained between the simulated and sample estimates for the geometric means of employment income.

Given the calibrated parameters for employment income, the calibration focused next upon matching the incidence of employment participation/non-employment. Here, the utility price of leisure α serves to reduce the preference for employment in general, and the learning-by-doing effects κ increase employment early in the working lifetime, relative to later life. The parameter adjustments necessary to match the model to employment participation, also serve to distort the match obtained to labour income, both through the direct effect that varying the parameters κ have on the intertemporal development of latent full-time wages, and indirectly through distributional heterogeneity in labour supply responses to employment incentives. Hence, the calibration process proceeded in an iterative loop to match the model to both the geometric mean of employment income and employment participation at the same time.

The calibration procedure focused next upon matching the model to rates of full-time and part-time employment. This aspect of the calibration proceeded in a very

Table 1: Model Parameters to Distinguish the Effects on Leisure and Labour income of Alternative Labour Supply Decisions

<i>employment option</i>	leisure cost	proportion of full-time wage
not employed	0.00	0.00
part-time employed	0.145	0.188
full-time employed	0.322	1.000

Source: authors' calculations on data from SILC 2005
Notes: based on population average statistics for full-time and part-time employed leisure cost assumes 18 allocatable hours per day and 7 days per week

similar fashion to that set out for employment participation, with the ratio of part-time to full-time wages replacing the utility price of leisure in the adjustment of parameters. It is important to note that this adjustment procedure has the very significant advantage that the wage parameters derived via the calibration take full account of the endogeneity of labour supply decisions, with which so much of the associated econometric literature has been concerned following the seminal contribution by James Heckman.

Having obtained a close match to moments of both employment income and labour supply, the calibration then focused upon matching the model to sample moments of household consumption. The model offers relatively blunt tools with which to achieve this match, and the associated calibration is somewhat more approximate as a result – in particular, we focused upon achieving a match between the peaks in consumption described by the simulated and sample data, and the general trend of age specific variances in consumption. In this regard, the discount rate δ tends to shift consumption into later periods of life, increasing the slope of the lifetime consumption profile. The parameter of relative risk aversion $1/\gamma$ motivates increased precautionary saving early in the working lifetime, which diminishes as the working lifetime proceeds. An alternative aspect that has been recognised as important here is the bearing that demographic needs have on consumption preferences; this aspect of the model was omitted from the calibration, due to the exogenous assumption of age specific demographics (reported in Section 5.4), and the revised OECD equivalence scale upon which the preference relation is based.

To summarise, the model parameters $\left\{ \mu_{n^a, t_0}, m \left(n_{i, t}^a, t \right) \right\}$ were then adjusted until a close match was obtained to the age and relationship specific geometric means for employment income. Given the parameters $\left\{ \mu_{n^a, t_0}, m \left(n_{i, t}^a, t \right) \right\}$, the model parameters $\left\{ \alpha, \kappa \left(n_{i, t-1}^a, t-1 \right) \right\}$ and the ratio of part-time to full-time wages were adjusted to match the simulated to sample rates of employment. This process was then repeated a number of times until the model obtained a reasonable match to both geometric means for employment income and rates of employment at the same time. The parameters $\left\{ \delta, 1/\gamma \right\}$ were then adjusted to obtain a better match to age specific geometric means for consumption described by survey data, and the parameters $\left\{ \sigma_{n^a, t_0}^2, \sigma_{\omega, n_{i, t-1}^a}^2 \right\}$ were adjusted to obtain an improved match to the age specific moments of both consumption and labour income. The entire process was then repeated to obtain the calibrated results that are reported in Section 6.

5. ESTIMATES FOR OBSERVABLE PARAMETERS

The model parameters for which exogenous estimates were obtained are principally concerned with four key issues: life expectancy, the terms of the available pension schemes, taxation, 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 are described in turn below.

5.1. Life expectancy

The survival probabilities assumed for calibrating the model are based upon CSO Population and Labour Force Projections, 2006-2036. These data are based upon observed survival rates between 2006 and 2007, and Official projections for improved longevity thereafter. The Official data permit survival rates to be calculated to age 99. Age specific survival probabilities between 100 and 120 were exogenously specified to obtain a smooth sigmoidal progression from the official estimate at age 99 to a 0 per cent survival probability at age 120. These probabilities are reported in Table 2.

Table 2: Exogenously estimated model parameters

rates of return / growth (% per annum)									
pension wealth:		4.1%	min. cost of debt:		6.0%	wage growth:		1.6%	
positive liquid wealth:		4.1%	max. cost of debt:		19.0%				
children			probability of mortality						
age	singles	couples	age	singles	couples	age	prob.	age	prob.
20	0.054	0.000	40	0.656	1.952	81	0.007	101	0.415
21	0.049	0.336	41	0.706	2.006	82	0.008	102	0.486
22	0.078	0.355	42	0.724	2.044	83	0.009	103	0.555
23	0.110	0.428	43	0.752	1.984	84	0.010	104	0.619
24	0.120	0.569	44	0.634	1.843	85	0.012	105	0.679
25	0.156	0.572	45	0.595	1.683	86	0.014	106	0.733
26	0.188	0.728	46	0.503	1.604	87	0.017	107	0.781
27	0.238	0.822	47	0.407	1.389	88	0.020	108	0.823
28	0.283	1.004	48	0.390	1.244	89	0.022	109	0.860
29	0.360	1.069	49	0.360	1.174	90	0.025	110	0.890
30	0.380	1.161	50	0.328	1.074	91	0.028	111	0.916
31	0.402	1.363	51	0.310	0.934	92	0.034	112	0.936
32	0.422	1.419	52	0.242	0.773	93	0.041	113	0.953
33	0.466	1.453	53	0.147	0.715	94	0.050	114	0.966
34	0.550	1.615	54	0.133	0.595	95	0.063	115	0.975
35	0.587	1.721	55	0.087	0.515	96	0.075	116	0.983
36	0.593	1.708	56	0.071	0.418	97	0.120	117	0.988
37	0.671	1.790	57	0.058	0.418	98	0.197	118	0.992
38	0.631	1.892	58	0.057	0.357	99	0.273	119	0.995
39	0.638	1.944	59	0.029	0.336	100	0.343	120	1.000

Source: age profiles for children equal to arithmetic averages calculated from ** survey data mortality probabilities calculated for couples where both members are aged 20 in 2005 on life-tables published *** return to pension wealth and positive balances of liquid wealth set equal to real growth observed for Irish GNP between 1970 and 2005 cost of debt exogenously assumed real wage growth calculated on data for workers in all industries between 1985 to 2006

5.2. The terms of private sector pension schemes

The terms of private sector pensions in Ireland are complex and diverse (see, for example, the *Pension Market Survey 2007*, IAPF). Defined benefit schemes remain important, but in the private sector, and especially for younger workers, defined contribution schemes have become much more common. In order to summarise Irish private sector pensions in a tractable fashion, we have opted to characterise the system in terms of a set of pension options which a worker may face. We represent DB schemes in terms of a DC scheme with a higher employer contribution – this helps to capture a key feature of DB schemes, while at the same time keeping the complexity of the problem to a manageable level.¹⁸

Three schemes are considered for the calibration, designed to reflect low, middle, and high pension contribution rates by employees and their employers. Employees are considered to be able to decide over whether to participate in these schemes, but not their respective rates of pension contributions, as is common for occupational pensions. The terms applied to each of the representative pension schemes are summarised in Table 3.

The top panel of Table 3 reports the rates of employee and employer pension contributions assumed for each alternative pension scheme. In each year between ages 20 and 64, households are allocated a pension scheme that they may choose to participate in during the respective year. The pension scheme to which a household is eligible in any given year is either carried over from the scheme that they chose to participate in during the preceding year, or – if they chose not to participate in a pension during the preceding year – then it is taken as a random draw with reference to the income specific probability distributions reported at the bottom panel of Table 3.

The statistics that are reported in Table 3 reflect the stylised observation that employer pension provisions tend to improve with employee wages, where pension support is virtually non-existent for employees on low wages – defined here as those with full-time wages worth less than €16,000. In contrast, many employees toward the top of the wage distribution tend to enjoy relatively generous pension support from their employers, while the majority of workers lie between these two extremes.

Furthermore, we ignore associated decisions regarding the portfolio allocation, and assume that all returns to investment are risk free. The rate of return to pension

¹⁸ Note that in our context, where there is no investment uncertainty and mortality rates are known, a career average DB scheme can be equivalently restated in terms of a DC scheme without loss of generality.

wealth is set to 4.1% per annum, equal to the average real growth of Gross National Product in Ireland during the period 1970 to 2005 (reported in the top panel of Table 2). Pension wealth is converted into an actuarially fair annuity at age 65 based on the assumed rate of return to pension wealth and the mortality rates discussed in Section 5.1. The value of this annuity is assumed to fall by 50% upon the death of a spouse.

Table 3: Terms Assumed for Private Sector Pensions: contribution rates and probabilities of eligibility

	scheme 1	scheme 2	scheme 3
contribution rates (% of labour income)			
employee	3%	4%	7%
employer	0%	7%	11%
eligibility probabilities by annual income band			
to € 16,000	90%	10%	0%
to € 38,400	10%	70%	20%
€ 38,400 and over	0%	40%	60%
Notes: authors' assumptions for terms of private sector pensions real return to pension wealth set to 4.1% p.a . income thresholds for probability distributions indexed to real wage growth of 1.6 % p.a .			

5.3. Taxes and benefits

We adopt a simplified representation of the tax/welfare system, which nevertheless captures some of the key features of interest. For the cohorts now entering the labour market, coverage of the State Contributory Pension scheme will be much higher than heretofore. We consequently adopt the simplifying assumption that, in future, all those aged above State Pension Age will be eligible for the State Contributory Pension. The model allows for both the State Pension Age and the level of payment to be varied.

For those of working age, we take account of the following schemes:

- Jobseekers' Allowance
- One Parent Family Payment
- Child income support via child benefit, qualified child increase and Family Income Supplement

On the income tax side, we allow for the basics of personal and PAYE tax credits, tax bands and rates, and for PRSI and levies (which may be structured along the lines of the Universal Social Charge). Special attention is given to alternative possible tax treatments of pensions, varying from the EET (exempt, exempt, taxed) structure which approximates that in place until recent years to a potential new system with

tax relief at a single hybrid rate, as per the recommendations of the National Pension Framework.

An additional issue of concern in relation to the simulated tax and benefits system is the way that it is assumed to evolve with time. Given the wage growth that is used to adjust the financial statistics against which the model is calibrated (reported in Section 4), ignoring indexation would result in “fiscal drag” (or tax bracket creep) and a decline in the relative value of benefits. Three main approaches to this issue can be distinguished. One is to allow for full indexation of tax parameters and welfare payment levels with respect to wage growth. This has the merit of ensuring that the ratio of tax to income remains constant, and that welfare incomes rise in line with general wage growth. This approach is in line with the distributionally neutral benchmark adopted in analysis of budgetary impact.

An alternative approach would be to project indexation of tax parameters and of welfare rates in line with past experience; the indexation parameters applying to tax and welfare might then differ from each other, and from wage indexation. Data for the period 1987 to 2005 indicate benefit parameters, and especially tax parameters, were adjusted by more than the growth in wages. Projecting forward on the basis of this experience does not seem advisable. It must be remembered that the public finance situation in 2005 was boosted by revenues arising from the house price bubble. When projecting forward on a very long term basis, it would be desirable to incorporate the adjustment currently under way to bring a sustainable fiscal balance. Part of the challenge, therefore, is to construct a scenario which takes account of the required adjustment, while not imposing an excessive adjustment over the very long term. We have consequently adopted the former approach here, adjusting tax thresholds and benefits in line with wage growth of 1.6% per annum.

5.4. Household demographics

The calibration that is reported here assumes that a household can be comprised of one or two adults between ages 20 and 95, where the number of adults is considered to be uncertain between adjacent years. From age 96, all households are comprised of a single adult. The logit model considered to describe the evolution of adults in a household is described by equation (15):

$$s_{i,t+1} = \alpha_0^A + \alpha_1^A t + \alpha_2^A t^2 + \alpha_3^A t^3 + \alpha_5^A s_{i,t} \quad (15)$$

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. This logit equation was estimated using data derived from waves 7 and 8 of the Living in Ireland survey. Regression statistics are reported in Table 4.

Table 4: Regression Statistics for Logit Model of Relationship Status

variable	coefficient	std. error
single(t-1)	9.031	0.371
age	-0.679	0.138
age_2	1.34E-02	2.96E-03
age_3	-7.53E-05	1.94E-05
Constant	4.450	1.933
sample size		6137
proportion single		0.418
correct predictions		0.985

Source: authors' calculations on data from waves 7 and 8 of the Living in Ireland survey

Dependant children were modelled deterministically when calibrating the model, based on age and relationship specific averages reported in Living in Ireland survey data. These age specific averages are reported in Table 2.

6. CALIBRATED MODEL PARAMETERS

Our calibrated model parameters are reported in Table 5, and the associated fit between the simulated and sample moments is reported in Figures 1 to 4. We begin by interpreting the calibrated model parameters, and the key considerations underlying the parameter values that we settled upon. We then discuss the ways in which our calibration could be improved, which remain for further research.

6.1. Interpreting the calibrated parameters

Starting with the parameter of relative risk version, the calibrated value of 3.1 implies an intertemporal elasticity of substitution for consumption calculated at population averages of 0.16, which lies firmly within the range of estimates reported in the associated empirical literature. Grossman & Shiller (1981), Mankiw (1985) and Hall (1988), for example, report econometric estimates for the intertemporal elasticity between 0 and 0.4, Blundell et al. (1994) report an estimate of 0.75, while Hansen & Singleton (1983) and Mankiw et al. (1985) report estimates just over 1. Although values of the coefficient of risk aversion required to explain the equity premium puzzle (Mehra & Prescott (1985)) are large by comparison, evidence from attitudinal surveys suggest that the value is unlikely to larger than 5 (Barsky et al. (1997)).

The relative values of the intra-temporal elasticity (ϵ) and relative risk aversion ($1/\gamma$) imply that consumption and leisure are direct substitutes, which has been suggested as a potential explanation for the fall in consumption that is commonly observed about retirement (e.g. Heckman (1974)). The discount factor indicates more impatience than the assumed real rate of return (5.0% c.f. 4.1% per annum),

and the utility price of leisure for singles and couples is in the region of 1.0 by construction.¹⁹ The probability of a low wage offer is 20% at any age between 20 and 75 for single adults, and 1% for couples. These parameters appear to display a passable level of internal consistency, given the imperfect correlation associated with the likelihood of involuntary unemployment for a husband and wife.²⁰ The ratio of part-time to full-time hours of employment was set equal to the associated sample averages reported in SILC data, as defined in Table 1. In contrast, the ratio of part-time to full-time wages was reduced by one third, relative to the associated survey data, to dampen incentives to take up part-time employment.

Turning to the age specific model parameters reported in Table 5, the calibration produced experience effects that tend to decline with age, where the experience effects identified for singles exceed those for couples throughout the simulated working lifetime. This difference is most pronounced early in the simulated lifetime, where the population tends to be primarily comprised of singles adults: at age 20, a single adult who chooses to work full-time can expect to earn 25% more by age 21 than they would have done had they chosen not to work at all at age 20. This compares with a 2.5% expected wage premium for couples at age 20. This focus of experience effects early in the working lifetime is consistent with the use of an experience effect as a tool for motivating employment participation early in the simulated lifetime.

The parameters that describe the age dependent component of the intertemporal evolution of latent full-time wages are best interpreted taking account of the experience effects that are described in the preceding paragraph. For singles, these parameters imply positive wage growth of 6% per annum on average for individuals who work full-time between ages 20 and 40, relative to wage decline of 12% per annum for individuals who do not work – part-time employment falls between these two extremes. In contrast, full-time employment implies an approximately flat wage profile (in real terms) between ages 40 and 65, relative to real wage decline of 6% per annum in respect of non-employment. A similar profile is described for adult couples, subject to smaller experience effects. In the case of full-time employed couples, for example, average wage growth between ages 20 and 40 is 1% per annum, relative to an average wage decline of 1% where employment is not supplied. From age 65, wages tend to fall quite sharply for both singles and couples, even where full-time employment is maintained.

¹⁹ This is achieved by multiplying the equivalence scale by 550, to normalise equivalised consumption (c/θ).

²⁰ If there is a 20% probability of any adult being unemployed, and the probability of employment is independent of spouse labour status, then there would be a 4% probability of both members of a couple being unemployed at the same time. The calibration is not sufficiently precise to distinguish between a 1% and 4% probability of unemployment.

Table 5: Model Parameters Calibrated to Match Simulated to Sample Moments

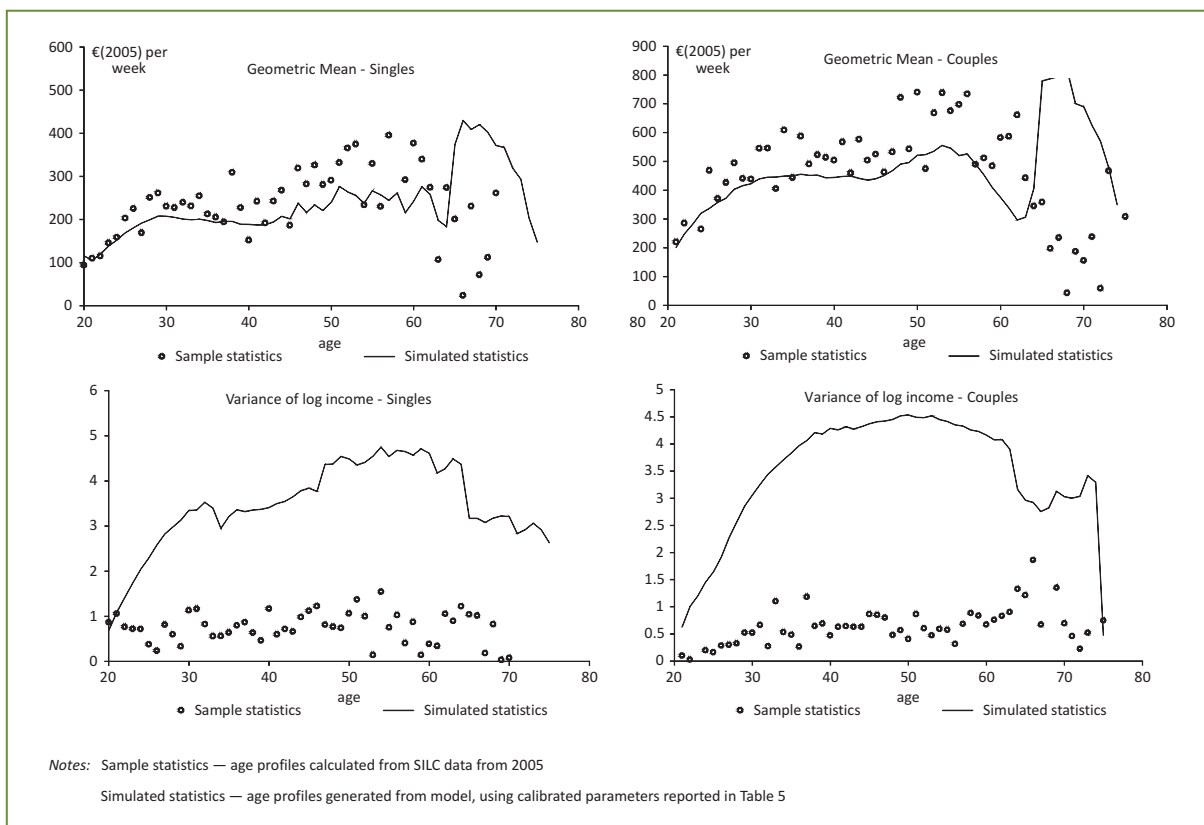
<i>Preference Parameters</i>									
relative risk aversion (1/gamma)		3.10	elasticity of substitution (epsilon)				0.55		
discount factor (delta)		0.95	utility price of leisure (alpha)				0.55*		1.7**
<i>Wage Parameters</i>									
prob of low wage offer - singles		0.20	part-time to full-time leisure cost				0.463		
prob of low wage offer - couples		0.01	part-time to full-time wage ratio				0.400		
<i>Age Specific Parameters</i>									
		trend income		experience effect		trend income		experience effect	
age	singles	couples	singles	couples	age	singles	couples	singles	couples
20	1000.000	1000.000	0.250	0.025	48	34.514	662.168	0.082	0.015
21	857.021	1333.375	0.244	0.025	49	33.264	664.906	0.076	0.014
22	781.609	1547.855	0.238	0.024	50	32.015	685.152	0.070	0.014
23	770.084	1653.848	0.232	0.024	51	30.765	683.423	0.064	0.013
24	729.782	1603.066	0.226	0.024	52	29.515	687.220	0.058	0.013
25	684.510	1601.240	0.220	0.023	53	28.266	703.668	0.052	0.013
26	628.312	1590.876	0.214	0.023	54	27.016	681.389	0.046	0.012
27	565.330	1594.449	0.208	0.022	55	25.767	650.762	0.040	0.012
28	494.661	1650.370	0.202	0.022	56	24.517	636.596	0.034	0.012
29	417.112	1539.240	0.196	0.022	57	23.268	614.606	0.028	0.011
30	358.136	1446.184	0.190	0.021	58	22.018	581.439	0.022	0.011
31	287.453	1390.094	0.184	0.021	59	20.768	548.271	0.016	0.010
32	240.418	1298.615	0.178	0.021	60	19.519	515.103	0.010	0.010
33	183.914	1212.764	0.172	0.020	61	18.269	481.935	0.010	0.010
34	147.897	1142.045	0.166	0.020	62	17.020	448.767	0.010	0.010
35	138.391	1051.007	0.160	0.019	63	15.770	415.599	0.010	0.010
36	117.694	976.720	0.154	0.019	64	14.521	382.431	0.010	0.010
37	104.831	889.670	0.148	0.019	65	13.271	349.263	0.010	0.010
38	93.217	839.271	0.142	0.018	66	12.022	316.096	0.010	0.010
39	78.333	768.300	0.136	0.018	67	10.772	282.928	0.010	0.010
40	67.889	749.428	0.130	0.018	68	9.522	249.760	0.010	0.010
41	62.398	723.956	0.124	0.017	69	8.273	216.592	0.010	0.010
42	54.896	709.681	0.118	0.017	70	7.023	183.424	0.010	0.010
43	51.315	676.203	0.112	0.016	71	5.774	150.256	0.010	0.010
44	50.139	654.010	0.106	0.016	72	4.524	117.088	0.010	0.010
45	44.813	641.729	0.100	0.016	73	3.275	83.920	0.010	0.010
46	37.013	652.456	0.094	0.015	74	2.025	50.752	0.010	0.010
47	35.763	652.747	0.088	0.015	75	0.776	17.585	0.010	0.010
* singles; ** couples									

6.2. The match between simulated and sample moments

We discuss the match obtained between the simulated and sample moments in the same order in which we conducted the model calibration, as described in Section 4.2. The top two panels of Figure 1 indicate that the model obtains a close match to the

age profiles described by survey data for the geometric mean of private non-property (employment) income, for both single adults and couples. Given the discussion provided in Section 4.2, it is reasonable to expect that this aspect of the calibration should obtain a close match to survey data, because the number of associated model parameters is exactly equal to the number of calibration moments. The most obvious anomaly is the jump up in the geometric mean of employment income that is evident for couples at age 65, which is the pensionable age considered for the calibration. This jump up is not generated for any household taken in isolation – indeed, as noted above, household full-time potential wages tend to fall late in the working lifetime – rather, the increase in the geometric mean of employment income later in the working lifetime reflects a mass departure from employment of lower wage households after they gain access to their accrued pension wealth. Although this jump up in the geometric mean of non-property income is not evident in the survey data, it is important not to overstate the importance of this disparity, as relatively few adults choose to be employed after pension age.

Figure 1: Private Non-Property Income Profiles by Age – simulated versus sample moments



This brings us neatly to the profiles for employment that are displayed in Figure 2, delaying for a moment discussion of the variances of employment income that are reported in the lower panels of Figure 1. The two panels of Figure 2 indicate that the model does a very good job of capturing observed rates of employment participation for both singles and couples, with the most substantial disparity

between the simulated and sample statistics being the relatively high rates of full-time employment observed among single adults between ages 35 and 45. This disparity is attributable to the application of taxes and benefits, as the simulation model assumes that single adults without children receive no benefits if they are in full-time employment, but can receive Jobseeker's Allowance if they work part-time. Single adults with no children and low full-time wage potentials can therefore significantly increase their disposable income by electing to work part-time. Single adults with children, however, are eligible to OPFP and Child Benefit irrespective of their employment status, and FIS if they work full-time, so that full-time employment obtains an unambiguous increase in disposable income when a single adult cares for at least one child. The balance between these countervailing incentives switches from part-time to full-time employment during peak child-rearing years.

Two important factors underly the close match between the simulation model and the data that is otherwise reported for employment statistics. First, it was necessary to assume that the wage earned from part-time employment returns a smaller fraction of the full-time wage that the population average statistics imply; without this assumption, the model tended to generate too much part-time employment, relative to the incidence described by the associated sample statistics. A possible explanation for this is that the sample data are affected by selection effects, so that those who take up full-time employment tend to have poorer options if they were to work part-time than current part-time workers, and vice versa for current part-time employees. Indeed, qualitative data give some credence to this view. Fagan (2003), for example, reports that approximately 1 in 5 employed people in Europe work full-time when they would prefer to work part-time. The reasons most commonly given for the mis-match include the perception that it would not be possible to do a desired job part-time, that part-time employment is not offered by a desired employer, and that it would damage career prospects.

The second factor underlying the match obtained between the model and sample moments of employment is the allowance for an experience effect on future prospects for the latent full-time wage. In the absence of this experience effect, the model tended to generate too little employment participation at the outset of the working lifetime, as is common in the associated literature. Low (2005) points out that a casual inspection of the data shows that young workers tend to command a low wage but have high participation rates, whereas older workers have a higher wage but lower participation rates. It is difficult to reconcile these stylised facts with an intertemporal model of labour supply in which the age-earnings profile is deterministic. Both Low (2005) and French (2005) suggest that a possible explanation for the apparent inconsistency is the self-insurance motive where incomes are stochastic. Individuals work hard when young to accumulate assets, which insure them against wage uncertainty in later life. Yet the careful simulations of Low (2005)

suggest that this motive can only partially reconcile a life-cycle model with survey data. We find the same applies here.

One other anomaly that does present itself in Figure 2 is the jump up in the proportions of adults choosing not to be employed at age 65 (pension age), which is particularly evident for couple households, and not as evident in the associated survey data. In this regard, it is important to note that the model assumes that all households access their pension wealth at age 65. In practice, however, many pension schemes make provisions for both early and late retirement, so that the impact of pension eligibility tends to be concentrated in the simulation model in a way that it is not in the survey data.

As noted in Section 4.2, the age profiles for the geometric means of consumption generated by the model were amended by adjusting the discount rate δ and the parameter of relative risk aversion $1/\gamma$. Increasing the former of these tends to tilt the consumption profile up with age, and increasing the latter depresses consumption early in the working lifetime (when prospective wage uncertainty is high), and raises consumption toward the end of the working lifetime (as prospective wage uncertainty declines). The top two panels of Figure 3 indicate that the model broadly matches the sample moments for geometric mean consumption by age calculated from survey data. For singles, the sample moments suggest that the age profile of consumption starts out flat, and then falls away in retirement; and for couples, it rises to a discrete peak about age 50. We focused most of our effort here in trying to capture the peak at age 50 described for couples by the sample moments, achieved by increasing both δ and $1/\gamma$. Our scope for increasing δ was limited by our desire not to obtain monotonically increasing consumption profiles with age; in relation to relative risk aversion we were also limited by what is considered 'reasonable' by the wider empirical literature. We return to alternatives that might help to improve this aspect of the calibration below.

The final set of moments considered for the calibration were the age specific measures of the variance of log employment income and log consumption, which are reported respectively in the lower halves of Figures 1 and 3. Note that the same set of parameters were adjusted to match the simulation model to these two separate series of sample statistics; the parameters controlling the variance of (log) latent full-time wages. The lower panels of Figures 1 and 3 indicate that it was not possible to match to both the variances described for employment income and the variances of consumption at the same time. We consequently opted to focus primarily upon the variances for consumption, bearing in mind that these have the most important bearing on household saving and welfare. Hence, while the model substantially overstates age specific variances of employment income described by survey data, it broadly matches variances for household consumption.

Figure 2: Employment Rates — simulated versus sample moments


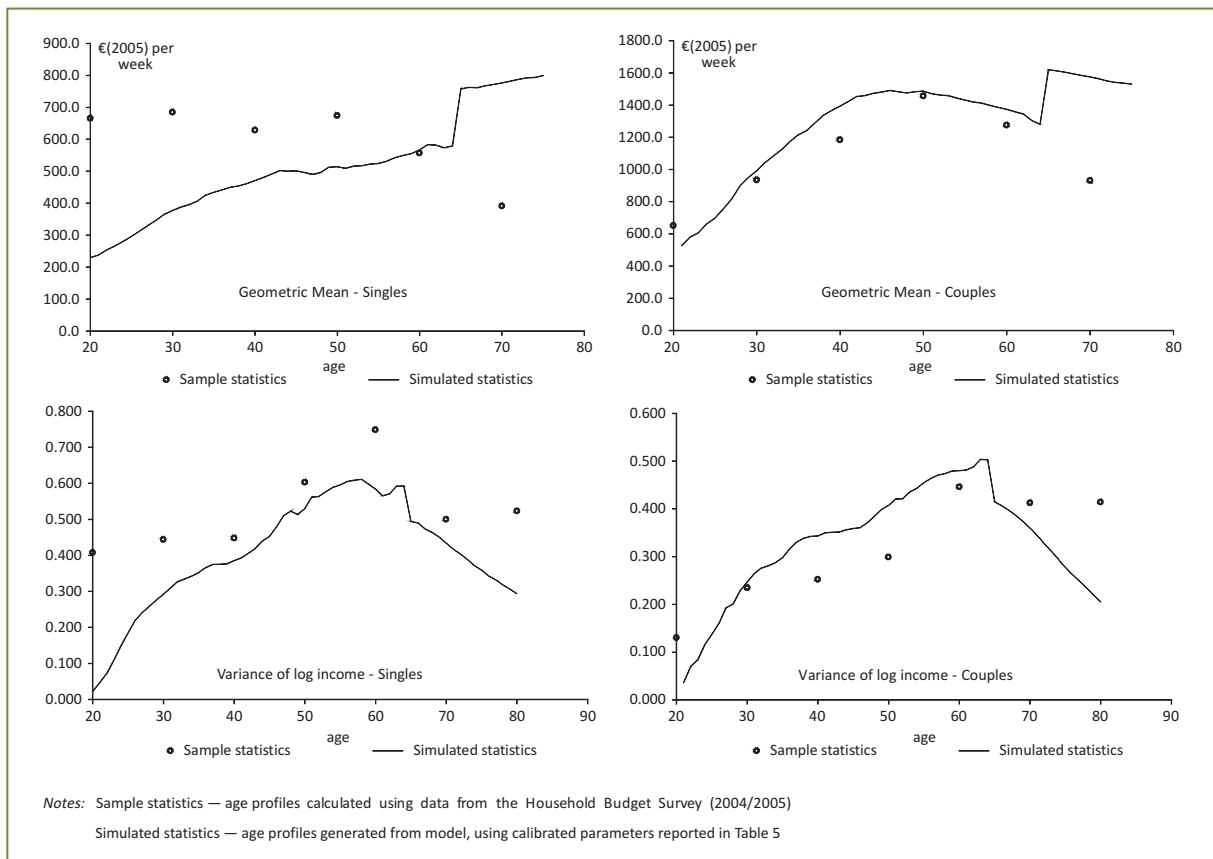


Figure 3: Consumption Profiles by Age — simulated versus survey data



6.3. Improving the model fit

Three aspects of the model calibration that is reported above appear to warrant further attention. First, the model obtains a very approximate fit to the profile of age specific geometric means for consumption. Second, the model fails to capture the variances of employment income and consumption simultaneously – we can calibrate it only to one or the other of these two sets of statistics. And third, the model generates discrete shocks to labour supply and employment income about the assumed pension age that are not evident in survey data. The last of these appears to be the most straightforward to address, and least concerning of the three; this disparity is likely to disappear after allowance is made for the more flexible timing of pension dispersals that is often possible in practice. We consequently devote the remainder of this section to discussing the other two concerns that are noted above.


The relationship between employment income and disposable income generated by the model has a strong bearing upon the mismatch between the simulated and sample moments for consumption. Consider, for example, the associated statistics for disposable income that are reported in Figure 4. First, it is encouraging to note that the simulated and sample geometric means are closely aligned to state pension age, suggesting that the way that we have described the tax and benefits function during the working lifetime provides a decent reflection of the practical reality. Two key discrepancies do, however, emerge. First, although a close match is obtained to the geometric means of employment income and disposable income, the variances associated with these two distributions suggest that the model tax function produces greater redistribution than is achieved in practice. And secondly, there is a substantial jump up in disposable income generated by the simulation model at pension age that is not displayed by the sample data. These two disparities between the model and the survey data are clearly consistent with the disparities reported above for consumption, suggesting that a common set of distortional factors is responsible for both.

Disposable income in the model is comprised of employment income, property income, and the influence of taxes and benefits. As the disparity between the dispersion of employment and disposable income is large throughout the simulated lifetime – irrespective of the temporal aspect of property income accrual – the associated departure of the model from the statistical record is likely attributable to the application of taxes and benefits. Indeed, there is good reason to suppose that this is true, given that the stylised nature of the simulation model is ill-adapted to providing a comprehensive description of the Irish tax and benefits system. At the most basic level, the model does not include a very wide range of population characteristics that have an important bearing on the transfer payments to which individuals are eligible in practice. These characteristics include sickness, injury, disability, and the number and age of dependant children in a household.

Furthermore, the focus of the model on annual time increments means that it is not possible to capture heterogeneity that depends on shorter time intervals, such as part year unemployment and so on. Hence, it is to be expected that this population heterogeneity that is unaccounted for in the model should result in greater homogeneity of disposable income than is observed in practice. Addressing this disparity in a way that is computationally feasible is an issue that remains for further research.

In relation to the spike up in disposable income that is generated by the simulation model at pension age, some progress may be made by allowing for the more flexible terms of pension fund dispersals, as discussed in the preceding subsection. Yet, the scale of the jump in disposable income suggests that the model is also generating excessive saving through pension assets (which produce an annuity income stream from retirement). We intend to explore how the model matches to pension fund participation in future work.

Figure 4: Disposable Income Profiles by Age — simulated versus sample moments



7. EXPLORING POLICY ISSUES

The model has been constructed in such a way as to allow a range of pension policy issues to be analysed. These include the introduction of mandatory DC pensions, changes in the indexation of State pensions, changes in the State pension age,

and changes in the level of State pension benefits. In this section we illustrate the application of the model by exploring trade-offs between changes in the level of State pension benefits and changes in the age at which the State pension becomes payable. Increases in the State pension age are widely viewed as a potential element of a strategy to cope with the demands of demographic ageing. If State Pension Age does not rise to some extent in response to longer life expectancies, fiscal constraints will imply that State pension benefits will be lower than if pension ages do adjust to life expectancy increases.

The model allows this trade off to be identified more precisely, taking account of economic responses to the changed incentives which arise from differing combinations of the level of benefit and the age at which the State Pension becomes payable. Here, we focus on the aggregate impact of a “grid” of policy choices. It should be emphasised that none of these should be regarded as a policy proposal or recommendation. The purpose of this approach is to identify what implications different combinations of benefit levels and State Pension Age have for the exchequer and for society. The level of the State pension payment is allowed to vary between a high of € 230 per week – close to the current, 2011 values – and a low of about € 170 per week. This “low” figure is about 5% below the value of the State Pension in 2005, the data year, which is taken as the “base case” for the analysis, and about 25% below the “high” value.

Until recently, the State Pension Age had been set at 66 – but with a special “retirement” or “transitional” pension available at age 65, making 65 the effective age at which a State Pension could be obtained. Under legislation implementing aspects of the National Pensions Framework, the State Pension Age is set to rise to 66 in 2014, to 67 in 2021 and to 68 in 2028. Given these pending changes in State Pension Age, we explore combinations of different benefit levels with a State Pension Age between 65 and 68.

The implications of differing combinations of State Pension Age and pension payment levels for the government’s budget, for net private saving, for consumption and for employment are set out in Tables 6 to 9 below. It should be noted that the model is geared towards generating the long-term implications of a policy change, and the statistics reported here provide only a qualitative indication of short run incentive effects.

Looking first at the impact on the net budgetary position of the government (Table 6), we see that the net tax take increases with state pension age, and decreases with the generosity the SCP. Increasing state pension age by a year raises the net tax intake by between € 180m and € 255m per year. The higher are benefit levels, the greater the saving from an increase in the State Pension Age.

Table 6: Impact on Net Government Budget, Relative to Base Policy Scenario (Euro millions, 2005 prices, per annum)

<i>value of State Contributory Pension (€ per week)</i>	<i>State Pension Age</i>			
	<i>65</i>	<i>66</i>	<i>67</i>	<i>68</i>
170	178.4	363.5	559.0	743.8
190	-179.3	33.3	251.8	459.7
210	-389.7	-177.0	59.1	287.6
230	-614.7	-381.3	-127.1	115.8

Note: Base simulation assumes State Pension Age of 65 and value of contributory pension of 180 per week

Table 7: Impact on Consumption, Relative to Base Policy Scenario (Euro millions, 2005 prices, per annum)

<i>value of State Contributory Pension (€ per week)</i>	<i>State Pension Age</i>			
	<i>65</i>	<i>66</i>	<i>67</i>	<i>68</i>
170	-139.2	-301.2	-457.6	-610.0
190	139.7	-36.8	-208.7	-381.0
210	428.0	216.1	29.7	-155.3
230	686.6	469.0	268.1	68.2

Note: Base simulation assumes State Pension Age of 65 and value of contributory pension of 180 per week

Increasing the payment rate of the SCP by €10 per week leads to a net increase in exchequer costs of between €85m and €180m per year. The cost increase is naturally lower when the State Pension Age is higher; an equal-valued increase in benefit also turns out to be less expensive, in terms of exchequer cost, at higher levels of benefit.

The broad import of these findings is that, in a steady state situation, a given budgetary envelope for pensions can be compatible with a high pension age and a low payment rate, or a higher payment rate with a low pension age.. For example, a rise from the base case (2005 levels) of €180 per week to €190 per week, combined with an increase in the pension age from 65 to 66, would be broadly neutral for the Exchequer. Thus, raising the pension age by a single year is compatible a higher payment level in a budgetary neutral policy change – or can allow the maintenance of existing levels when public finances are under sustained short and medium term pressure.

Comparing Tables 6 and 7 reveals that the impact of the considered policy counterfactuals on aggregate domestic consumption is an inverted relation of the impact on the government budget. Hence, increased government saving can be interpreted as a form of enforced private saving, and vice versa, an observation that is particularly evident in wake of the 2008 financial crisis.

The effects on aggregate labour supply of the policy counterfactuals, reported in Table 8, have the expected signs, indicating increased employment as the generosity of state pensions is reduced, and as the state pension age is increased. It is noteworthy that the (long-run) effects on the duration of the working lifetime of increasing state pension age that are projected by the model are much smaller than is commonly assumed in the policy debate. The duration of the working lifetime is projected to increase by between 0.1 and 0.2 years for each year that state pension age is increased, with the effect rising with the generosity of state pensions. This is in contrast to the common assumption that a one year increase in state pension age will lead to an equivalent shift in labour market exit rates (implying an equivalent increase in the average duration of the working lifetime).²¹

An important factor underlying this result is that it reflects a long-run effect in which individuals are considered to foresee the higher state pension age from the beginning of their working lifetimes, and adapt their savings behaviour to accommodate the change (returned to below). This highlights the capacity of a dynamic model to reveal trade-offs in behaviour that might not be obvious at first glance, and to provide some quantitative detail around those trade-offs. When interpreting this interesting result, however, a number of additional factors should also be borne in mind:

1. The model does not account for the “signal” effect that the state pension age may have on individual expectations and planning in practice.
2. The model is calibrated to declining wages later in the working lifetime, and this decline is projected on current work profiles, so that these profiles may feed indirectly into work incentives in the policy counterfactual. We have attempted to control for this type of effect, both in the general approach to calibration (which is designed to take endogenous account of selection effects), and by imposing smooth age trends from reasonably early on in the working lifetime (age 45 for singles and 55 for couples).
3. The employment profiles of singles, in particular, do not respond very strongly to the policy environment.

The sensitivity of the model-based results shown here needs to be tested using alternative assumptions about the formation of decisions on retirement, and how they may be influenced by increases in State Pension Age (SPA). There are, however, a number of considerations that argue against the simple assumption that average

²¹ In an assessment of the impacts of pension reforms enacted in the UK in 2011, for example, the Department for Work and Pensions assumed that the announcement of a one year increase in the state pension age of men would “increase the age at which males would exit the labour market from age 55 onwards; for instance, a 66 year-old man would adopt the exit rate from the labour market currently adopted by a 65-year old” Department for Work & Pensions (2011), p. 11, paragraph 28.

retirement ages will increase 1 for 1 with increases in SPA. Such considerations include:

1. If labour market productivity and/or wages tend to decline later in life.
2. If the desire and/or capacity to undertake work tends to decline later in life.
3. If the system of private (personal and occupational) pensions and alternative retirement saving vehicles provide sufficient funds to meet the expenditure needs of individuals later in life, and may be drawn upon prior to state pension age.

In the long run, these considerations suggest that the burden of an increase in state pension age is likely to be shared between lower consumption and decreased leisure.

Table 8: Impact on Employment, Relative to Base Policy Scenario (average years)

value of State Contributory Pension (€ per week)	State Pension Age			
	65	66	67	68
170	0.09	0.22	0.37	0.50
190	-0.09	0.07	0.24	0.40
210	-0.25	-0.06	0.13	0.30
230	-0.35	-0.16	0.05	0.25

Note: Base simulation assumes State Pension Age of 65 and value of contributory pension of 180 per week

Table 9: Impact on Net Private Saving, Relative to Base Policy Scenario (Euro billions, 2005 prices)

value of State Contributory Pension (€ per week)	State Pension Age			
	65	66	67	68
170	3.9	3.1	1.7	1.0
190	-4.2	-4.7	-5.6	-6.9
210	-11.2	-12.4	-13.5	-14.8
230	-18.6	-19.2	-20.4	-22.2

Note: Base simulation assumes State Pension Age of 65 and value of contributory pension of 180 per week

Private savings responses to the policy counterfactuals (Table 9) are the product of two considerations:

1. The need for private saving is reduced by increases in the generosity of state retirement benefits; in this case state benefits act as a substitute for private saving.
2. The need for private saving is reduced as the age of retirement increases.

These considerations work in opposite directions when increases in state pension age are coupled with higher benefits. Changes in the State Pension Age operate at

the margin, while changes in levels of benefit affect all pensioners. As a result, the latter effect tends to dominate.

8. CONCLUSION

In this paper we describe how the National Institute's NIBAX model – a model well suited to the analysis of issues regarding household savings, including pensions, labour supply and asset allocation – has been adapted and calibrated to Irish circumstances. The model is an advanced tool, which has been tried and tested both in the policy sphere (in work for the UK Revenue authorities) and in top-level academic journals. Given the large infrastructural investment in building such models, adaptation via calibration represents a promising way of making this analytical approach available to a wider policy context.

Dynamic microsimulation models are essential in tracing the impact of changes in pension policy over the life course. The nature of these models is quite different from the more familiar static tax benefit models. The need to model decisions over an individual's lifetime means that the characterisation of the policy and labour market environment needs to be much more streamlined than in simpler "snapshot" or cross-section analyses.

The data and methods used to calibrate the model to Irish circumstances have been described, and the strategic simplifications used in characterising the tax and welfare systems, pension regimes and the labour market have been outlined. Calibration results indicate that the model does now capture many key features of the Irish system, including the patterns of labour market participation and of wage income over the life-course.

A brief policy analysis suggests that increasing the state retirement age in absence of other labour market reforms may deliver a smaller improvement in the government budget than is commonly assumed in the contemporary policy literature. Although the significance of this finding is difficult to overstate in the current policy environment, the analysis reported here only touches upon the subjects of study that are made possible by the model. We look forward to expanding upon this analysis in future work. One of the priorities for future analysis is the impact of changes in the tax treatment of superannuation contributions, along the lines proposed in the National Pensions Framework (Department of Social and Family Affairs, 2010). Analysis using the present model can help to examine the potential impact of such changes on pension coverage at different income levels, taking into account the ways in which both age and income tend to influence decisions regarding pensions.

REFERENCES

- Andreoni, J. (1989), 'Giving with impure altruism: applications to charity and Ricardian equivalence', *Journal of Political Economy* **97**, 1447–1458.
- Attanasio, O. P. & Rohwedder, S. (2003), 'Pension wealth and household saving: evidence from pension reforms in the United Kingdom', *American Economic Review* **93**, 1499–1521.
- Attanasio, O. P. & Weber, G. (1995), 'Is consumption growth consistent with intertemporal optimization?', *Journal of Political Economy* **103**, 1121–1157.
- Balcer, Y. & Sadka, E. (1986), 'Equivalence scales, horizontal equity and optimal taxation under utilitarianism', *Journal of Public Economics* **29**, 79–97.
- Barsky, R. B., Kimball, M. S., Juster, F. T. & Shapiro, M. D. (1997), 'Preference parameters and behavioral heterogeneity: an experimental approach in the health and retirement survey', *Quarterly Journal of Economics* **112**, 537–579.
- Blundell, R., Browning, M. & Meghir, C. (1994), 'Consumer demand and the life cycle allocation of household expenditure', *Review of Economic Studies* **61**, 57–80.
- 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.
- Connolly, S. & Gregory, M. (2008), 'Moving down: Women's part-time work and occupational change in Britain 1991-2001', *Economic Journal* **118**, F52–F76.
- Deaton, A. & Muellbauer, J. (1980), *Economics and Consumer Behavior*, Cambridge University Press, Cambridge.
- Department for Work & Pensions (2011), A sustainable State Pension: when the State Pension age will increase to 66. Pensions Act 2011 Impacts Annex A.
- Diamond, P. & Köszegi, B. (2003), 'Quasi-hyperbolic discounting and retirement', *Journal of Public Economics* **87**, 1839–1872.
- Disney, R., Emmerson, C. & Wakefield, M. (2007), Tax reform and retirement saving incentives: Evidence from the introduction of Stakeholder pensions in the UK. Institute for Fiscal Studies Working Paper WP19/07.
- Eckstein, Z. & Wolpin, K. I. (1989), 'Dynamic labour force participation of married women and endogenous work experience', *Review of Economic Studies* **56**, 375–390.
- Epstein, L. & Zin, S. (1989), 'Substitution, risk aversion, and the temporal behavior of consumption and asset returns: a theoretical framework', *Econometrica* **57**, 937–969.
- 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.
- Fernandez-Villaverde, J. & Krueger, D. (2006), Consumption over the life cycle: facts from Consumer Expenditure Survey Data. Working Paper, University of Pennsylvania.

- French, E. (2005), 'The effects of health, wealth, and wages on labour supply and retirement behaviour', *Review of Economic Studies* **72**, 395–427.
- Gallant, A. R. & Tauchen, G. E. (1996), 'Which moments to match?', *Econometric Theory* **12**, 657–681.
- Gourieroux, C., Monfort, A. & Renault, E. (1993), 'Indirect inference', *Journal of Applied Econometrics* **8**, S85–S118.
- Grossman, S. J. & Shiller, R. J. (1981), 'The determinants of the variability of stock market prices', *American Economic Review* **71**, 222–227.
- 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.
- Hall, R. E. (1988), Intertemporal substitution in consumption. NBER Working Paper W0720.
- Hansen, P. & Singleton, K. (1983), 'Stochastic consumption, risk aversion and the temporal behaviour of asset returns', *Journal of Political Economy* **91**, 249–266.
- Heckman, J. (1974), 'Life cycle consumption and labor supply: an explanation of the relationship between income and consumption over the life cycle', *American Economic Review* **64**, 188–194.
- Hicks, J. R. (1939), *Value and Capital*, Clarendon Press, Oxford.
- Keys, R. G. (1981), 'Cubic convolution interpolation for digital image processing', *IEEE Transactions on Acoustics, Speech, and Signal Processing* **29**, 1153–1160.
- Laibson, D. (1997), 'Golden eggs and hyperbolic discounting', *Quarterly Journal of Economics* **112**, 443–477.
- 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.
- Mankiw, N. G. (1985), 'Consumer durables and the real interest rate', *Review of Economics and Statistics* **62**, 353–362.
- Mankiw, N. G., Rotemberg, J. J. & Summers, L. H. (1985), 'Intertemporal substitution in macroeconomics', *Quarterly Journal of Economics* **100**, 225–251.
- 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.
- Mincer, J. & Ofek, H. (1982), 'Depreciation and restoration of human capital', *Journal of Human Resources* **17**, 3–24.
- Muellbauer, J. & van de Ven, J. (2004), Estimating Equivalence Scales for Tax and Benefits Systems. NIESR Working Paper 229.
- Myck, M. & Paull, G. (2004), The role of employment experience in explaining the gender wage gap. Institute for Fiscal Studies Working Paper WP04/16.

- Pickard, L., Comas-Herrera, A., i Font, J. C., Gori, C., di Maio, A., Rothgang, H. & Wittenberg, R. (2007), 'Modelling an entitlement to long-term care in Europe: Projections for long-term care expenditure to 2050.', *Journal of European Social Policy* **17**, 33–48.
- Press, W. H., Flannery, B. P., Teukolsky, S. A. & Vetterling, W. T. (1986), *Numerical Recipes: the art of scientific computing*, Cambridge University Press, Cambridge.
- Sefton, J. & van de Ven, J. (2004), 'Simulating household savings and labour supply: an application of dynamic programming', *National Institute Economic Review* **188**, 56–72.
- Sefton, J., van de Ven, J. & Weale, M. (2008), 'Means testing retirement benefits: fostering equity or discouraging savings?', *Economic Journal* **118**, 556–590.
- Stern, S. (1997), 'Simulation based estimation', *Journal of Economic Literature* **35**, 2006–2039.
- van de Ven, J. (2010), An empirical investigation of quasi-hyperbolic discounting. National Institute of Economic and Social Research Discussion Paper 355.
- van de Ven, J. (2011), 'A structural dynamic microsimulation model of household savings and labour supply', *Economic Modelling* **28**, 2054–2070.
- Waldfogel, J. (1998), 'The family gap for young women in the United States and Britain: can maternity leave make a difference?', *Journal of Labor Economics* **16**, 505–545.

A. List of Variables and Parameters

Preference Parameters

$U_{i,t}$	expected lifetime utility of household i at age t
E_t	expectations operator evaluated at age t
$1/\gamma$	coefficient of risk aversion
ϵ	elasticity of substitution between equivalised consumption and leisure
$\beta_{1,2}$	short-run quasi-hyperbolic discount parameters
δ	long-run (exponential) discount rate
$\zeta_{a,b}$	warm-glow bequest parameters

Ages

T	maximum possible duration of life
t_{SPA}	state pension age
t_{ER}	early retirement age
t_D	age at which all debts must be repaid

Consumption and Demographics

$c_{i,t}$	discretionary composite consumption of household i at age t
$l_{i,t}$	proportion of time spent in leisure of household i at age t
$l_{i,t}^{FT}$	leisure time of single adult, full-time employed
$l_{i,t}^{PT}$	leisure time of single adult, part-time employed
$l_{i,t}^{2FT}$	leisure time of adult couple, both full-time employed
$l_{i,t}^{FtPt}$	leisure time of adult couple, one full-time and one part-time employed
$l_{i,t}^{FtNe}$	leisure time of adult couple, one full-time and one not employed
$l_{i,t}^{PtNe}$	leisure time of adult couple, one part-time and one not employed
$\theta_{i,t}$	equivalence scale of household i at age t
$n_{i,t}^a$	number of adults in household i at age t
$n_{i,t}^c$	number of children in household i at age t
$\varphi_{j-t,t}$	probability of surviving to age j , given survival to age t

Wealth and Pensions

$w_{i,t}$	net liquid wealth of household i at age t
$w_{i,t}^+$	non-negative net liquid wealth balance of household i at age t
$w_{i,t}^s$	safe liquid assets of household i at age t
$w_{i,t}^r$	risky liquid assets of household i at age t
$nu_{i,t}$	proportion of liquid wealth invested in the risky asset by household i at age t
$w_{i,t}^{p/o}$	wealth held in personal / occupational pension of household i at age t
$r_{i,t}^s$	return on safe liquid assets of household i at age t
r_t^r	return on risky assets at age t
$r_t^{p/o}$	return to personal / occupational pension wealth at age t
π_{div}	proportion of liquid wealth lost upon marital dissolution prior to t_{SPA}
D_t	credit constraint on liquid net worth at age t
$\pi^{pc/oc}$	private contribution rate to personal / occupational pensions
$\pi_l^{p/o}$	lower bound on labour income to contribute to personal / occupational pensions
π_u^p	upper bound on labour income to contribute to personal pensions
$\pi_{ec}^{p/o}$	employer (and government) contribution rate to personal / occupational pensions
$\pi_a^{l/p/o}$	propn of liquid / personal pension / occupational pension wealth annuitised at t_{SPA}
π^{pe}	proportion of pension contributions that is tax exempt
π^{pt}	proportion of pension annuity income that is taxable
$\pi_{penalty_a}^p$	“account opening” cost on first contributions to personal pension
$\pi_{penalty_b}^p$	“investment cost” for choosing a contribution rate different from $\pi_{c_default}^p$

Income

$\tau_{i,t}$	net tax and benefit (disposable) income of household i at age t
$x_{i,t}$	non-property income of household i at age t
$y_{i,t}$	property income of household i at age t
$pc_{i,t}^{p/o}$	private contributions to private / occupational pensions of household i at age t
$g_{i,t}$	labour income of household i at age t
$p_{i,t}$	pension annuity income of household i at age t
$h_{i,t}$	latent wage of household i at age t
$m_{i,t}$	wage growth parameter of household i at age t
$\psi_{i,t}$	intertemporal persistence of earnings of household i at age t
$\kappa_{i,t}$	experience effect on earnings of household i at age t
$wo_{i,t}$	wage offer identifier
$p_{i,t}^{wo}$	probability of household i receiving a wage offer at age t

B. Maximisation of Expected Lifetime Utility

The intertemporal preference relation, as defined by equation (1), is:

$$U_t = \frac{1}{1-1/\gamma} \left\{ u_t^{1-1/\gamma} + E_t \left[\beta_1 \delta \left(\varphi_t u_{t+1}^{1-1/\gamma} + (1-\varphi_t) (\zeta_a + \zeta_b w_{t+1}^+)^{1-1/\gamma} \right) + \beta_1 \beta_2 \sum_{j=t+2}^T \delta^{j-t} \left(\varphi_{j-t,t} u_j^{1-1/\gamma} + (1-\varphi_{j-t,t}) (\zeta_a + \zeta_b w_j^+)^{1-1/\gamma} \right) \right] \right\} \quad (16)$$

where the individual subscripts i have been suppressed, $\varphi_t = \varphi_{1,t}$, and $u_t = u\left(\frac{c_{i,t}}{\theta_{i,t}}, l_{i,t}\right)$. Define:

$$W_t = \frac{1}{1-1/\gamma} \left[u_t^{1-1/\gamma} + E_t \sum_{j=t+1}^T \delta^{j-t} \left(\varphi_{j-t,t} u_j^{1-1/\gamma} + (1-\varphi_{j-t,t}) (\zeta_a + \zeta_b w_j^+)^{1-1/\gamma} \right) \right]$$

$$= \frac{1}{1-1/\gamma} u_t^{1-1/\gamma} + \delta E_t \left(\varphi_t W_{t+1} + \frac{(1-\varphi_t)}{(1-1/\gamma)} (\zeta_a + \zeta_b w_{t+1}^+)^{1-1/\gamma} \right)$$

Then:

$$U_t = \frac{1}{1-1/\gamma} u_t^{1-1/\gamma} + E_t \left[\frac{\beta_1 \delta}{(1-1/\gamma)} \left(\varphi_t u_{t+1}^{1-1/\gamma} + (1-\varphi_t) (\zeta_a + \zeta_b w_{t+1}^+)^{1-1/\gamma} \right) + \beta_1 \beta_2 \delta^2 \left(\varphi_{2,t} W_{t+2} + \frac{(1-\varphi_{2,t})}{(1-1/\gamma)} (\zeta_a + \zeta_b w_{t+2}^+)^{1-1/\gamma} \right) \right]$$

and if $\beta_1 = \beta_2 = 1$, then:

$$U_t = W_t = \frac{1}{1-1/\gamma} u_t^{1-1/\gamma} + \delta E_t \left(\varphi_t U_{t+1} + \frac{(1-\varphi_t)}{(1-1/\gamma)} (\zeta_a + \zeta_b w_{t+1}^+)^{1-1/\gamma} \right)$$

The value function describes expected lifetime utility as a function of the state variables at any given time, t , on the assumption that optimal choices are made and conditional on being alive at the start of time t .²² Define \hat{u}_t as the value of intra-temporal utility at time t , specified as a function of the state variables in time t , and conditional on the optimising decisions at time t . Similarly, define \hat{W}_t as the value of W_t , given the state variables at time t , evaluated at the sequence of optimising decisions for all $t \leq j \leq T$. Then the value function at time t is defined by:

²² Hence, the value function is a functional of the optimised decision stream.

$$V_t = \max_{\psi} \left\{ \frac{1}{1-1/\gamma} u_t^{1-1/\gamma} + E_t \left[\frac{\beta_1 \delta}{(1-1/\gamma)} \left(\varphi_t \hat{u}_{t+1}^{1-1/\gamma} + (1-\varphi_t) (\zeta_a + \zeta_b w_{t+1}^+)^{1-1/\gamma} \right) + \beta_1 \beta_2 \delta^2 \left(\varphi_{2,t} \hat{W}_{t+2} + \frac{(1-\varphi_{2,t})}{(1-1/\gamma)} (\zeta_a + \zeta_b w_{t+2}^+)^{1-1/\gamma} \right) \right] \right\}$$

where ψ is the set of decision alternatives available at time t .

To ensure that the value function is positive, it is recorded by the model in the form of the following monotonic transformation of equation (16):

$$Z_t = (1-1/\gamma) V_t^{1-1/\gamma} \quad (17)$$

B.1. Final period of life: $t = T$

Variables

In the final period of life, the household's decision is limited to their period specific consumption. As the opportunity to work or to invest in risky assets is not permitted, and as death in the following period is certain, there is no uncertainty associated with the maximisation problem in this period. Here, we have:

$$\begin{aligned} \text{state variables } t: & \quad w_t, p_t, n_t^a \\ \text{control variables } t: & \quad c_t \\ \text{state variables } t+1: & \quad w_{t+1} \end{aligned}$$

Value function

The value function, V_t , is defined by:

$$V_t = \max_{c_t} U_t \text{ subject to:} \quad (18)$$

$$w_{t+1} = w_t + \tau_t - c_t \quad (19)$$

$$c_t \leq c_t^{\max} \quad (20)$$

$$c_t^{\max} = w_t + \tau_t \quad (21)$$

where c_t^{\max} enforces the lower limit of zero considered for net liquid wealth where a household is subject to a certain probability of death.

Euler condition (if $\beta_1 = \beta_2 = 1$)

Euler conditions are only calculated if quasi-hyperbolic discounting is suppressed ($\beta_1 = \beta_2 = 1$). In this case, the Euler condition associated with the period T maximisation problem is:

$$\frac{\partial V_t}{\partial c_t} = \frac{\partial u_t}{\partial c_t} u_t^{-1/\gamma} - \delta \zeta_b (\zeta_a + \zeta_b w_{t+1}^+)^{-1/\gamma} \geq 0 \quad (22)$$

Backward induction intermediates

If solutions are based upon Euler conditions, then the following differential terms are also calculated for reference by the backward induction procedure that is used to evaluate solutions at $t < T$:

$$\frac{\partial V_t}{\partial w_t} = \frac{\partial V_t}{\partial c_t} \frac{\partial c_t^{\max}}{\partial w_t} + \delta \zeta_b (\zeta_a + \zeta_b w_{t+1})^{-1/\gamma} \frac{\partial w_{t+1}}{\partial w_t} \quad (23)$$

$$\frac{\partial c_t^{\max}}{\partial w_t} = (1 + ptr_t) \quad (24)$$

$$\frac{\partial w_{t+1}}{\partial w_t} = (1 + ptr_t) \quad (25)$$

$$\frac{\partial V_t}{\partial p_t} = \frac{\partial V_t}{\partial c_t} \frac{\partial c_t^{\max}}{\partial p_t} + \delta \zeta_b (\zeta_a + \zeta_b w_{t+1})^{-1/\gamma} \frac{\partial w_{t+1}}{\partial p_t} \quad (26)$$

$$\frac{\partial c_t^{\max}}{\partial p_t} = (1 - \pi^{pt} + \pi^{pt} mpy_t) \quad (27)$$

$$\frac{\partial w_{t+1}}{\partial p_t} = (1 - \pi^{pt} + \pi^{pt} mpy_t) \quad (28)$$

where ptr_t denotes the post-tax return to savings received by the household in period t , and mpy_t denotes marginal post-tax income ($\partial \tau / \partial x$).

If quasi-hyperbolic preferences are considered and $\beta_2 \neq 1$, then $\hat{X}_t = E_t \left(\frac{1}{(1-1/\gamma)} (\zeta_a + \zeta_b w_{t+1}^+)^{1-1/\gamma} \right)$ and \hat{u}_t are stored separately to permit evaluation of the value function in period $t = T - 1$. Otherwise, if $\beta_1 \neq 1$ but $\beta_2 = 1$, then $\hat{Y}_t = \hat{u}_t + \delta E_t \left((\zeta_a + \zeta_b w_{t+1}^+)^{1-1/\gamma} \right)$ is calculated and stored.

B.2. From State Pensionable Age: $t_{SPA} \leq t < T$

Variables

During this period, a household can choose their consumption, labour supply, and the proportion of their liquid wealth that is invested in a risky asset, v_t . Here, we have:²³

state variables t : w_t, p_t, h_t, n_t^a

control variables t : c_t, v_t, l_t

state variables $t + 1$: $w_{t+1}, p_{t+1}, h_{t+1}, n_{t+1}^a$

²³ In period T , h_T is omitted from the decision problem.

Value function

$$V_t = \max_{c_t, v_t, l_t} \left\{ \frac{1}{1-1/\gamma} u_t^{1-1/\gamma} + E_t \left[\frac{\beta_1 \delta}{(1-1/\gamma)} \left(\varphi_t \hat{u}_{t+1}^{1-1/\gamma} + (1-\varphi_t) (\zeta_a + \zeta_b w_{t+1}^+) \right)^{1-1/\gamma} \right] + \beta_1 \beta_2 \delta^2 \varphi_t \hat{X}_{t+1} \right\} \quad (29)$$

$$w_{t+1} = \min \{ w_{t+1}^{\max}, \max \{ w_{t+1}^{\min}, \hat{w}_{t+1} \} \} \quad (30)$$

$$p_{t+1} = \min \{ p_{t+1}^{\max}, \hat{p}_{t+1} \} \quad (31)$$

$$h_{t+1} = \min \{ h_{t+1}^{\max}, \max \{ h_{t+1}^{\min}, H(t, h_t, n_t^a, n_{t+1}^a, l_t, \omega_t) \} \} \quad (32)$$

$$n_{t+1}^a = N^a(t, n_t^a, \epsilon_t^a) \quad (33)$$

$$0 \leq v \leq 1 \quad (34)$$

$$c_t \leq c_t^{\max} \quad (35)$$

$$\hat{w}_{t+1} = w_t + \tau_t - c_t \quad (36)$$

$$\hat{p}_{t+1} = \left(\frac{\pi^s + (1-\pi^s)(n_{t+1}^a - 1)}{\pi^s + (1-\pi^s)(n_t^a - 1)} \right) p_t \quad (37)$$

$$c_t^{\max} = w_t + \tau_t^{\min} - w_{t+1}^{\min} \quad (38)$$

where $H(\cdot)$ denotes the intertemporal evolution of a household's latent wage (as defined by equation (10)), and $N^a(\cdot)$ defines the intertemporal development of relationship status (as defined by equation (11)). We do not impose an upper bound on the loss that may be incurred when investing in the risky asset. Two assumptions ensure that net liquid wealth never falls below the maximum debt that is considered for analysis. First, we assume that the consumption decision is subject to an upper bound, c_t^{\max} , to limit the probability that the upper bound on debt will be breached.²⁴ Second, given the upper limit on consumption, we assume that the government provides an income top-up to enforce the lower bound on net liquid wealth (as is implied by the specification of the intertemporal evolution of wealth defined by equation (30)).

Equation (30) also makes clear that we assume that a 100% wealth tax is levied on any wealth accrued beyond the maximum threshold w_{t+1}^{\max} . Similar assumptions are made in relation to the pension annuity and wage state variables, so that the evaluation of expected utility does not require grid extrapolations. We avoid extrapolating outside of the state-space defined by the grids considered for analysis, as we have found this to be an important source of error in previous work. The upper bound on wealth

²⁴ Hence τ_t^{\min} is set so that the probability of $\tau_t < \tau_t^{\min}$ is small. In practice, we evaluate τ_t^{\min} on the basis of the worst case scenario implied by the abscissae of the Gaussian quadrature that is used to evaluate expectations in the model.

by age is determined endogenously by the model, based on the grid limits that are assumed for labour income. The user should consequently specify an upper bound on labour income that is sufficiently high to capture extremes observed in practice.

Euler conditions (if $\beta_1 = \beta_2 = 1$)

A local optimum over the decision set (c_t, v_t) is calculated for each discrete labour option, and the decision set (c_t, v_t, l_t) is then selected to maximise the associated value function. Solution for each (c_t, v_t) combination is based upon the Euler conditions defined by:

$$\frac{\partial V_t}{\partial c_t} = \frac{\partial u_t}{\partial c_t} u_t^{-1/\gamma} + \delta E_t \left\{ \left[\varphi_t \frac{\partial V_{t+1}}{\partial w_{t+1}} + (1 - \varphi_t) \zeta_b (\zeta_a + \zeta_b w_{t+1}^+)^{-1/\gamma} \Phi_{t+1} \right] \frac{\partial w_{t+1}}{\partial c_t} \right\} \geq 0 \quad (39)$$

$$\frac{\partial w_{t+1}}{\partial c_t} = \begin{cases} -1 & \text{if } w_{t+1}^{\min} \leq \hat{w}_{t+1} \leq w_{t+1}^{\max} \\ 0 & \text{otherwise} \end{cases} \quad (40)$$

$$\frac{\partial V_t}{\partial v_t} = \frac{\partial V_t}{\partial c_t} \frac{\partial c_t^{\max}}{\partial v_t} + \delta E_t \left\{ \left[\varphi_t \frac{\partial V_{t+1}}{\partial w_{t+1}} + (1 - \varphi_t) \zeta_b (\zeta_a + \zeta_b w_{t+1}^+)^{-1/\gamma} \Phi_{t+1} \right] \frac{\partial w_{t+1}}{\partial v_t} \right\} \leq 0 \quad (41)$$

$$\frac{\partial c_t^{\max}}{\partial v_t} = mpy_t^{\min} (r_t^{r, \min} - r_t^s) w_t \quad (42)$$

$$\frac{\partial w_{t+1}}{\partial v_t} = \begin{cases} mpy_t (r_t^r - r_t^s) w_t & \text{if } w_{t+1}^{\min} \leq \hat{w}_{t+1} \leq w_{t+1}^{\max} \\ 0 & \text{otherwise} \end{cases} \quad (43)$$

where Φ_{t+1} is a dummy variable that is equal to one when $w_{t+1} \geq 0$, and is zero otherwise. mpy_t^{\min} and $r_t^{r, \min}$ define the values of mpy and r^r used to calculate τ^{\min} .²⁵ Note that the Euler conditions used to identify the (locally) optimal values of c and v make reference to $\partial V_{t+1}/\partial w_{t+1}$; this term is evaluated by interpolation with reference to the model solutions obtained for period $t + 1$.

Backward induction intermediates

Where Euler conditions are considered then the following terms are calculated and stored:

$$\frac{\partial V_t}{\partial w_t} = \frac{\partial V_t}{\partial c_t} \frac{\partial c_t^{\max}}{\partial w_t} + \delta E_t \left\{ \left(\varphi_t \frac{\partial V_{t+1}}{\partial w_{t+1}} + (1 - \varphi_t) \zeta_b (\zeta_a + \zeta_b w_{t+1}^+)^{-1/\gamma} \Phi_{t+1} \right) \frac{\partial w_{t+1}}{\partial w_t} \right\} \quad (44)$$

²⁵ Where c_t is not bound, then the evaluation of $\partial c_t^{\max}/\partial v_t$ does not influence equation (41), as $\partial V/\partial c = 0$.

$$\frac{\partial c_t^{\max}}{\partial w_t} = (1 + ptr_t^{\min}) \quad (45)$$

$$\frac{\partial w_{t+1}}{\partial w_t} = \begin{cases} (1 + ptr_t) & \text{if } w_{t+1}^{\min} \leq \hat{w}_{t+1} \leq w_{t+1}^{\max} \\ 0 & \text{otherwise} \end{cases} \quad (46)$$

$$\begin{aligned} \frac{\partial V_t}{\partial p_t} = & \frac{\partial V_t}{\partial c_t} \frac{\partial c_t^{\max}}{\partial p_t} + \delta E_t \left\{ \left(\varphi_t \frac{\partial V_{t+1}}{\partial w_{t+1}} + (1 - \varphi_t) \zeta_b (\zeta_a + \zeta_b w_{t+1}^+)^{-1/\gamma} \Phi_{t+1} \right) \frac{\partial w_{t+1}}{\partial p_t} + \right. \\ & \left. + \varphi_t \frac{\partial V_{t+1}}{\partial p_{t+1}} \frac{\partial p_{t+1}}{\partial p_t} \right\} \end{aligned} \quad (47)$$

$$\frac{\partial c_t^{\max}}{\partial p_t} = (1 - \pi^{pt} + \pi^{pt} mpy_t^{\min}) \quad (48)$$

$$\frac{\partial w_{t+1}}{\partial p_t} = \begin{cases} (1 - \pi^{pt} + \pi^{pt} mpy_t) & \text{if } w_{t+1}^{\min} \leq \hat{w}_{t+1} \leq w_{t+1}^{\max} \\ 0 & \text{otherwise} \end{cases} \quad (49)$$

$$\frac{\partial p_{t+1}}{\partial p_t} = \begin{cases} \left(\frac{\pi^s + (1 - \pi^s)(n_{t+1}^a - 1)}{\pi^s + (1 - \pi^s)(n_t^a - 1)} \right) & \text{if } \hat{p}_{t+1} \leq p_{t+1}^{\max} \\ 0 & \text{otherwise} \end{cases} \quad (50)$$

where ptr_t^{\min} denotes the post-tax rate of return to net liquid wealth used to evaluate the upper bound imposed on c_t .

As above, if quasi-hyperbolic preferences are considered and $\beta_2 \neq 1$, then \hat{u}_t and

$$\hat{X}_t = E_t \left(\varphi_t \hat{W}_{t+1} + \frac{(1 - \varphi_t)}{(1 - 1/\gamma)} (\zeta_a + \zeta_b w_{t+1}^+)^{1-1/\gamma} \right)$$

are stored separately to permit evaluation of the value function in period $t - 1$. Otherwise, if $\beta_1 \neq 1$ but $\beta_2 = 1$, then

$$\hat{Y}_t = \hat{u}_t + \delta E_t \left((1 - 1/\gamma) \varphi_t \hat{W}_{t+1} + (1 - \varphi_t) (\zeta_a + \zeta_b w_{t+1}^+)^{1-1/\gamma} \right)$$

is calculated and stored.

B.3. Period prior to State Pensionable Age: $t = t_{SPA} - 1$

Variables

During this period, a household can choose their consumption, labour supply, private pension contribution, and investment in a risky asset. Here, we have:

state variables t :	$w_t, w_t^o, w_t^p, h_t, n_t^a, n_t^k$
control variables t :	$c_t, v_t, \pi_t^{pc}, l_t$
state variables $t + 1$:	$w_{t+1}, p_{t+1}, h_{t+1}, n_{t+1}^a$

Value function

$$V_t = \max_{c_t, v_t, \pi_t^{pc}, l_t} \left\{ \frac{1}{1-1/\gamma} u_t^{1-1/\gamma} + E_t \left[\frac{\beta_1 \delta}{(1-1/\gamma)} \left(\varphi_t \hat{u}_{t+1}^{1-1/\gamma} + (1-\varphi_t) (\zeta_a + \zeta_b w_{t+1}^+) \right)^{1-1/\gamma} \right] + \beta_1 \beta_2 \delta^2 \varphi_t \hat{X}_{t+1} \right\} \quad (51)$$

$$w_{t+1} = \min \{ w_{t+1}^{\max}, \max \{ w_{t+1}^{\min}, \bar{w}_{t+1} \} \} \quad (52)$$

$$p_{t+1} = \min \{ p_{t+1}^{\max}, \bar{p}_{t+1} \} \quad (53)$$

$$h_{t+1} = \min \{ h_{t+1}^{\max}, \max \{ h_{t+1}^{\min}, H(t, h_t, n_t^a, l_t, n_{t+1}^a, \omega_t) \} \} \quad (54)$$

$$n_{t+1}^a = N^a(t, n_t^a, \epsilon_t^a) \quad (55)$$

$$0 \leq v_t \leq 1 \quad (56)$$

$$\pi_l^{pc} \leq \pi_t^{pc} \leq \pi_u^{pc} \quad (57)$$

$$c_t \leq c_t^{\max} \quad (58)$$

$$\hat{w}_{t+1} = w_t + \tau_t - c_t \quad (59)$$

$$\bar{w}_{t+1} = \min \{ \hat{w}_{t+1}, 0 \} + (1 - \pi_a^l) \max \{ 0, \hat{w}_{t+1} \} + (1 - \pi_a^o) w_{t+1}^o + (1 - \pi_a^p) w_{t+1}^p \quad (60)$$

$$\bar{p}_{t+1} = \chi(\pi_a^l \max \{ 0, \hat{w}_{t+1} \} + \pi_a^o w_{t+1}^o + \pi_a^p w_{t+1}^p) \quad (61)$$

$$w_{t+1}^o = (1 + r_t^o) w_t^o + (\pi^{oc} + \pi_{ec}^o) g_t \Phi_t'' \quad (62)$$

$$w_{t+1}^p = (1 + r_t^p) w_t^p + (\pi_t^{pc} + \pi_{ec}^p) (g_t - \pi_l^p) \Phi_t' \quad (63)$$

$$c_t^{\max} = w_t + \tau_t^{\min} + (1 - \pi_a^o) w_{t+1}^{o, \min} + (1 - \pi_a^p) w_{t+1}^{p, \min} - w_{t+1}^{\min} \quad (64)$$

where Φ_t' is a dummy variable that equals one when $\pi_l^p \leq g_t \leq \pi_u^p$ and zero otherwise, and Φ_t'' is a dummy variable that equals one when $\pi_l^o \leq g_t$ and zero otherwise.

Euler conditions (if $\beta_1 = \beta_2 = 1$)

A local optimum is calculated with respect to the decision set (c_t, v_t, π_t^{pc}) for each discrete labour option, and the decision set $(c_t, v_t, \pi_t^{pc}, l_t)$ is selected to maximise the associated value function. Solution for each (c_t, v_t, π_t^{pc}) combination is based upon the Euler conditions defined by:

$$\frac{\partial V_t}{\partial c_t} = \frac{\partial u_t}{\partial c_t} u_t^{-1/\gamma} + \delta E_t \left\{ \varphi_t \frac{\partial V_{t+1}}{\partial w_{t+1}} \frac{\partial w_{t+1}}{\partial c_t} + \varphi_t \frac{\partial V_{t+1}}{\partial p_{t+1}} \frac{\partial p_{t+1}}{\partial c_t} + (1 - \varphi_t) \zeta_b (\zeta_a + \zeta_b w_{t+1}^+)^{-1/\gamma} \Phi_{t+1} \frac{\partial w_{t+1}}{\partial c_t} \right\} \geq 0 \quad (65)$$

$$\frac{\partial w_{t+1}}{\partial c_t} = \begin{cases} 0 & \text{if } \bar{w}_{t+1} < w_{t+1}^{\min} \\ -1 & \text{if } w_{t+1}^{\min} \leq \bar{w}_{t+1} \leq w_{t+1}^{\max} \text{ and } \hat{w}_{t+1} < 0 \\ -(1 - \pi_a^l) & \text{if } w_{t+1}^{\min} \leq \bar{w}_{t+1} \leq w_{t+1}^{\max} \text{ and } \hat{w}_{t+1} \geq 0 \\ 0 & \text{if } w_{t+1}^{\max} < \bar{w}_{t+1} \end{cases} \quad (66)$$

$$\frac{\partial p_{t+1}}{\partial c_t} = \begin{cases} 0 & \text{if } \bar{p}_{t+1} \leq p_{t+1}^{\max} \text{ and } \hat{w}_{t+1} < 0 \\ -\chi \pi_a^l & \text{if } \bar{p}_{t+1} \leq p_{t+1}^{\max} \text{ and } \hat{w}_{t+1} \geq 0 \\ 0 & \text{if } p_{t+1}^{\max} < \bar{p}_{t+1} \end{cases} \quad (67)$$

$$\begin{aligned} \frac{\partial V_t}{\partial v_t} &= \frac{\partial V_t}{\partial c_t} \frac{\partial c_t^{\max}}{\partial v_t} + \delta E_t \left\{ \varphi_t \frac{\partial V_{t+1}}{\partial w_{t+1}} \frac{\partial w_{t+1}}{\partial v_t} + \varphi_t \frac{\partial V_{t+1}}{\partial p_{t+1}} \frac{\partial p_{t+1}}{\partial v_t} + \right. \\ &+ \left. (1 - \varphi_t) \zeta_b (\zeta_a + \zeta_b w_{t+1}^+)^{-1/\gamma} \Phi_{t+1} \frac{\partial w_{t+1}}{\partial v_t} \right\} \begin{matrix} \geq 0 \\ \leq 0 \end{matrix} \quad (68) \end{aligned}$$

$$\frac{\partial c_t^{\max}}{\partial v_t} = mpy_t^{\min} (r_t^{r,\min} - r_t^s) w_t \quad (69)$$

$$\frac{\partial w_{t+1}}{\partial v_t} = \begin{cases} 0 & \text{if } \bar{w}_{t+1} < w_{t+1}^{\min} \\ mpy_t (r_t^r - r_t^s) w_t & \text{if } w_{t+1}^{\min} \leq \bar{w}_{t+1} \leq w_{t+1}^{\max} \text{ and } \hat{w}_{t+1} < 0 \\ (1 - \pi_a^l) mpy_t (r_t^r - r_t^s) w_t & \text{if } w_{t+1}^{\min} \leq \bar{w}_{t+1} \leq w_{t+1}^{\max} \text{ and } \hat{w}_{t+1} \geq 0 \\ 0 & \text{if } w_{t+1}^{\max} < \bar{w}_{t+1} \end{cases} \quad (70)$$

$$\frac{\partial p_{t+1}}{\partial v_t} = \begin{cases} 0 & \text{if } \bar{p}_{t+1} \leq p_{t+1}^{\max} \text{ and } \hat{w}_{t+1} < 0 \\ \chi \pi_a^l mpy_t (r_t^r - r_t^s) w_t & \text{if } \bar{p}_{t+1} \leq p_{t+1}^{\max} \text{ and } \hat{w}_{t+1} \geq 0 \\ 0 & \text{if } p_{t+1}^{\max} < \bar{p}_{t+1} \end{cases} \quad (71)$$

$$\begin{aligned} \frac{\partial V_t}{\partial \pi_t^{pc}} &= \frac{\partial V_t}{\partial c_t} \frac{\partial c_t^{\max}}{\partial \pi_t^{pc}} + \delta E_t \left\{ \varphi_t \frac{\partial V_{t+1}}{\partial w_{t+1}} \frac{\partial w_{t+1}}{\partial \pi_t^{pc}} + \varphi_t \frac{\partial V_{t+1}}{\partial p_{t+1}} \frac{\partial p_{t+1}}{\partial \pi_t^{pc}} + \right. \\ &+ \left. (1 - \varphi_t) \zeta_b (\zeta_a + \zeta_b w_{t+1}^+)^{-1/\gamma} \Phi_{t+1} \frac{\partial w_{t+1}}{\partial \pi_t^{pc}} \right\} \begin{matrix} \geq 0 \\ \leq 0 \end{matrix} \quad (72) \end{aligned}$$

$$\frac{\partial c_t^{\max}}{\partial \pi_t^{pc}} = \frac{\partial \tau_t^{\min}}{\partial \pi_t^{pc}} \quad (73)$$

$$\frac{\partial w_{t+1}}{\partial \pi_t^{pc}} = \begin{cases} 0 & \text{if } \bar{w}_{t+1} < w_{t+1}^{\min} \\ \frac{\partial \tau_t}{\partial \pi_t^{pc}} + (1 - \pi_a^p) (g_t - \pi_l^p) \Phi_t' & \text{if } w_{t+1}^{\min} \leq \bar{w}_{t+1} \leq w_{t+1}^{\max} \text{ and } \hat{w}_{t+1} < 0 \\ (1 - \pi_a^l) \frac{\partial \tau_t}{\partial \pi_t^{pc}} + (1 - \pi_a^p) (g_t - \pi_l^p) \Phi_t' & \text{if } w_{t+1}^{\min} \leq \bar{w}_{t+1} \leq w_{t+1}^{\max} \text{ and } \hat{w}_{t+1} \geq 0 \\ 0 & \text{if } w_{t+1}^{\max} < \bar{w}_{t+1} \end{cases} \quad (74)$$

$$\frac{\partial p_{t+1}}{\partial \pi_t^{pc}} = \begin{cases} \chi \pi_a^p (g_t - \pi_l^p) \Phi_t' & \text{if } \bar{p}_{t+1} \leq p_{t+1}^{\max} \text{ and } \hat{w}_{t+1} < 0 \\ \chi \left[\pi_a^l \frac{\partial \tau_t}{\partial \pi_t^{pc}} + \pi_a^p (g_t - \pi_l^p) \Phi_t' \right] & \text{if } \bar{p}_{t+1} \leq p_{t+1}^{\max} \text{ and } \hat{w}_{t+1} \geq 0 \\ 0 & \text{if } p_{t+1}^{\max} < \bar{p}_{t+1} \end{cases} \quad (75)$$

Backward induction intermediates

Where Euler conditions are considered then the following terms are calculated and stored:

$$\begin{aligned} \frac{\partial V_t}{\partial w_t} = & \frac{\partial V_t}{\partial c_t} \frac{\partial c_t^{\max}}{\partial w_t} + \delta E_t \left\{ \varphi_t \frac{\partial V_{t+1}}{\partial w_{t+1}} \frac{\partial w_{t+1}}{\partial w_t} + \varphi_t \frac{\partial V_{t+1}}{\partial p_{t+1}} \frac{\partial p_{t+1}}{\partial w_t} + \right. \\ & \left. + (1 - \varphi_t) \zeta_b (\zeta_a + \zeta_b w_{t+1}^+)^{-1/\gamma} \Phi_{t+1} \frac{\partial w_{t+1}}{\partial w_t} \right\} \end{aligned} \quad (76)$$

$$\frac{\partial c_t^{\max}}{\partial w_t} = (1 + ptr_t^{\min}) \quad (77)$$

$$\frac{\partial w_{t+1}}{\partial w_t} = \begin{cases} 0 & \text{if } \bar{w}_{t+1} < w_{t+1}^{\min} \\ (1 + ptr_t) & \text{if } w_{t+1}^{\min} \leq \bar{w}_{t+1} \leq w_{t+1}^{\max} \text{ and } \hat{w}_{t+1} < 0 \\ (1 - \pi_a^l) (1 + ptr_t) & \text{if } w_{t+1}^{\min} \leq \bar{w}_{t+1} \leq w_{t+1}^{\max} \text{ and } \hat{w}_{t+1} \geq 0 \\ 0 & \text{if } w_{t+1}^{\max} < \bar{w}_{t+1} \end{cases} \quad (78)$$

$$\frac{\partial p_{t+1}}{\partial w_t} = \begin{cases} 0 & \text{if } \bar{p}_{t+1} \leq p_{t+1}^{\max} \text{ and } \hat{w}_{t+1} < 0 \\ \chi \pi_a^l (1 + ptr_t) & \text{if } \bar{p}_{t+1} \leq p_{t+1}^{\max} \text{ and } \hat{w}_{t+1} \geq 0 \\ 0 & \text{if } p_{t+1}^{\max} < \bar{p}_{t+1} \end{cases} \quad (79)$$

$$\begin{aligned} \frac{\partial V_t}{\partial w_t^p} = & \frac{\partial V_t}{\partial c_t} \frac{\partial c_t^{\max}}{\partial w_t^p} + \delta E_t \left\{ \varphi_t \frac{\partial V_{t+1}}{\partial w_{t+1}} \frac{\partial w_{t+1}}{\partial w_t^p} + \varphi_t \frac{\partial V_{t+1}}{\partial p_{t+1}} \frac{\partial p_{t+1}}{\partial w_t^p} + \right. \\ & \left. + (1 - \varphi_t) \zeta_b (\zeta_a + \zeta_b w_{t+1}^+)^{-1/\gamma} \Phi_{t+1} \frac{\partial w_{t+1}}{\partial w_t^p} \right\} \end{aligned} \quad (80)$$

$$\frac{\partial c_t^{\max}}{\partial w_t^p} = (1 - \pi_a^p) (1 + r_t^{p,\min}) \quad (81)$$

$$\frac{\partial w_{t+1}}{\partial w_t^p} = \begin{cases} (1 - \pi_a^p) (1 + r_t^p) & \text{if } w_{t+1}^{\min} \leq \bar{w}_{t+1} \leq w_{t+1}^{\max} \\ 0 & \text{otherwise} \end{cases} \quad (82)$$

$$\frac{\partial p_{t+1}}{\partial w_t^p} = \begin{cases} \chi \pi_a^p (1 + r_t^p) & \text{if } \bar{p}_{t+1} \leq p_{t+1}^{\max} \\ 0 & \text{otherwise} \end{cases} \quad (83)$$

If quasi-hyperbolic preferences are considered and $\beta_2 \neq 1$, then \hat{u}_t and

$$\hat{X}_t = E_t \left(\varphi_t \hat{W}_{t+1} + \frac{(1 - \varphi_t)}{(1 - 1/\gamma)} (\zeta_a + \zeta_b w_{t+1}^+)^{1-1/\gamma} \right)$$

are stored separately to permit evaluation of the value function in period $t - 1$. Otherwise, if $\beta_1 \neq 1$ but $\beta_2 = 1$, then

$$\hat{Y}_t = \hat{u}_t + \delta E_t \left((1 - 1/\gamma) \varphi_t \hat{W}_{t+1} + (1 - \varphi_t) (\zeta_a + \zeta_b w_{t+1}^+)^{1-1/\gamma} \right)$$

is calculated and stored.

B.4. Period $t < t_{SPA} - 1$

Variables

Here, we have:

$$\begin{aligned} \text{state variables } t: & \quad w_t, w_t^o, w_t^p, h_t, n_t^a, n_t^k \\ \text{control variables } t: & \quad c_t, v_t, \pi_t^{pc}, l_t \\ \text{state variables } t + 1: & \quad w_{t+1}, w_{t+1}^o, w_{t+1}^p, h_{t+1}, n_{t+1}^a, n_{t+1}^k \end{aligned}$$

Value function

$$V_t = \max_{c_t, v_t, \pi_t^{pc}, l_t} \left\{ \frac{1}{1 - 1/\gamma} u_t^{1-1/\gamma} + E_t \left[\frac{\beta_1 \delta}{(1 - 1/\gamma)} \left(\varphi_t \hat{u}_{t+1}^{1-1/\gamma} + (1 - \varphi_t) (\zeta_a + \zeta_b w_{t+1}^+) \right)^{1-1/\gamma} \right] + \beta_1 \beta_2 \delta^2 \varphi_t \hat{X}_{t+1} \right\} \quad (84)$$

$$w_{t+1} = \min \{ w_{t+1}^{\max}, \max \{ w_{t+1}^{\min}, \hat{w}_{t+1} \} \} \quad (85)$$

$$w_{t+1}^o = \min \{ w_{t+1}^{o, \max}, \hat{w}_{t+1}^o \} \quad (86)$$

$$w_{t+1}^p = \min \{ w_{t+1}^{p, \max}, \hat{w}_{t+1}^p \} \quad (87)$$

$$h_{t+1} = \min \{ h_{t+1}^{\max}, \max \{ h_{t+1}^{\min}, H(t, h_t, n_t^a, l_t, n_{t+1}^a, \omega_t) \} \} \quad (88)$$

$$n_{t+1}^a = N^a(t, n_t^a, \epsilon_t^a) \quad (89)$$

$$n_{t+1}^k = N^k(t, n_{t+1}^a, n_t^k, \epsilon_t^k) \quad (90)$$

$$0 \leq v_t \leq 1 \quad (91)$$

$$\pi_l^{pc} \leq \pi_t^{pc} \leq \pi_u^{pc} \quad (92)$$

$$c_t \leq c_t^{\max} \quad (93)$$

$$\hat{w}_{t+1} = w_t + \tau_t - c_t \quad (94)$$

$$\hat{w}_{t+1}^o = (1 + r_t^o) w_t^o + (\pi^{oc} + \pi_{ec}^o) g_t \Phi_t'' \quad (95)$$

$$\hat{w}_{t+1}^p = (1 + r_t^p) w_t^p + (\pi_t^{pc} + \pi_{ec}^p) (g_t - \pi_l^p) \Phi_t' \quad (96)$$

$$c_t^{\max} = w_t + \tau_t^{\min} - w_{t+1}^{\min} \quad (97)$$

Euler conditions (if $\beta_1 = \beta_2 = 1$)

A local optimum is calculated with respect to the decision set (c_t, v_t, π_t^{pc}) for each discrete labour option, and the decision set $(c_t, v_t, \pi_t^{pc}, l_t)$ is selected to maximise the associated value function. Solution for each (c_t, v_t, π_t^{pc}) combination is based

upon the Euler conditions defined by:

$$\frac{\partial V_t}{\partial c_t} = \frac{\partial u_t}{\partial c_t} u_t^{-1/\gamma} + \delta E_t \left\{ \varphi_t \frac{\partial V_{t+1}}{\partial w_{t+1}} \frac{\partial w_{t+1}}{\partial c_t} + (1 - \varphi_t) \zeta_b (\zeta_a + \zeta_b w_{t+1}^+)^{-1/\gamma} \Phi_{t+1} \frac{\partial w_{t+1}}{\partial c_t} \right\} \geq 0 \quad (98)$$

$$\frac{\partial w_{t+1}}{\partial c_t} = \begin{cases} -1 & \text{if } w_{t+1}^{\min} \leq \hat{w}_{t+1} \leq w_{t+1}^{\max} \\ 0 & \text{otherwise} \end{cases} \quad (99)$$

$$\frac{\partial V_t}{\partial v_t} = \frac{\partial V_t}{\partial c_t} \frac{\partial c_t^{\max}}{\partial v_t} + \delta E_t \left\{ \varphi_t \frac{\partial V_{t+1}}{\partial w_{t+1}} \frac{\partial w_{t+1}}{\partial v_t} + (1 - \varphi_t) \zeta_b (\zeta_a + \zeta_b w_{t+1}^+)^{-1/\gamma} \Phi_{t+1} \frac{\partial w_{t+1}}{\partial v_t} \right\} \geq 0 \quad (100)$$

$$\frac{\partial c_t^{\max}}{\partial v_t} = mpy_t^{\min} (r_t^{r,\min} - r_t^s) w_t \quad (101)$$

$$\frac{\partial w_{t+1}}{\partial v_t} = \begin{cases} mpy_t (r_t^r - r_t^s) w_t & \text{if } w_{t+1}^{\min} \leq \hat{w}_{t+1} \leq w_{t+1}^{\max} \\ 0 & \text{otherwise} \end{cases} \quad (102)$$

$$\begin{aligned} \frac{\partial V_t}{\partial \pi_t^{pc}} &= \frac{\partial V_t}{\partial c_t} \frac{\partial c_t^{\max}}{\partial \pi_t^{pc}} + \delta E_t \left\{ \varphi_t \frac{\partial V_{t+1}}{\partial w_{t+1}} \frac{\partial w_{t+1}}{\partial \pi_t^{pc}} + \varphi_t \frac{\partial V_{t+1}}{\partial w_{t+1}^p} \frac{\partial w_{t+1}^p}{\partial \pi_t^{pc}} + \right. \\ &+ \left. (1 - \varphi_t) \zeta_b (\zeta_a + \zeta_b w_{t+1}^+)^{-1/\gamma} \Phi_{t+1} \frac{\partial w_{t+1}}{\partial \pi_t^{pc}} \right\} \geq 0 \end{aligned} \quad (103)$$

$$\frac{\partial c_t^{\max}}{\partial \pi_t^{pc}} = \frac{\partial \tau_t^{\min}}{\partial \pi_t^{pc}} \quad (104)$$

$$\frac{\partial w_{t+1}}{\partial \pi_t^{pc}} = \begin{cases} \frac{\partial \tau_t}{\partial \pi_t^{pc}} & \text{if } w_{t+1}^{\min} \leq \hat{w}_{t+1} \leq w_{t+1}^{\max} \\ 0 & \text{otherwise} \end{cases} \quad (105)$$

$$\frac{\partial w_{t+1}^p}{\partial \pi_t^{pc}} = \begin{cases} (g_t - \pi_t^p) \Phi_t' & \text{if } \hat{w}_{t+1}^p \leq w_{t+1}^{p,\max} \\ 0 & \text{otherwise} \end{cases} \quad (106)$$

Backward induction intermediates

Where Euler conditions are considered then the following terms are calculated and stored:

$$\frac{\partial V_t}{\partial w_t} = \frac{\partial V_t}{\partial c_t} \frac{\partial c_t^{\max}}{\partial w_t} + \delta E_t \left\{ \varphi_t \frac{\partial V_{t+1}}{\partial w_{t+1}} \frac{\partial w_{t+1}}{\partial w_t} + (1 - \varphi_t) \zeta_b (\zeta_a + \zeta_b w_{t+1}^+)^{-1/\gamma} \Phi_{t+1} \frac{\partial w_{t+1}}{\partial w_t} \right\} \quad (107)$$

$$\frac{\partial c_t^{\max}}{\partial w_t} = (1 + ptr_t^{\min}) \quad (108)$$

$$\frac{\partial w_{t+1}}{\partial w_t} = \begin{cases} (1 + ptr_t) & \text{if } w_{t+1}^{\min} \leq \hat{w}_{t+1} \leq w_{t+1}^{\max} \\ 0 & \text{otherwise} \end{cases} \quad (109)$$

$$\frac{\partial V_t}{\partial w_t^p} = \delta E_t \left\{ \varphi_t \frac{\partial V_{t+1}}{\partial w_{t+1}^p} \frac{\partial w_{t+1}^p}{\partial w_t^p} \right\} \quad (110)$$

$$\frac{\partial w_{t+1}^p}{\partial w_t^p} = \begin{cases} (1 + r_t^p) & \text{if } \hat{w}_{t+1}^p \leq w_{t+1}^{p,\max} \\ 0 & \text{otherwise} \end{cases} \quad (111)$$

If quasi-hyperbolic preferences are considered and $\beta_2 \neq 1$, then \hat{u}_t and

$$\hat{X}_t = E_t \left(\varphi_t \hat{W}_{t+1} + \frac{(1 - \varphi_t)}{(1 - 1/\gamma)} (\zeta_a + \zeta_b w_{t+1}^+)^{1-1/\gamma} \right)$$

are stored separately to permit evaluation of the value function in period $t - 1$. Otherwise, if $\beta_1 \neq 1$ but $\beta_2 = 1$, then

$$\hat{Y}_t = \hat{u}_t + \delta E_t \left((1 - 1/\gamma) \varphi_t \hat{W}_{t+1} + (1 - \varphi_t) (\zeta_a + \zeta_b w_{t+1}^+)^{1-1/\gamma} \right)$$

is calculated and stored.

C. Data Sources

C.1. Household Budget Survey

The Household Budget Survey of 2004/2005 gathered data on the expenditure patterns, and socio-demographic composition, of just under 6,900 households. (A household was defined as a single person or group of people who regularly reside together in the same accommodation and who share the same catering arrangements.).

For the purposes of our model, the composite consumption good of interest is best interpreted as expenditure on all goods and services, including rent and mortgage interest on the family's residence, but excluding any mortgage capital repayment. This variable was constructed using the version of the HBS lodged at the ISSDA archive, and consumption expenditure classified by age group and partnership status (single/couple) was derived and used in calibration.

C.2. CSO Survey on Income and Living Conditions, 2005

The Survey on Income and Living Conditions (SILC) is the Irish element of the EU SILC. It is used both nationally and in an EU context as a tool for monitoring issues related to poverty and social inclusion. At the ESRI, the SILC data is reshaped into family units (single persons or married couples together with their dependent children) in order to provide the database for SWITCH, the ESRI tax-benefit model. In this form the data are well suited to provide a basis for the dynamic microsimulation model. A sub-population, excluding the self-employed and public sector employees, is defined, as explained in Section 1 of the paper. Special analyses of employment and disposable income are then used as part of the calibration process.

C.3. CSO Quarterly National Household Survey, 2005 (Quarter 2)

The QNHS is a very large scale survey (39,000 households) which gathers detailed information on employment and labour market participation. The version lodged at ISSDA was used to define the relevant subpopulation and conduct special analyses of participation in full-time and part-time employment by age, for use in the calibration process.

C.4. Living in Ireland Survey, 2000-2001

In order to estimate probabilities of transition from single to couple status, it was necessary to use data from a panel study. We used data from the last two waves of the Living in Ireland panel study (2000-2001) to estimate a logit regression for this purpose.

