

Have households under-saved for retirement? *

Rowena Crawford[†] Cormac O’Dea[‡]

31st October 2015

Abstract

There is a common perception that households are under-saving for retirement – indeed this has motivated many recent policy reforms in the UK. In this paper we use a structural model, specified to closely match the UK policy environment, to investigate this ‘fact’ for the cohort born in the 1940s. Using the model and administrative data linked with survey data from the English Longitudinal Study of Ageing, an optimal level of wealth is calculated for each household. This is then compared to the levels of wealth observed in the data. Our results show that the vast majority of households in this cohort actually hold far greater wealth than is necessary to maintain their living standards into and through retirement.

JEL Classification: D31, D91, E21, D12

Keywords: Lifecycle model; Wealth; Dynamic Programming; Savings

1 Introduction

There is a perception in the popular press, and one implicit in much of policy formation in many countries, that households have under-saved and continue to under-save for retirement. In this paper, using a structural model in which households can save in non-pension wealth and in a risky, tax-advantaged illiquid pension fund, we investigate whether this ‘fact’ is true for one particular cohort - those born in the 1940s who approached retirement in England through the 2000s. We assess whether the wealth holdings of these households can be rationalised as the outcome of optimal behaviour as predicted by a lifecycle model of consumption, saving and portfolio choice.

A result from Scholz et al. (2006) shows that in the US, where conventional wisdom also held that households under-save, levels of wealth accumulation for the vast majority of households in a slightly older cohort than those we study are more than sufficient to allow ‘optimal’ levels of consumption into and through retirement. Our contribution relative to that work is in two areas. First, the model in that paper allows households only to save in the form a safe asset. In reality, in the UK, the US and in most developed countries, substantially greater levels of wealth are held in the form of private pensions than are held in safe assets. Wealth held in private pensions differs

*Thanks to James Banks, Mariacristina De Nardi, Rachel Griffith and Hamish Low for very helpful comments and suggestions and to the Economic and Social Research Council (ESRC-NCRM Node ‘Programme Evaluation for Policy Analysis’ (PEPA) - ref: ES/I02574X/1; Centre for Microeconomic Analysis of Public Policy - ref: RES-544-28-50001) and to the Joseph Rowntree Foundation for funding. Correspondence to rowena_c@ifs.org.uk or cormac.odea@ifs.org.uk. Any errors are our own.

[†]Institute for Fiscal Studies

[‡]Institute for Fiscal Studies and University College London

from that held in cash for a number of reasons. Cash savings are flexible, bear no risk and are liquid though they earn only a low return, while private pension saving facilitates higher expected returns and a degree of tax advantage at the cost of illiquidity and riskiness. Second, the setting we study (the UK) is an environment where the public pension is not particularly generous by international standards (we illustrate this with a comparison to the (US) Social Security in Section 2). Substantially less retirement saving is therefore done ‘automatically’ - more is left to individuals (or employers on behalf of individuals). Under these conditions, there is a greater capacity for an under-saving problem to develop.

Our model incorporates, in a structural dynamic setting, each of three principal savings vehicles that facilitate consumption in retirement. These are the state pension system, cash savings and private pension holdings. Each of these three assets provide different opportunities for saving and consumption in retirement. The state pension system provides an income in retirement that is only weakly related to earnings during working life. Cash and pensions provide, as noted above, different levels of risk, return and liquidity.

Formally, our model predicts optimal cash and private pension holdings, conditional on state pension income. Households are assumed to be ex-ante heterogeneous on entering working life. This heterogeneity is over their lifetime productivity, their entitlements under the state pension system and the number, and timing of, their children. This heterogeneity implies that consumption and saving policy functions differ for each household. Simulating behaviour using each household’s own decision rules yields a prediction for optimal wealth. We then compare this optimal level of wealth for each household with that observed in the data.

Intuitively, modelled optimal consumption and saving decisions will be those which most effectively smooth consumption over the life cycle.¹ Those with less than optimal levels of wealth will be forced to reduce their (equivalised) consumption through retirement; those with more than optimal levels of wealth will have the capacity to increase their (equivalised) consumption in retirement or leave a bequest. The level of wealth that exactly facilitates consumption smoothing is optimal wealth. An alternative, more descriptive, manner of assessing whether wealth in retirement is adequate to facilitate some consumption smoothing is to assess whether (and if so by how much) income, including that coming from assets, is likely to fall on retirement. Analyses in that spirit have been carried out for the US (see e.g. Munnell et al. (2007, 2012)) and for the UK (see e.g. Banks et al. (2005) and Crawford and O’Dea (2012)).

The substantial data requirements involved in solving and simulating our model are met with the use of linked survey and administrative data. We use survey data from the English Longitudinal Study of Ageing (ELSA) linked to administrative data that allows us to observe respondents’ current wealth holdings, calculate their actual state pension entitlements, and estimate accurate measures of their entire lifetime earnings histories.

Our results show that for the vast majority of respondents (over 90%), observed wealth levels are greater than optimal levels. The excesses over optimal wealth tend to be large. The median surplus among those who have oversaved is over £225,000 which, if annuitised would facilitate consumption of approximately £7,000 a year more than optimal levels. Deficits, where they exist, are small - the median deficit is less than £40,000. The proportion oversaving for retirement is higher than found in the US by Scholz et al. (2006) for an older cohort using a model where households can

¹Formally, consumption in any period is chosen, subject to a borrowing constraint, to set marginal utility in that period equal to expected discounted marginal utility in the next.

either consume or save in cash but do not have the choice of saving into a pension fund. That we find a stronger result is perhaps surprising for two reasons. First, the inclusion in the model of a tax-advantaged, illiquid, higher return-earning (albeit risky) asset increases households' optimal wealth levels relative to those levels suggest by a one-asset model.² Second, the UK's system of state pensions is less generous than that in the US, in particular at higher levels of income. Therefore, all else equal, optimal *private* wealth would be higher in the UK than in the US.

This paper proceeds as follows. Section 2 outlines the model. Section 3 introduces the data that we use. Section 4 discusses the estimation and parameterisation of features of the model. Section 5 discusses our results. Section 6 concludes.

2 A model of wealth accumulation

In order to assess whether a household's observed wealth holdings are optimal given their lifetime earnings trajectories, we solve a lifecycle model of consumption, saving and portfolio choice. Details are given in this section. Briefly, its key features include decisions made, collectively, by the household; a careful specification of the tax, benefit and state pension system; uncertainty over employment, earnings, returns on a pension fund and mortality; and exogenous heterogeneity over the earnings process, state pension entitlements and fertility. For reasons we discuss in section 3 we consider couple households only.

2.1 Preferences and the economic environment

Preferences

Household utility in each period (one year in the model) is assumed to exhibit constant relative risk aversion in equivalised consumption, multiplied by the number of equivalised adults in the household:

$$U(c_t) = n_t \frac{\left(\frac{c_t}{n_t}\right)^{1-\gamma}}{1-\gamma}$$

where $\left(\frac{c_t}{n_t}\right)$ is household equivalised consumption and n_t is the number of equivalised adults in the household. The subscript t refers to the age of the man in the household.

Employment and earnings

Households are assumed to start working when the man in the couple reaches age 20 and can remain in the labour force until 65 (the age at which the public pension in the UK can be claimed), at which point they no longer work. In each period employment arrives with a certain probability and earnings, conditional on employment, are stochastic. The log earnings of a household if it is employed (\tilde{e}) are given by the sum of a fixed effect, a quadratic in age (t) and a stochastic process:

$$\ln \tilde{e}_{it} = \alpha_i + \delta_1^j t + \delta_2^j t^2 + u_{it} \quad (1)$$

²The discount rate that we use is equal to the interest rate on the risk free asset. The (risky) pension asset has an expected return substantially greater than the discount rate.

The coefficients on the quadratic in age vary with education group (j). There are three education groups - those with only compulsory education, those with more than compulsory but no post-secondary education and those with some post-secondary education.

The stochastic component of earnings follows a first order autoregressive process with normally distributed innovation:

$$u_{it} = \rho^j u_{it-1} + \xi_{it} \quad (2)$$

$$\xi \sim N\left(0, \sigma_j^2\right) \quad (3)$$

where the persistence of the process (ρ^j) and the variance of the innovation (σ_j^2) each vary by education.

Employment occurs in each period with probability π . Earnings (e_{it}) are therefore given by:

$$e_{it} = \begin{array}{ll} \tilde{e}_{it} & w.p. \quad \pi \\ 0 & w.p. \quad 1 - \pi \end{array}$$

π can be interpreted as the probability of spells of long-term unemployment (lasting at least a year) for all household members.

Private pensions

Households in the model can, each period, pay a proportion of their *pre-tax* income into a Defined Contribution (DC, i.e. 401k-style) pension. These funds earn an expected rate of return that is greater than that earned on cash, but is risky. The stock of wealth held in DC funds is modelled as illiquid in two senses. First, the wealth cannot be obtained until retirement. Second, there is a compulsion in the model, as there was in reality for decades in the UK³, to use three quarters of the fund to purchase an annuity. We assume that annuitisation happens as an employee turns 65.⁴ While the flow of income from an annuity is taxed, the non-annuitised quarter of the DC fund can be taken in cash tax free (both in the model and in reality), giving this form of saving a degree of tax advantage.

The evolution of the stock of wealth in the DC fund (DC) depends on flows into the fund (dc) and the return on the fund in each year (ϕ).

$$DC_{t+1} = (1 + \phi_{t+1})(DC_t + dc_t)$$

The return on DC funds is assumed to be normally distributed (with a truncation at -100%).

$$\phi_t \sim N(\bar{\phi}, \sigma_\phi^2)$$

Each household observed in a given year earns the same return.

Annuity rates (q) are assumed to be actuarially fair after a proportion (l) has been deducted to meet the administrative costs (including profits) associated with the provision of annuities. Private pension income at each age post-retirement (pp_t) is therefore given by an actuarially fair annuity rate multiplied by three quarters of value of funds held in DC wealth at age 65, scaled down to

³This compulsion was removed in March 2014.

⁴In reality, the current rules allow annuitisation from the age of 55.

allow for administrative load:

$$pp_t = q(0.75)DC_{65}(1 - z)$$

The annuity rates are calculated using the survival probabilities for men and women (which differ) born in 1945 - the middle year of birth for our cohort of interest.

Pensions in the model and reality differ in two important ways. First, in reality, some contributions to pension funds are made by the employer rather than the employee. This does not, of course, imply that it is the employer and not the employee who ultimately bears the incidence of these payments. However, if the employer remits the payment to the fund (regardless of whether such a remittance represents an actual flow of funds or whether, in the case of a less than fully-funded scheme, it is simply a promise to pay a pension in the future), our earnings data will not capture this remuneration. We adjust our earnings measure to take this into account (discussed in Section 3.4, after we introduce our data). Second, there is no Defined Benefit (DB) pension in the model, while, in reality many in this cohort have large stocks of DB wealth. The optimal level of wealth that the model yields is that level which would have been optimal for households had they only had access to a DC fund, but had they been given their total remuneration (including employer pension contributions to DB and DC funds) to freely allocate between current consumption and saving.

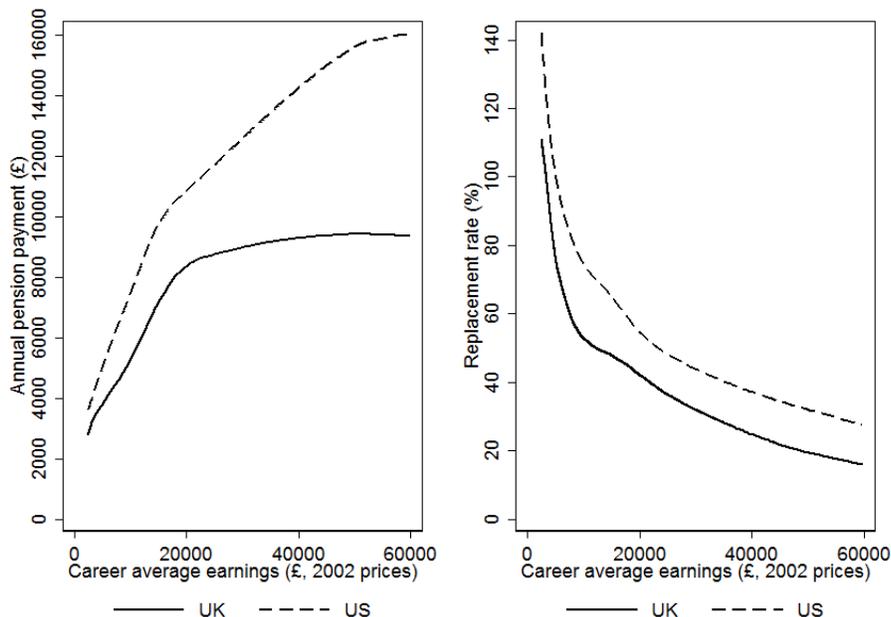
Systems of state pension(s)

The UK state pension system for the cohort we consider has two main components. The first is the Basic State Pension - a payment to those over the State Pension Age that depends on number of years worked but not on earnings in those years. The second is an earnings-related payment (the name and form of which has changed over time - it has been variously known as Graduated Retirement Benefit, the State Earnings Related Pension Scheme and the Second State Pension). Households in the model are assumed to know, from the start of working life, their future entitlement to these benefits.

It is worth describing briefly how the UK pension system differs from its counterpart in the US - Social Security. In the UK, the state pension system is related to earnings in a much weaker sense than is Social Security in the US. We have estimated, for our sample of households, the pension they would receive if they had earned (and paid taxes and social contributions) in the US.⁵ Figure 1 shows, in the left hand panel, the average UK state pension payment that households with various levels of lifetime average income can expect to receive, compared to what they would receive under Social Security. The figure shows that the level of the payment is higher across the entire distribution of lifetime earnings in the US, and is substantially higher for those with middle and higher levels of lifetime earnings. This implies (as shown in the right hand panel which shows what proportion of average lifetime earnings are replaced by the state in retirement) that the US system replaces a substantially greater proportion of working-life income than does the system in the UK.

⁵For the US system, we use the Social Security notch points and maximum insurable earnings for 2002, convert US dollar figures to pounds using the average exchange rate in that year (\$1 = £0.66), and index wages by UK earnings growth rather than US earnings growth. For the UK, we use actual entitlements calculated from administrative data and graph them against earnings estimated from that same data. In calculating entitlements we assume, for this figure, no 'contracting out'. 'Contracting out' is a device whereby individuals can divert some of their National Insurance contributions into a private sector pension fund and forgo some of their state pension entitlements.

Figure 1: Comparison of UK State Pensions with US Social Security



This figure differs from Figure ?? in a number of ways (though the overall pattern shown is the same). First, the UK pension system modelled here is that of 2002/2003 rather than 2012/13. Second the US Social Security notch points and maximum insurable earnings are for the earlier year. Third, the exchange rate used to convert dollars to pounds is the average in 2002 rather than 2012.

Taxes and Transfers

We specify an exogenous, time-invariant tax and benefit system that is based on the UK tax system in 2002/2003 - the year represented by our main data source. Household net income is a function (τ) of gross income. This function is described in detail in Appendix A. Here we give a brief overview. The modelled components of the tax and benefit system are income tax, National Insurance (the UK equivalent of the US Payroll Tax), Jobseekers' Allowance (a payment to the unemployed), Child Benefit and the Minimum Income Guarantee for pensioners. The function which returns net income:

$$y_t = \tau(e_t, ra_t, pp_t, sp_t, h_t, k_t, dc_t, t)$$

depends on household earnings (e_t), interest income (ra_t), private pension payments (pp) state pension payments (sp), number of adults still alive (h_t), number of dependent children (k_t), chosen contributions to the pension fund (dc_t) (since those attract tax relief) and finally on the age of the household (t) (the UK tax system taxes the elderly to a lesser extent than those of working age).

Jobseekers' Allowance is paid to unemployed households at a rate which depends on the number of adults and children in the household. Child Benefit is paid on the basis of the number of dependent children that they have and is paid at a more generous rate for the first child than for subsequent children.

Households aged over 60 are additionally entitled to an income- and asset-tested transfer that aims to insure them against destitution (similar to Supplemental Security Income in the US). This

was known in 2002 as the Minimum Income Guarantee (MIG).⁶ Entitlement to the MIG is based on current circumstances only and does not depend on a households' history of tax payments or national insurance contributions. The MIG simply tops net income up to a minimum level (f , which was £5,184 per year for singles and £7,790 for couples in 2002/03) so the benefit is withdrawn at an effective tax rate of 100% as private income increases.⁷ The formula for the MIG is (where y is net income)⁸:

$$mig = \max(0, f - y)$$

Intertemporal budget constraint

Cash assets earn an interest rate of r and evolve in the following manner for, respectively, an already-retired household, a retiring household and a pre-retirement household (for the case of the retiring household, recall that they will receive a cash lump sum to the value of a quarter of their pension fund value):

$$\begin{aligned} a_{t+1} &= (1+r)(a_t + y_t - c_t) & \forall t > 65 \\ a_{65} &= (1+r)(a_{64} + y_{64} - c_{64}) + (0.25)DC_{65}(1-z) \\ a_{t+1} &= (1+r)(a_t + y_t - c_t - dc_t) & \forall t \leq 64 \end{aligned}$$

Borrowing is not allowed.

Household heterogeneity

Households are ex-ante heterogeneous and therefore the optimisation problem that they face differs. They differ in both the fixed effect in their earnings process (α_i) and the number of children in the household at each age ($\{k_{it}\}_{t=20}^{100}$). These two elements of household heterogeneity are also present in the model of Scholz et al. (2006). Additionally, households are heterogeneous in their future, exact entitlements to state pensions ($\{sp_{it}(h_{it})\}_{t=65}^{100}$). The level of the state pension entitlement depends on who is alive at household age t - this is summarised by household composition (h_t). In what follows we summarise the 'type' of household i by $\theta_i = (\alpha_i, \{k_{it}\}_{t=20}^{100}, \{sp_{it}(h_{it})\}_{t=65}^{100})$.

Uncertainty

Four features of the model are uncertain. These are mortality, employment, the level of earnings (if employed) and the return on the pension fund. The probability distribution over mortality is summarised by the sequence of probabilities $\{s_t^m, s_t^f\}_{t=20}^{100}$ which give the probability of surviving

⁶In 2003, the MIG was changed in name and in form and is now known as Pension Credit - the benefit that is studied in Chapter 2. For a comprehensive discussion of the current UK State Pension system and its history see Bozio et al. (2010).

⁷This unattractive feature of the benefit was a key reason for a reform in 2003 - and most of the means-tested payment is now withdrawn at a rate of 40%.

⁸The concept of income assessed under the MIG, in reality, includes an imputed flow of income coming from non-housing assets. There is a small disregard, above which all non-housing assets are assumed to yield an annual income flow of 10% of the capital value. As our model does not distinguish between housing and non-housing assets, we do not include this imputed income from assets when assessing entitlement to MIG in our baseline model. However, we do also run a version of our model where we include an imputed income from wealth held in the safe asset when assessing entitlement, and our results are not materially affected.

to age t , conditional on being alive at age $t - 1$. The joint distribution over earnings, employment and pension fund returns (the joint distribution of which) is given by $F(\cdot)$. Households have rational expectations.

2.2 Household maximisation problem and value functions

We can now outline the household maximisation problem and value functions. The period of life where decisions are modelled starts at age 20 and all individuals are assumed to have died by the age of 100 at the latest. We do not model marriage or divorce - couples are assumed to start their productive life already married, and stay together until death. Our sample selection will be discussed later - however, it is worth noting here that we drop from our sample those who are separated or divorced by the time we observe them in our data and the small share of the sample who have never married. We now discuss in turn the optimisation problem facing retired households and working age households.

Retired household's problem Households in the model enter retirement when the male reaches the age of 65. Retirement here is associated with two distinct events. The first is withdrawal from the labour market. The second is the conversion of three quarters of the stock of wealth held in the DC fund into a life annuity, with the remaining quarter taken as a tax-free lump sum. The state variables that summarise the household's problem in retirement are age (t), cash assets (a), private pension income (pp) and household composition (h). This last variable takes a value of 1, 2, or 3 indicating, respectively, that both spouses are still alive, only the male is alive and only the female is alive. Households in retirement make a single choice in each period - their level of consumption (c).

The problem facing retired households with both spouses alive at time t is therefore (Appendix B gives the corresponding problem for households in which one spouse has died) :

$$\begin{aligned}
V_t(a_t, pp_t, h_t = 1; \theta_i) = \max_{c_t} & \left(u(c_t) \quad + \quad \beta s_{t+1}^m s_{t+1}^f V_{t+1}(a_{t+1}, pp_{t+1}, h_{t+1} = 1; \theta_i) \right. \\
& + \quad \beta s_{t+1}^m (1 - s_{t+1}^f) V_{t+1}(a_{t+1}, pp_{t+1}, h_{t+1} = 2; \theta_i) \\
& \left. + \quad \beta (1 - s_{t+1}^m) (s_{t+1}^f) V_{t+1}(a_{t+1}, pp_{t+1}, h_{t+1} = 3; \theta_i) \right) \\
s.t. \quad y_t & = \tau(e_t, r a_t, pp_t, sp_t, h_t, dc = 0, t) \\
a_{t+1} & = (1 + r)(a_t + y_t - c_t) \quad \forall t \neq 64 \\
a_{65} & = (1 + r)(a_{64} + y_{64} - c_{64}) + (1 - z)(0.25)DC_{65} \\
pp_t & = q(1 - z)(0.75)DC_{65}
\end{aligned} \tag{4}$$

where $u(\cdot)$ is an instantaneous utility function, $V_t(\cdot)$ is the value function in period t which is a function of the state variables and the type of the household, β is a geometric discount factor and the rest of the variables have been defined earlier in this section. That the value function (and therefore the maximisation problem and decision rules) faced by each household differs is indicated by the inclusion of θ_i as an argument.

Working age household's problem The state variables that summarise the household's problem during working life are age (t), cash assets (a), pension wealth (DC), earnings (e) and household composition (h).

At each age during working life households make two choices: their level of consumption (c_t) and their payments into a DC fund (dc_t). The balance of their resources is saved in a risk-free asset. The problem facing working households (again with both spouses alive at time t , with the corresponding value function for a single adult household shown in Appendix B) is:

$$\begin{aligned}
V_t(a_t, DC_t, e_t, h_t = 1; \theta_i) = \max_{c_t, dc_t} & \left(\begin{aligned}
& u(c_t) + \beta s_{t+1}^m s_{t+1}^f \int V_{t+1}(a_{t+1}, DC_{t+1}, e_{t+1}, 1; \theta_i) dF(\phi_{t+1}, e_{t+1} | e_t) \\
& + \beta s_{t+1}^m (1 - s_{t+1}^f) \int V_{t+1}(a_{t+1}, DC_{t+1}, e_{t+1}, 2; \theta_i) dF(\phi_{t+1}, e_{t+1} | e_t) \\
& + \beta (1 - s_{t+1}^m) (s_{t+1}^f) \int V_{t+1}(a_{t+1}, DC_{t+1}, e_{t+1}, 3; \theta_i) dF(\phi_{t+1}, e_{t+1} | e_t)
\end{aligned} \right) \\
s.t. \quad y_t & = \tau(e_t, ra_t, pp_t, sp_t, h_t, dc_t, t) \\
a_{t+1} & = (1 + r)(a_t + y_t - c_t - dc_t) \\
DC_{t+1} & = (1 + \phi_{t+1})(DC_t + dc_t)
\end{aligned}$$

2.3 Model solution and the simulation of consumption and savings behaviour

There is no analytical solution to the maximisation problem outlined which contains two choice variables (consumption and pension saving) and six state variables (age, cash, pension wealth, earnings, household composition and pension income). Full details on the methods used in solving the households' problem and simulating their behaviour are given in Appendix C. In short, we obtain decision rules numerically by solving the households' problem, first for the final period of life, storing the value functions for this period and then iterating backwards. After having obtained decision rules (the pension saving function and consumption function), we simulate wealth levels using these decision rules. This procedure requires us to use some realisations for the stochastic variables. Following Scholz et al. (2006), instead of using draws from a pseudo-random generator, we use actual realisations from the stochastic components drawn from the data on these households. That is, in each year in which we simulate behaviour using the calculated decision rules, we use households' actual employment status, realised earnings and the average return observed on DC assets in that year.

3 Data

The data used in this paper come from two sources. These are the English Longitudinal Study of Ageing (ELSA) and linked administrative National Insurance records for a subset of ELSA respondents.

ELSA is a biennial longitudinal survey, started in 2002/03, that contains a representative sample of the English private household population aged 50 and over. It is similar in form and purpose to ageing surveys in other countries, including the Health and Retirement Study (HRS) in the US and the Survey of Health, Ageing and Retirement in Europe (SHARE) in 20 European countries.

ELSA contains detailed data on demographics, labour market circumstances, health and, most importantly for our purposes, income and the level and composition of wealth holdings.

We complement the ELSA data with administrative data on ELSA respondents. Respondents were asked for their National Insurance number (equivalent to Social Security number) and permission to link to their history of National Insurance contributions. Data on these contributions allows us (subject to two particular issues that we discuss in more detail below) to obtain a detailed history of the earnings of ELSA respondents.

Almost 80% of ELSA respondents agreed to the linking of their survey records with their administrative data. Linked administrative and survey data of this kind has been used before in the US (by for, example, Gustman and Steinmeier (2005), Scholz et al. (2006) and Bound et al. (2010)), but has only recently been made available in the UK (see Bozio et al. (2011) for more detail on this data source).

3.1 Sample

We restrict our sample to couple households that contain at least one man born between 1940 and 1949. There are 1,615 couples of this type in the 2002/03 ELSA data. These individuals would be aged between 52 and 63 and therefore approaching the state pension age (of 65) when observed in 2002. We exclude households who refused permission to link to their administrative data, since for these households we cannot obtain lifetime earnings. We also exclude households where the man is observed in the NI data for fewer than 5 years. After applying these restrictions, 996 couples, or approximately 63% of couples, remain.

Table 1 provides descriptive statistics for our sample of households (for whom a long period of linked NI data is available) and all ELSA couples with men in the relevant cohort. The average age of men in our sample when they are observed in 2002 is just under 57, with wives on average being slightly younger. Nearly 70% of men in our sample reported still being in work, and only 15% defined themselves as retired. Home ownership was the norm for this cohort - in our sample over 90% of households own their home (either outright or still mortgaged).

A comparison of the descriptive statistics for our sample and the descriptives for all couples in ELSA with a man in the relevant cohort suggests that sample selection issues are not a major concern. An exception to this is with regard to self-employed individuals who are less likely to grant permission to link to their administrative records. Those with a history of self-employment are under-represented in our sample. We return to the issue of self-employment in our discussion of the calculation of measures of earnings in the Section 3.3.

3.2 Wealth measures in the English Longitudinal Study of Ageing

The model will yield an ‘optimal’ level for private wealth. ELSA data will be used to construct an empirical analogue to that measure. This is the sum of i) private pension wealth ii) housing wealth and iii) non-pension non-housing wealth.

The first of these, private pension wealth, is calculated differently for each of DC and DB pensions. DC pension wealth is the value of funds held - a question asked of respondents. We calculate DB pension *wealth* as the capital sum that would be required to purchase the projected stream of DB income (see Appendix F for more on this calculation). This future income is calculated from respondents’ reported accrual of rights to date, the rules of that scheme, and the projected

Table 1: Descriptive statistics for couples in ELSA with male born 1940-1950

<i>Individual characteristics</i>	Our sample		All	
	<i>Husband</i>	<i>Wife</i>	<i>Husband</i>	<i>Wife</i>
Mean age	56.8	54.1	56.9	51.6
Low education	35.3%	43.6%	36.5%	41.9%
Mid education	27.3%	32.1%	26.6%	31.1%
High education	37.3%	24.3%	36.8%	22.4%
Employee	60.6%	59.3%	55.8%	52.8%
Self-employed	10.7%	3.7%	14.9%	4.7%
Retired	14.9%	11.0%	14.3%	10.8%
Other	13.8%	25.9%	15.0%	31.7%
<i>Household characteristics</i>				
Owner occupier	89.6%		87.6%	
Median total income	£22,542		£22,142	
Median employment income	£17,862		£17,191	
Median asset income	£187		£182	
Median private pension wealth	£164,797		£166,100	
Median housing wealth	£120,000		£120,000	
Median non-pension, non-housing wealth	£32,140		£30,625	
<i>Sample size</i>	<i>996</i>		<i>1,583</i>	

labour market behaviour between the date of the survey and the date at which income will be received (see Banks et al. (2005)). Housing wealth is the sum of all gross housing wealth (on owner-occupied and other properties) less any mortgage debt. Non-pension non-housing wealth is the sum of net financial wealth (cash, stocks and shares less any outstanding non-mortgage debt), business wealth and physical wealth (land, antiques and collectibles). The sum of each of these three components (referred to below as ‘private wealth’) will be compared to modelled ‘optimal’ wealth.

The composition of wealth holdings (at the mean) among our cohort of couples is described in Table 2. The mean level of total private wealth is £451,777, of which approximately 40% is in each of private pension wealth and housing wealth and 20% in other forms of wealth. The table also includes a value for mean state pension wealth. This is calculated from state pension *income* in a manner analogous to how projected DB income is converted into DB wealth (see Appendix F).

Table 2: Composition of observed wealth holdings in 2002/03

	£	% of private wealth	% of total wealth
Private pension wealth	187,281	41%	33%
Housing wealth	171,020	38%	30%
Non-pension non-housing wealth	93,476	21%	16%
Total private wealth	451,777	100%	
State pension wealth	122,271		21%
Total wealth	574,048		100%

3.3 Administrative data

The national insurance (NI) data are the administrative records of individuals' national insurance contributions, and the dataset that is used by the UK government to establish individuals' rights to claim contributory benefits such as the state pension. We use this data to estimate ELSA respondents history of earnings. The NI records cover the years 1948 to 2003, though there are different levels of information for each of three sub-periods: 1948-1974, 1975-1996 and 1997 to 2003.

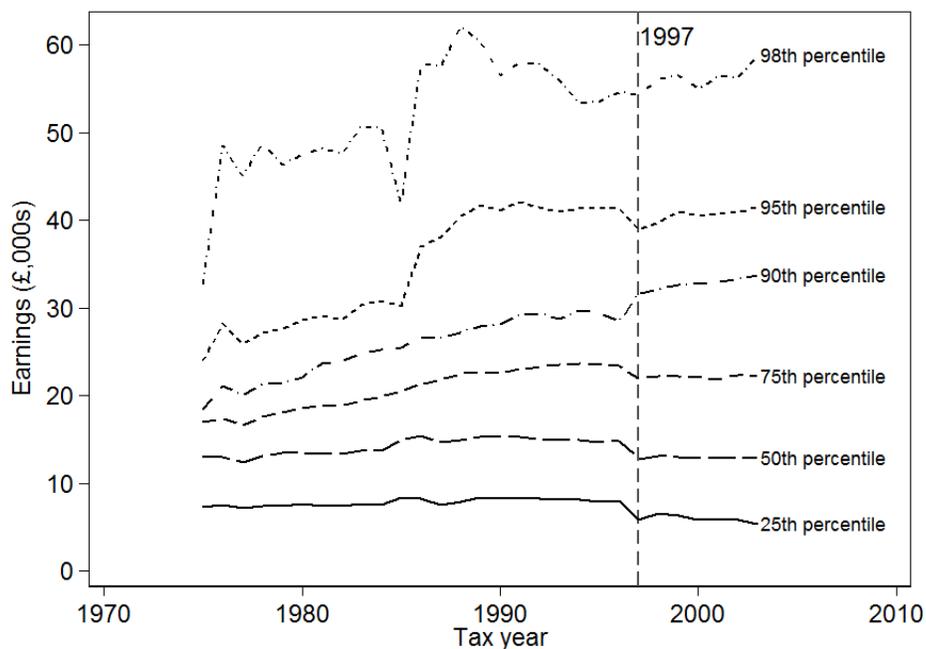
Taking the most recent period first, the NI records contain uncensored data on annual earnings as, in these years, employers were required to report the total earnings of their employees. For the middle period - years between 1975 and 1996 - the NI records contain data on employee National Insurance contributions. National Insurance payments in that interval were levied as a proportion of earnings between two values which are known as the Lower Earnings Limit (LEL) and the Upper Earnings Limit (UEL). For the period under consideration these values have been located at approximately the 8th and 80th percentile of the distribution of (positive) earnings. This data on NI contributions therefore allow us to calculate earnings, subject to right-censoring at the UEL and conditional on there being some earnings above the LEL. Prior to 1975 the NI records contain only data on the number of weeks that an individual earned above the LEL (and therefore paid NI contributions) and not the level of earnings. (This is because during this period the level of earnings was not relevant to the accrual of rights to state benefits or the state pension.)

To predict censored earnings in the years 1975 to 1996, we estimate the coefficients of a fixed-effect Tobit on earnings from 1975 to 2003 with the censoring point in each year up to 1996 equal to UEL (from 1997 there is no censoring). We use these coefficients to predict earnings for those who are affected by the censoring. The fixed-effect Tobit, when the length of the panel is fixed, is known to yield inconsistent results due to the incidental parameters problem (see Neyman and Scott (1948) for a general discussion of this problem). However Greene (2004) investigates, using Monte Carlo methods, its properties and finds that parameters of the fixed effects Tobit model are little affected by this problem even with panel of lengths substantially shorter than our panel (which has length 29). Further, Figure 2 shows a plot of selected quantiles of earnings through time using the censored and imputed data prior to 1997 and the uncensored data from 1997 onwards. This shows only a very small discontinuity in 1997.

To simulate earnings before 1975 we follow broadly the methodology used by Bozio et al. (2010). Using the NI data, we calculate an individual's mean earnings over the years 1975 to 2003 in which they are observed working, and then estimate potential previous years' earnings by adjusting for average economy-wide earnings growth and individual level earnings growth given their age, sex and education level. Having obtained this measure of potential earnings in each year, we then need to predict the years in which the individuals were working. The NI data records how many weeks the individual made NI contributions between 1948 and 1975. For men we assume they worked those weeks immediately prior to 1975 (therefore any periods not working were at the start of working life). To take account of the diminished propensity for women to work after having children, we assume that they worked those weeks from the point of leaving full-time education (therefore any periods not working were immediately prior to 1975). The combination of the estimates of potential earnings in a particular year for each individual and the years in which they were working yields our earnings estimates for years prior to 1975.

Household earnings are calculated by summing in each year the earnings for each individual in

Figure 2: Selected quantiles of earnings



the household.

The discussion above relates only to earnings in employment and not income earned in self-employment. National insurance payments are levied on self-employment income - but in a different manner than on earnings. As a result, the NI records enable us to identify years in which self-employment income was earned, but not the level of that income. Our measure of earnings therefore excludes income from self-employment. However, our results are not affected by the exclusion of the 13% of households with more than 5 years of self-employment income.

3.4 Employer contributions to pensions

Our earnings data do not capture payments that employers (rather than employees) remit to pensions and our model assumes all contributions into pension funds come from the employee. To take account of the fact that some pension rights will have been purchased by the employer on behalf of the employee we adjust upwards our earnings data to reflect these employer pension contributions. To do this, we first make an assumption about the proportion of pension wealth holdings which derive from employer contributions (κ). κ differs by whether the pension is a DB scheme or a DC scheme. We then scale up individuals' observed earnings by a constant factor x , where, if e_t^d represents the level of earnings that we observe in the data, x is such that if xe_t^d were saved in a pension each year, the resulting accumulated pension would amount to the part of observed wealth that is assumed to have arisen from employer pension contributions. The earnings data that we use to estimate the earnings process that feeds into our model is given by $e_t = (1 + x)e_t^d$. x is given by the following formula:

$$x = \frac{\kappa P_S}{\sum_t^{S-1} e_t^d (\prod_t^S (1 + \phi_t))}$$

where P_s is the pension wealth observed in survey period S and ϕ_t is the return on DC funds in the year the particular household is of age t . For defined benefit pensions we assume that $\kappa = 0.7$, while for DC pensions we assume $\kappa = 0.5$ (i.e. we assume 70% of observed DB pension wealth and 50% of observed DC pension wealth is the result of employer contributions).⁹

4 Estimation and parameterisation

In order to solve the model we must estimate (or set) a number of parameter values. The parameters we estimate are those of the earnings process, while parameters we set with reference to the literature include preference parameters (the coefficient of relative risk aversion and the discount factor), the rates of return on assets and the process for returns on DC funds and survival probabilities.

4.1 Estimation of earnings processes

To estimate the parameters of the earnings process we aggregate individual earnings histories into household earnings histories. We then divide households into three groups according to the education of the man in the couple (indexed by j). The parameters to be estimated parameters are $\{\alpha_i\}_{i=1}^N$ and $\{\beta_1^j, \beta_2^j, \rho^j, \sigma_j^2\}_{j=1}^3$.

To allow for measurement error in earnings, we augment equation (1) with an *iid* measurement error term m_{it} . The assumed data generating process for our earnings data is given in equations (5) to (7) (where, to keep notation concise, the fixed effect and earnings parameters for household i of education type j are contained in vector ψ_{ij} , and the respective variables are contained in vector x_{it}):

$$\ln e_{it}^{data} = \mathbf{x}_{it}\psi_{ij} + v_{it} \quad (5)$$

$$v_{it} = u_{it} + m_{it} \quad (6)$$

$$u_{it} = \rho u_{it-1} + \xi_{it} \quad (7)$$

The approach to estimation is a standard one (see, for example, Low et al. (2010)). It involves first running a fixed effects regression and estimating ψ_{ij} - the household fixed effect and quadratic in age. Residuals (r) are then obtained:

$$r_{it} = \ln e_{it}^{data} - x_{it}\hat{\psi}_{ij} \quad (8)$$

The parameters of the wage process are obtained by choosing those that minimise the distance between the empirical covariance matrix of differences in these residuals and the theoretical covariance matrix implied by equation (6).¹⁰ The theoretical variances and autocovariances of the differences in the stochastic component of earnings (Δv) are (the derivation of these is given in Appendix D):

⁹This is based on a comparison of the current mean employer contribution with the mean employee contribution in occupational DB and DC pension schemes. See ONS (2013).

¹⁰The covariance matrix of the *levels* of residuals could also be used, as in Guvenen (2009) or in Chapter 2 of this thesis.

$$\text{var}(\Delta v_t \Delta v_t) = \rho^2 \text{var}(\Delta u_{t-1}) + 2(1 - \rho)\sigma_\xi^2 + 2\sigma_m^2 \quad (9)$$

$$\text{cov}(\Delta v_t, \Delta v_{t+1}) = \rho \text{var}(\Delta u_t) - \sigma_\xi^2 - \sigma_m^2 \quad (10)$$

$$\text{cov}(\Delta v_t, \Delta v_{t+j}) = \rho \text{cov}(\Delta u_t \Delta u_{t+j-1}) \quad \forall j \geq 2 \quad (11)$$

Estimates of the parameters of the earnings process for each of three education groups are given in Table 3.

Table 3: Estimates of earnings process parameters

	Education group		
	Low	Middle	High
ρ	0.8468 (0.0838)	0.9727 (0.0153)	0.9527 (0.0025)
σ_ξ^2	0.0413 (0.0026)	0.0417 (0.0033)	0.0422 (0.0026)
σ_m^2	0.0024 (0.0021)	0.0029 (0.0026)	0.0066 (0.0016)

4.2 Parameterisation

This section discusses the model parameterisation - first the parameters that set the economic conditions faced by the cohort of interest and then their preference parameters.

4.2.1 Economic parameters

Return on the safe asset The safe asset in the model accounts for all non-pension holdings, of which cash and housing will be the largest components. The average real return on cash balances was 1.6% between 1952 and 2012 (see Table 1 of Barclays Capital (2012)). Using the Nationwide House Price Index for the years it is available, the average real increase in house prices between 1974 and 2013 was 2.8%. We use the midpoint between these two numbers - 2.2% as the return on the safe asset.

Distribution of pension fund returns We base the mean and standard deviation of pension fund returns on an index known as the ‘‘DCisions index’’¹¹. This is an index of total fund return that reflects the asset allocation decisions made by leading DC pension plans in their default investment strategies. This index provides information on returns stretching back to 1994. For years prior to 1994 when the DCisions index is not available, we estimate ϕ_t using the FTSE all-share index (on which data is available back to the early 1960s) and the ratio between the FTSE all-share index and the DCisions index over the period where both are available (1994 - 2010). We discuss how this is estimated in Appendix E. We use the mean and standard deviation of this time series in our model. These parameters are, respectively, $\bar{\phi} = 3.97\%$ and $\sigma_\phi = 13.8\%$.

Unemployment rate The period in our model is a year. Households will be observed as unemployed in our data only if both spouses are unemployed for at least a full financial year. When we

¹¹http://www.ftse.com/Indices/FTSE_DCisions_Index_Series/

observe very low earnings in the data the likelihood is that the individuals in the household have been unemployed for part of the year. The unemployment rate assumed in the model is 6.2%. This is the incidence in our data of a household between the ages of 25 and 50 having total earnings of less than £4,402 - the level of unemployment benefit payable to an unemployed couple in 2002/03.

Survival probabilities The survival probabilities, which differ for men and women, were obtained from the Office for National Statistics. These are survival probabilities for the cohort born in 1945.

Administrative load on annuities We assume that 10% of the value of the DC fund to be annuitised is taken by the provider to cover administrative costs and profits. We take this estimate from Murthi et al. (1999) who apply the methodology of Mitchell et al. (1999) to the UK.

4.2.2 Preference parameters

Discount factor The assumed level of patience is pivotal. One can rationalise almost any observed level of wealth as resulting from optimal behaviour by choosing a particular discount rate. Indeed, Samwick (1998), solves a simple lifecycle model separately for each household and estimates the distribution of discount rates by obtaining a discount rate that exactly rationalises each household's level of wealth. Our approach is to set the discount rate here equal to the risk free interest rate (so that the discount factor, β , is equal to $\frac{1}{1+r} = 0.978$). At this rate consumption and saving decisions will be made such that marginal utility in each period will be equal to expected marginal utility in the next period. This is due to the fact that, at an optimum and when borrowing constraints do not bind, consumption is chosen to satisfy the (risk-free asset) Euler equation (where the expectation operator is over mortality, investment returns, employment and earnings):

$$u'(c_t) = \beta(1+r)E[u'(c_{t+1})] = E[u'(c_{t+1})]$$

where the second equality holds due to the selection of β .¹² This choice lends a particular interpretation to our results. Optimal wealth will be such that expected marginal utility will be held constant - not falling as it would if β was greater than $1+r$, nor rising as it would be if β was less than $1+r$.

Coefficient of relative risk aversion We set the coefficient of relative risk aversion to 1.5, which is within the range estimated by Attanasio and Weber (1993) using UK data.

Equivalence scale We use the 'modified OECD equivalence scale' (see Anyaegbu (2010) for a discussion). The first adult in a household gets a weight of 1. Subsequent adults and children aged 14 and over get a weight of 0.5. Children aged 13 or younger get a weight of 0.3.

Summary These parameters of the model and values assigned are summarised in Table 4.

¹²The model has two assets - cash and the DC pension fund - and therefore two Euler equations. The second Euler equation is: $u'(c_t) = \beta E[(1 + \phi_t)u'(c_{t+1})]$.

Table 4: Parameterisation

Parameter	Symbol	Value/Source
Unemployment rate	π	6.2%
Return on safe asset	r	2.2%
Mean return on DC fund	$\bar{\phi}$	3.97%
Variance of return on DC fund	σ_{ϕ}^2	13.8%
Survival probabilities	s_t^m, s_t^f	ONS Life Tables
Administrative load on annuities	q	10%
Discount factor	β	$\frac{1}{1+r} = 0.978$
Coefficient of relative risk aversion	γ	1.5
Equivalence scale	n	Modified OECD scale

5 Results

The optimal levels of total net private wealth simulated by our model for each household in 2002/03 are summarised in Table 5. All quantities are expressed in 2002 prices. Median optimal wealth is estimated to be £76,990 (mean is £147,158). On average, across the whole sample, optimal levels of pension wealth are substantially higher than optimal levels of the safe asset. This is not surprising as pension wealth is tax favoured and exhibits higher average returns than cash, and while the assumptions of our model mean that it cannot be accessed until age 65, that is at most 14 years (and sometimes is as little as two years) in the future for these households.

Optimal private wealth is, of course, larger on average for households with higher lifetime earnings than for households with lower lifetime earnings. The median optimal wealth level among households in the lowest quintile of the lifetime earnings distribution is less than £643, several orders of magnitude lower than the £392,272 simulated for the median household in the highest lifetime earnings quintile. This last fact is the result of two main factors. First, those with higher lifetime earnings will have higher levels of lifetime consumption, and will therefore need to accumulate a greater stock of wealth in order to smooth marginal utility through retirement. Second, the optimal wealth simulated by the model is optimal *private* wealth, conditional on a household's state pension wealth. The state pension provides greater replacement of lifetime earnings for households in the lowest lifetime earnings quintile than households in the highest lifetime earnings quintile (recall Figure 1), leaving less of a role for private wealth to fund retirement. Table 6 illustrates that optimal private wealth is 2% of the discounted (by the risk-free interest rate) sum of lifetime earnings among households in the lowest earnings quintile, and 22% among households in the highest earnings decile. However, taking both optimal private wealth and state pension wealth together, this amounts to 24.7% of lifetime earnings across households in the lowest quintile, compared to 28.4% in the highest quintile.

5.1 Have households saved optimally?

Figure 3 gives a scatter plot of observed wealth against the optimal net wealth predicted by our model, for households with optimal and observed net wealth of less than £750,000. If observed wealth were exactly the same as optimal wealth, all the points on the scatter plot would lie on the 45-degree line. In fact, the majority of points lie above the 45-degree line, suggesting that most households in this cohort have saved more wealth than our model suggests is optimal. Just

Table 5: Optimal private net wealth, by quintile of lifetime earnings

	Optimal total wealth		Optimal “non-pension” wealth		Optimal pension wealth	
	Median	Mean	Median	Mean	Median	Mean
All	£76,990	£147,158	£27,708	£38,935	£21,339	£108,223
Lowest lifetime earnings quintile	£643	£9,706	£489	£6,567	£0	£3,139
2	£28,863	£38,900	£22,283	£26,654	£0	£12,245
3	£73,100	£82,072	£31,354	£38,688	£14,968	£43,384
4	£152,010	£167,678	£34,598	£39,887	£124,944	£127,791
Highest lifetime earnings quintile	£392,272	£438,126	£54,597	£83,040	£325,968	£355,086

Table 6: State pension wealth, lifetime earnings, and implied average lifetime savings rates, by quintile of lifetime earnings

	Mean state pension wealth	Mean lifetime earnings	Mean optimal private wealth / mean lifetime earnings	Mean (optimal private wealth + state pension wealth) / mean lifetime earnings
All	£122,271	£1,090,338	13.5%	24.7%
Quintile				
1 (Lowest)	£107,537	£483,363	2.0%	24.3%
2	£123,050	£793,739	4.9%	20.4%
3	£124,479	£970,357	8.5%	21.3%
4	£129,306	£1,219,221	13.8%	24.4%
5 (Highest)	£127,055	£1,988,058	22.0%	28.4%

7.9% of households in our sample are ‘under-saving’ in that they have observed wealth of less than that which our model suggests would be optimal for them.¹³ A univariate regression of observed on optimal wealth yields an R^2 of 31%. Despite the model not matching the levels of wealth accumulated by households, it can explain almost a third of the heterogeneity in these wealth levels (though this is less than the R^2 of 86% found by Scholz et al. (2006)).

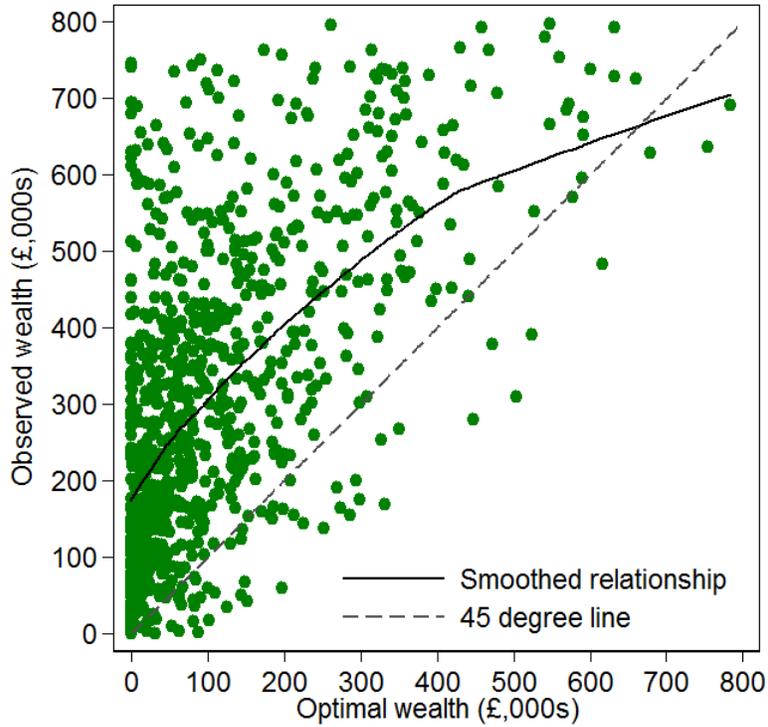
Table 7 sets out the proportion of under-savers in each lifetime earnings quintile in our sample. The median deficit among the 8% of under-savers (at less than £40,000) is substantially less than the median surplus among over-savers (at more than £225,000). The proportion of under-savers is higher among those at the extremes of the lifetime earnings distribution than among those in the middle.

¹³If we exclude the 13% of households which have some self-employment income in more than 5 years (the level of which we do not observe) the proportion under-saving increases to only 8.6%.

Table 7: Optimal wealth and the proportion of under-savers

	Median optimal wealth	Proportion under-saving	Median deficit (conditional)	Median surplus (conditional)	Median observed wealth
All	£76,990	7.9%	£38,747	£226,491	£324,135
Lowest lifetime earnings quintile	£643	9.5%	£8,461	£126,000	£119,340
2	£28,863	4.5%	£11,337	£188,841	£213,013
3	£73,100	6.5%	£27,608	£231,551	£293,185
4	£152,010	8.5%	£78,800	£283,214	£392,251
Highest lifetime earnings quintile	£392,272	10.6%	£94,401	£329,136	£690,328

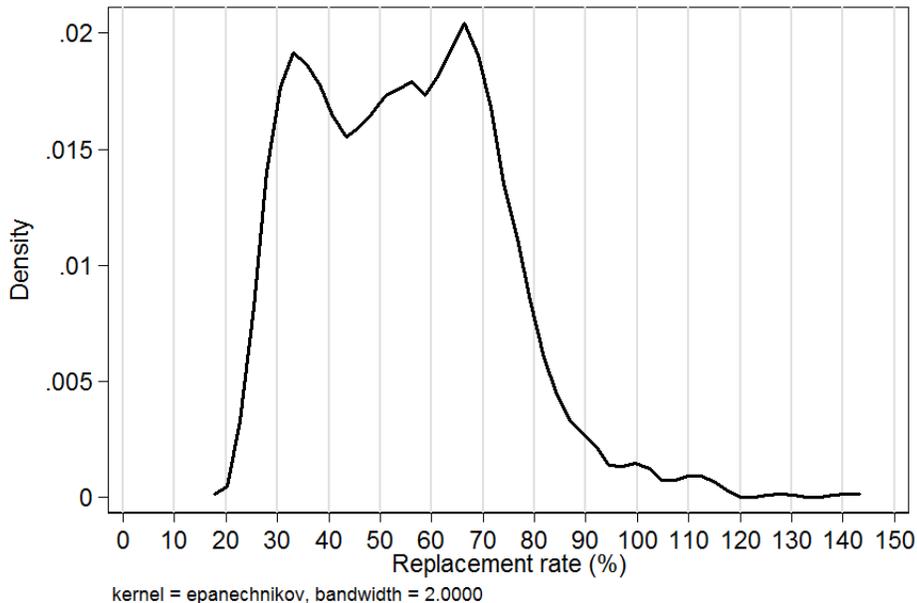
Figure 3: Scatter plot of observed and optimal wealth



Our data on lifetime earnings can, along with the model-predicted income in retirement, be used to calculate what the optimal replacement rate is for each household. Figure 4 shows the distribution of the ratio of optimal income at age 65 to average (uprated) earnings between the ages of 20 and 50. The numerator in this quantity includes (in addition to private pension income and interest income) state pension income - though we do not include means-tested benefit income. Most households have optimal replacement rates of between 30% and 70%.

These optimal replacement rates beg the question of how households can smooth expected marginal utility while only replacing (for many of them) half or less of their average earnings. There are four principal reasons for this. First, households in retirement no longer have any need to save for retirement and so need less income as a result. Second, for households with children a given level of marginal utility will be more expensive to obtain in working life than in retirement when children will (for most) be financially independent. Third, households in retirement pay lower rates of taxation than they do during working life. This is primarily due to the fact that National Insurance is only levied on earned income, and only below the state pension age. National Insurance rates are not small - the main rate paid by most employees was 10% for most of the period we consider. Tax rates are also lower in retirement because income falls (optimally, for reasons already noted), and therefore the progressivity of the tax system reduces average tax rates. This reduces the level of gross income needed in retirement to obtain a given quantity of net income. Fourth, once they reach the age of 60 households have access to a system of means-tested support that is more generous and comes with fewer conditions than welfare during working life. For households at certain points in the lifetime income distribution, this further reduces the need to have (private or state) gross pension in retirement.

Figure 4: Optimal replacement rates



Sensitivity We turn now to an assessment of whether our strong result is driven by some particular choice of parameters. We consider a number of sensitivity tests. In the first set, (panel (a) in Table 8), we consider alternative assumptions on some of our modelling inputs that are, arguably, as plausible as those in our baseline case. These are that retirement happens at 60 rather than at 65, that the coefficient of relative risk aversion is 3 (and not 1.5) and that asset-tests are applied to wealth when determining eligibility to the Minimum Income Guarantee. The results in each case are similar to those in our baseline specification with between 8.2% and 10.2% of households under-saving.

We next turn to assess the sensitivity of our results to three more extreme assumptions. The first sensitivity analysis reported in panel (b) is a comparison of modelled *total* private wealth with observed *non-housing* wealth. We do this to assess whether the very large returns some households in our cohort will have earned on their property purchases explain our results. In making the comparison between modelled total wealth and observed non-housing wealth, we are treating housing as both costless and of no value (the capital value is not available for consumption in retirement, nor does owning a house give an implied consumption value). Even under this extreme assumption, 75% of our sample have more than optimal wealth, while the surpluses are much larger than the deficits (£121,475 versus £43,701). The particular characteristics of housing and the large appreciation in house prices that have benefited homeowners in this cohort therefore may play a part in explaining our result but cannot fully account for it.¹⁴

We next turn to the implications of an assumption that households are fully patient (that is the

¹⁴Developing and solving a model where households can purchase housing in addition to saving in a pension and in cash would facilitate a number of interesting research questions. However, the computational constraints would be very large. Solving and simulating the model here is computationally expensive. Even with the speed of Fortran (which is an order of magnitude faster than, for example, Matlab) a single solution takes 3 minutes. A separate solution must be undertaken for each household and so the time taken to generate results for the entire sample is 50 hours. Our data cannot be taken from a secure data room and so we cannot make use of HPCs.

discount factor (β) is equal to 1). A discount rate of this level is much higher than those estimated by most of the literature on household discounting behaviour (see a review of this literature in Frederick et al. (2002)). Even in the absence of discounting, fewer than half of households (42.9%) are considered to have undersaved for retirement. We also test the sensitivity of our findings to timing effects. The simulations from our model suggest that household wealth holdings should increase until age 64 - after which households retire and convert the majority of their pension wealth into an income stream, and start to decumulate their holdings of the safe asset. Households in our sample, born in the 1940s, are aged between 51 and 63 when we observe them in the data. When we compare observed wealth holdings with optimal wealth at age 64, we still find that 70% of households hold more wealth than our model suggests they would optimally hold even on the eve of retirement.

Our final set of results considers the implication of removing the option to save in a pension and allowing households only the option of saving in the risk-free asset. When we do this while keeping all other aspects of the model unchanged, optimal wealth falls due, largely, to the lower return available on cash than on the pension fund and the proportion of under-savers falls to 4.5% (see the first row of panel (c)). Finally, in addition to shutting down the pension fund, we also assume that Defined Benefit pension income is exogenous and its ultimate level is known from the start of working life. The question the model answers here is how much should households save, over and above their DB and state pension entitlements. This makes the model solved closer to that in Scholz et al. (2006), where optimal saving is conditional on the stream of DB income. Here we also change our preference parameters and interest rate to match those chosen in that paper. In particular, we set the interest rate to 4%, the discount factor (β) to 0.96 and the coefficient of relative risk aversion (γ) to 3.¹⁵ In this case the comparison is between optimal wealth and observed wealth excluding that held in Defined Benefit pensions. Here, we find 12.5% of households are under-saving.

5.2 Discussion

The model set out here, while explaining 31% of the heterogeneity in wealth holdings, clearly cannot rationalise the *level* of wealth held by most individuals in this cohort. Nevertheless, many of the features that are sometimes included in lifecycle models that we do not include would, we argue, further reduce optimal wealth and strengthen the result we find. We discuss three of these in turn (non-separabilities between consumption and leisure in the utility function, a changing price of consumption in retirement and inaccurate expectations over life-expectancies) before turning to a number of potential partial explanations for the greater than optimal saving behaviour that we observe.

First, our utility function is separable in consumption and leisure. The evidence strongly rejects separability and finds that consumption and leisure are substitutes (see, for example, Browning and Meghir (1991), Blundell et al. (1994)). Introducing this fact into the utility function will cause consumption to optimally fall on retirement as leisure increases. In this case, necessary wealth for retirement will be less than that found in our separable case. Therefore, our utility function, we would argue, is biased against finding the result that we do. Non-separabilities between consumption

¹⁵One remaining difference between our model and that of Scholz et al. (2006) is that, in the latter, agents are assumed to face out-of-pocket medical expenses. We do not include such a feature in our model, largely because households in the UK are insured against such expenses by the National Health Service.

Table 8: Sensitivity analysis

	Median optimal wealth	Mean optimal wealth	Prop. under-saving	Median deficit (conditional)	Median surplus (conditional)	R squared
Baseline	76,990	147,158	7.9%	38,747	226,491	0.3088
(a)						
Retire at 60	80,660	174,969	10.2%	58,385	207,752	0.2759
$\gamma = 3$	75,041	151,326	8.2%	33,725	223,293	0.2991
Asset test in MIG	110,424	166,167	9.4%	46,899	205,659	0.3373
(b)						
Excl. housing wealth	76,990	147,158	25.1%	43,701	121,475	0.3227
$\beta = 1$	300,946	335,886	42.9%	94,241	137,657	0.3795
Comparing wealth to optimal wealth at 64	153,609	258,260	28.8%	105,016	190,828	0.1906
(c)						
One-asset	53,322	79,622	4.5%	11,221	273,488	0.2628
One-asset; Exog. DB; Scholz et al. (2006) params	47,954	68,591	12.5%	23,703	138,565	0.3373

and leisure are one of the explanations given by Banks et al. (1998) for substantial drops in spending as household members retire.

Second, we assume a constant ‘price’ of consumption throughout life. Aguiar and Hurst (2005) argue that a given level of consumption can be obtained at a lower level of spending in retirement. Once individuals have stopped working, they have more time to shop around for better value and have the ability to substitute home production for the purchase of some goods and services. Once again, this sort of dynamic would, if included in the model, reduce optimal wealth and strengthen the result that we find.¹⁶

Third, turning to life expectancies, our model assumes that individuals have accurate expectations over their survival probabilities. If they over-estimated their life expectancies they would be tempted to save more, reflecting the importance of a longer period in retirement. However, the evidence suggests that individuals (at least those of the age that we study) underestimate their life-expectancy (Ludwig and Zimper (2013)) rather than overestimate it, and therefore incorporating such biased expectations would reduce rather than increase desired saving.

Why then, have individuals in our sample accumulated so much wealth? Explicit bequest motives have been found to be a driver of saving behaviour though their quantitative importance for households outside those with the highest levels of lifetime earnings has been questioned (see Kopczuk and Lupton (2007), De Nardi (2004), De Nardi et al. (2010)). We make two comments here. First, housing makes up the vast majority of bequeathable wealth (pension wealth typically cannot be bequeathed to heirs other than a spouse). And recall that even excluding housing wealth over three quarters of households are oversaving. Second, to the extent that bequest motives are an important driver of savings behaviour (and it would be foolish to rule out that prospect completely), policymakers might want to consider the wealth levels of the elderly differently if such wealth is

¹⁶Skinner (2007), however, has suggested that, while that argument may be persuasive in thinking about a 65 year old, it is less so in considering an 85 year old, whose mobility, and therefore capacity for home production and shopping around, is diminished.

held to afford heirs large windfalls than if it was intended for their own consumption in retirement. This perspective might be especially important at the moment in the UK (and almost certainly elsewhere) as the elderly have been protected to a much greater extent than younger households from the fiscal measures that have been enacted in response to the public finance pressures arising from the financial crisis (see Adam (2012)).

Another potential explanation is that much of the saving of those households with large Defined Benefit pensions may have been, to an extent, ‘forced’ through compulsory (or highly incentivised) membership of employer schemes. This is potentially the case, but two factors will limit the importance of the argument. First, any ‘forced’ saving could be undone elsewhere by accumulating less in non-pension wealth. Second, in results not shown, we find substantial surpluses and almost equal incidence of over-saving among those with no Defined Benefit pension wealth (88% have over-saved relative to 92% in the baseline case).

Finally, some of the additional wealth may have been accumulated to insure against large medical bills in retirement. This has been shown to be a quantitatively important motive for saving behaviour in the US (see De Nardi et al. (2010)). This channel, while almost certainly less important in the UK given that health care is provided free by the state, is likely to be still present as long-term social care is not provided free. It is not necessarily the case, however, that incorporating the risk of social care expenses will increase optimal wealth accumulation for all households. In the UK social care is paid for by the state for those with very low assets. This intense asset-testing will mitigate (and possibly, for some types of households, reverse) the desire to hold substantial wealth at older ages that arises from the risk of social care costs. Further analysis of the interplay between the risk of large social care costs and savings behaviour in an environment where asset-tested assistance is provided is a fruitful avenue for further research.

6 Conclusions

This paper solves a dynamic consumption and saving problem for each household in a sample of those approaching the UK state pension age. The model we solve allows households to smooth consumption across the lifecycle using the state pension system and through two different assets - a risk-free asset and a Defined Contribution (401k style) pension.

Households are observed to have substantially greater wealth in reality than is considered optimal by our model - even when those households are endowed with utility functions and preference parameters that bias us against finding this result. Our result is, in terms of the proportion found to be over-saving, stronger than that found in the context of the one-asset model of Scholz et al. (2006) applied to an older cohort in the US. This is despite the inclusion in our model of a high-return pension fund which increases households’ desired holdings of wealth, and the fact that the UK’s state pension system replaces a substantially smaller share of pre-retirement income than does Social Security in the US.

Our findings should, perhaps, mitigate the concerns of successive UK governments that the wealth holdings among those currently approaching the state pension age are inadequate. However our result is not necessarily generalisable to those at the early and middle stage of working life for whom an emerging savings deficit may exist. The coverage of generous Defined Benefit pensions is much lower among these younger cohorts than it is among our cohort of interest, their

saving rates are lower and they are accumulating housing wealth at later ages (see Hood and Joyce (2013)). Assessing the wealth holdings of these younger cohorts is an important avenue for future research. However, in light of our results, the more modest saving behaviour of younger cohorts when compared to that of their parents' generation may be of lesser concern than has been thought to date.

References

- ADAM, S. (2012): “Personal Taxes and Benefits,” Presented as part of post-Autumn Statement 2012 Institute for Fiscal Studies press conference.
- AGUIAR, M. AND E. HURST (2005): “Consumption versus Expenditure,” *Journal of Political Economy*, 113, pp. 919–948.
- ANYAEGBU, G. (2010): “Using the OECD equivalence scale in taxes and benefits analysis,” *Economic & Labour Market Review*, 4, 49–54.
- ATTANASIO, O. P. AND G. WEBER (1993): “Consumption Growth, the Interest Rate and Aggregation,” *The Review of Economic Studies*, 60, pp. 631–649.
- BANKS, J., R. BLUNDELL, AND S. TANNER (1998): “Is There a Retirement-Savings Puzzle?” *American Economic Review*, 88, 769–88.
- BANKS, J., C. EMMERSON, Z. OLDFIELD, AND G. TETLOW (2005): “Prepared for retirement? The adequacy and distribution of retirement resources in England,” IFS Report 67, Institute for Fiscal Studies.
- BARCLAYS CAPITAL (2012): “The Equity Gilt Study 2012,” Tech. rep., Barclays Capital, London.
- BLUNDELL, R., M. BROWNING, AND C. MEGHIR (1994): “Consumer Demand and the Life-Cycle Allocation of Household Expenditures,” *The Review of Economic Studies*, 61, pp. 57–80.
- BOUND, J., T. STINEBRICKNER, AND T. WAIDMANN (2010): “Health, economic resources and the work decisions of older men,” *Journal of Econometrics*, 156, 106–129.
- BOZIO, A., R. CRAWFORD, AND G. TETLOW (2010): “The history of state pensions in the UK: 1948 to 2010,” Institute for Fiscal Studies Briefing Note 105, Institute for Fiscal Studies.
- BOZIO, A., C. EMMERSON, AND G. TETLOW (2011): “How much do lifetime earnings explain retirement resources?” Institute for Fiscal Studies Working paper 11/02, Institute for Fiscal Studies.
- BROWNING, M. AND C. MEGHIR (1991): “The Effects of Male and Female Labor Supply on Commodity Demands,” *Econometrica*, 59, pp. 925–951.
- CRAWFORD, R. AND C. O’DEA (2012): “The adequacy of wealth among those approaching retirement,” IFS Report 72, Institute for Fiscal Studies.
- DE NARDI, M. (2004): “Wealth Inequality and Intergenerational Links,” *The Review of Economic Studies*, 71, 743–768.
- DE NARDI, M., E. B. FRENCH, AND J. B. JONES (2010): “Why Do the Elderly Save? The Role of Medical Expenses,” *Journal of Political Economy*, 118, 39–75.
- FREDERICK, S., G. LOEWENSTEIN, AND T. O’DONOGHUE (2002): “Time Discounting and Time Preference: A Critical Review,” *Journal of Economic Literature*, 40, 351–401.
- GREENE, W. (2004): “The behaviour of the maximum likelihood estimator of limited dependent variable models in the presence of fixed effects,” *Econometrics Journal*, 7, 98–119.

- GUSTMAN, A. L. AND T. L. STEINMEIER (2005): “The social security early entitlement age in a structural model of retirement and wealth,” *Journal of Public Economics*, 89, 441–463.
- GUVENEN, F. (2009): “An Empirical Investigation of Labor Income Processes,” *Review of Economic Dynamics*, 12, 58–79.
- HOOD, A. AND R. JOYCE (2013): “The economic circumstances of cohorts born between the 1940s and the 1970s,” Tech. rep.
- KOPCZUK, W. AND J. P. LUPTON (2007): “To Leave or Not to Leave: The Distribution of Bequest Motives,” *Review of Economic Studies*, 74, 207–235.
- LOW, H., C. MEGHIR, AND L. PISTAFERRI (2009): “Wage Risk and Employment Risk over the Life Cycle,” NBER Working Papers 14901, National Bureau of Economic Research, Inc.
- (2010): “Wage Risk and Employment Risk over the Life Cycle,” *American Economic Review*, 100, 1432–67.
- LUDWIG, A. AND A. ZIMPER (2013): “A parsimonious model of subjective life expectancy,” *Theory and Decision*, 75, 519–541.
- MITCHELL, O. S., J. M. POTERBA, AND M. J. WARSHAWSKY (1999): “New Evidence on the Money’s Worth of Individual Annuities,” *American Economic Review*, 89, 1299–1318.
- MUNNELL, A. H., A. WEBB, AND F. N. GOLUB-SASS (2007): “Is There Really a Retirement Savings Crisis? An NRRI Analysis,” Brief IB7/11, Center for Retirement Research at Boston College.
- (2012): “The National Retirement Risk Index: An Update,” Brief IB12/20, Center for Retirement Research at Boston College.
- MURTHI, M., P. R. ORSZAG, AND P. R. ORSZAG (1999): “The Value for Money of Annuities in the UK: Theory, Experience and Policy,” .
- NEYMAN, J. AND E. L. SCOTT (1948): “Consistent Estimates Based on Partially Consistent Observations,” *Econometrica*, 16, pp. 1–32.
- ONS (2013): “Occupational Pension Schemes Survey, 2012,” Tech. rep., Office for National Statistics.
- SAMWICK, A. A. (1998): “Discount rate heterogeneity and social security reform,” *Journal of Development Economics*, 57, 117–146.
- SCHOLZ, J. K., A. SESHADRI, AND S. KHITATRAKUN (2006): “Are Americans Saving "Optimally" for Retirement?” *Journal of Political Economy*, 114, 607–643.
- SKINNER, J. (2007): “Are You Sure You’re Saving Enough for Retirement?” *Journal of Economic Perspectives*, 21, 59–80.
- TAUCHEN, G. (1986): “Finite state markov-chain approximations to univariate and vector autoregressions,” *Economics Letters*, 20, 177–181.

A Tax and transfer function

Income tax

Income tax is levied in the UK on quite a comprehensive definition of income which includes earnings, private and state pensions and interest income. In 2002/03, income was taxed in four bands, the smallest was exempt from tax, the second attracted tax at 10%, the third at 22% and the largest at 40%. The thresholds that define the bands vary with age, with a more generous treatment of older individuals. The equations below, together with Table 9 summarise the income tax system used in the model.

$$y^{gr} = e + pp + sp + ra$$

$$\begin{aligned}
 it(e, pp, sp, ra, t) &= 0 && \text{if } y^{gr} \leq \kappa_1 \\
 &= 0.1(y^{gr} - \kappa_1) && \text{if } \kappa_1 < y^{gr} \leq \kappa_2 \\
 &= 0.1(\kappa_2 - \kappa_1) + 0.22(y^{gr} - \kappa_2) && \text{if } \kappa_2 < y^{gr} \leq \kappa_3 \\
 &= 0.1(\kappa_2 - \kappa_1) + 0.22(\kappa_3 - \kappa_2) + 0.4(y^{gr} - \kappa_3) && \text{if } \kappa_3 > y^{gr}
 \end{aligned}$$

Table 9: Income tax thresholds

	Age		
	< 65	≥ 65, < 75	≥ 75
κ_1	4,615	6,100	6,370
κ_2	6,535	8,020	8,290
κ_3	36,435	37,920	38,190

National Insurance Contributions

National Insurance payments are levied on earnings (not on capital income or other forms of income) and only on those aged less than the state pension age (65). In 2002/03 it was levied at a rate of 10% of income between the ‘Lower Earnings Limit’ (LEL - £3,900) and the ‘Upper Earnings Limit’ (UEL £30,420).

$$\begin{aligned}
 ni(e, t) &= 0.1(\max(0, (\min(uel, e) - lel))) && \text{if } t < 65 \\
 ni(e, t) &= 0 && \text{if } t \geq 65
 \end{aligned}$$

Jobseekers’ Allowance

Jobseekers’ Allowance is paid to unemployed households under the age of 60 at a rate which depends on the number of adults and children in the household. In 2002/03 an unemployed couple were entitled to £4,401.80 with an additional payment of £1,924 for each dependent child.

$$\begin{aligned}
jsa(e, h, k, t) &= 4401.8 + 1,924k \text{ if } t < 60 \text{ and } e = 0 \\
jsa(e, h, k, t) &= 0 \text{ if } \geq 60 \text{ ore } > 0
\end{aligned}$$

Child Benefit

Child Benefit is paid on the basis of the number of dependent children that a household has. It is paid at a more generous rate (£834.6 per year) for the first child than for subsequent children (£559).

$$childben(k) = 834(1(k \geq 0)) + 559(max((k - 1), 0))$$

Minimum Income Guarantee

Households aged over 60 are entitled to a means-tested transfer (the Minimum Income Guarantee) that aims to ensure no older household faces destitution. Entitlement to the MIG is based on current circumstances only and does not depend on a household's history of tax or national insurance contributions. The MIG simply tops *net* income up to a minimum level (f , which was £5,184 per year for singles and £7,790 for couples in 2002/03). Define net income (before payment of any MIG) as:

$$y^{preMig} = e + ra + sp + pp - it - ni + childben + jsa$$

MIG is then:

$$\begin{aligned}
mig(e, ra, sp, pp, h, k, t) &= max(0, f - y^{preMig}) \text{ if } t \geq 60 \\
mig(e, ra, sp, pp, h, k, t) &= 0 \text{ if } t < 60
\end{aligned}$$

Net income

We can now summarise the net income function used in the model. Net income is:

$$\begin{aligned}
\tau(e, ra, pp, sp, h, k, dc, t) &= e + ra + sp + pp \\
&\quad - it(e, pp, sp, ra, t) \\
&\quad - ni(e, t) \\
&\quad + jsa(e, h, k, t) \\
&\quad + childben(k) \\
&\quad + mig(e, ra, sp, pp, h, k, t)
\end{aligned}$$

B Value functions

Section 2.2 gives the optimisation problems in both retirement and working life faced by households in which neither spouse has died. Here we give the corresponding value functions for households where the husband has died (the only two differences between these and the case where the wife has died are that the value functions in the latter case are conditional on $h_t = 2$ and the survival probabilities are those relating to the husband s_{t+1}^m).

Retired household's problem

$$\begin{aligned}
 V_t(a_t, pp_t, h_t = 3; \theta_i) &= \max_{c_t} \left(u(c_t) + \beta s_{t+1}^f V_{t+1}(a_{t+1}, pp_{t+1}, h_{t+1} = 3; \theta_i) \right) \\
 \text{s.t. } a_{t+1} &= (1+r)(a_t + y_t - c_t) \\
 y_t &= \tau(e_t, ra_t, pp_t, sp_t, h_t, 0, t) \\
 a_{t+1} &= (1+r)(a_t + y_t - c_t) \forall t \neq 64 \\
 a_{65} &= (1+r)(a_{64} + y_{64} - c_{64}) + (1-z)(0.25)DC_{65} \\
 pp_t &= q(1-z)(0.75)DC_{65}
 \end{aligned}$$

Working age household's problem

$$\begin{aligned}
 V_t(a_t, DC_t, e_t, h_t = 3; \theta_i) &= \max_{c_t, dc_t} \left(u(c_t) + \beta(1 - s_{t+1}^m)(s_{t+1}^f) \int V_{t+1}(a_{t+1}, DC_{t+1}, e_{t+1}, 3; \theta_i) dF(\phi_{t+1}, e_{t+1} | e_t) \right) \\
 \text{s.t. } y_t &= \tau(e_t, ra_t, pp_t, sp_t, h_t, dc_t, t) \\
 a_{t+1} &= (1+r)(a_t + y_t - c_t - dc_t) \\
 DC_{t+1} &= (1 + \phi_{t+1})(DC_t + dc_t)
 \end{aligned}$$

C Computational appendix: model solution and simulation of optimal behaviour

In this section we outline how we a) solve the households' maximisation problem to obtain decision rules (functions of the state variables which give optimal consumption and optimal pension saving) and b) use these decision rules, along with our data, to simulate the optimal saving behaviour of the households in our sample.

a) Solution

There is no analytical solution to the maximisation problem outlined. Decision rules are obtained numerically by iterating on the value function from the final period of life. Recall that the problem faced by the household at period 100 (where we do not repeat the constraints, which are listed as part of maximisation problem (4)):

$$\begin{aligned}
V_{100}(x_{100}, pp_{100}, h_{100} = 1; \boldsymbol{\theta}_i) = \max_{c_{100}} & \left(u(c_{100}) + \beta s_{101}^m s_{101}^f V_{101}(x_{101}, pp_{101}, h_{101} = 1; \boldsymbol{\theta}_i) \right. \\
& + \beta s_{101}^m (1 - s_{101}^f) V_{101}(x_{101}, pp_{101}, h_{t+1} = 2; \boldsymbol{\theta}_i) \\
& \left. + \beta (1 - s_{101}^m) (s_{101}^f) V_{101}(x_{101}, pp_{101}, h_{101} = 3; \boldsymbol{\theta}_i) \right).
\end{aligned}$$

Let us rewrite this concisely as:

$$V_{100}(\mathbf{X}_{100}; \boldsymbol{\theta}_i) = \max_{c_{100}} u(c_{100}) + \beta E[V_{101}(\mathbf{X}_{101}; \boldsymbol{\theta}_i) | \mathbf{X}_{100}] \quad (12)$$

where the vector \mathbf{X} contains the state variables of the problem and the expectation operator is over survival past the age of 100. For years before retirement, the expectation will additionally be over employment offers, earnings draws and returns on the DC fund. Our assumption that death in the next period is certain for those still alive at the age of 100 ($s_{101}^m = s_{101}^f = 0$), combined with the assumption on the absence of bequest motives means that the expectation in equation (12) evaluates to 0. At any particular point in \mathbf{X} , the maximisation is therefore possible and we can obtain $c_{100}(\mathbf{X}_{100}; \boldsymbol{\theta}_i)$, the consumption function, and $V_{100}(\mathbf{X}_{100}; \boldsymbol{\theta}_i)$, the associated value function at those points (we discuss below our procedure for maximisation). The knowledge of $V_{100}(\mathbf{X}_{100}; \boldsymbol{\theta}_i)$ at a subset of points in \mathbf{X} , combined with approximation methods (also discussed below), yields an approximation of $V_{100}(\mathbf{X}_{100}; \boldsymbol{\theta}_i)$ ($\hat{V}_{100}(\mathbf{X}_{100}; \boldsymbol{\theta}_i)$) at each point in \mathbf{X} .

With an approximation to $V_{100}(\mathbf{X}_{100})$ so obtained, we can solve for approximations to the true consumption function ($\hat{c}_{100}(\mathbf{X}_{100}; \boldsymbol{\theta}_i)$) and value function ($\hat{V}_{100}(\mathbf{X}_{100}; \boldsymbol{\theta}_i)$) for the particular household i at age 99, again at a subset of points in the state space in that period, by solving the following functional equation:

$$\hat{V}_{99}(\mathbf{X}_{99}; \boldsymbol{\theta}_i) = \max_{c_{99}} u(c_{99}) + \beta E[\hat{V}_{100}(\mathbf{X}_{100}; \boldsymbol{\theta}_i) | \mathbf{X}_{99}] \quad (13)$$

and obtain $\hat{c}_{99}(\mathbf{X}_{99}; \boldsymbol{\theta}_i)$ and ($\hat{V}_{99}(\mathbf{X}_{99}; \boldsymbol{\theta}_i)$). This iterative process is repeated until we get to the beginning of working life (at age 20). For periods before retirement, a second decision rule - the quantity paid into the pension fund ($\hat{d}c_t(\mathbf{X}_t; \boldsymbol{\theta}_i)$) is also calculated and stored.

Four particular features of the solution procedure will be detailed in the following discussion. These are the i) the discretisation of the continuous variables, ii) the process by which the integral in the functional equation is evaluated, iii) the manner in which the value function is approximated at points outside the discretised state space and iv) how the optimisation is carried out.

Discretisation of state and control variables We have four continuous state variables that need to be discretised. These are earnings, cash assets, pension wealth and pension income. Earnings are placed on a grid (that has 45 elements) using a procedure suggested by Tauchen (1986). Assets, DC stocks and pensions are discretised in a manner that gives smaller gaps between successive entries on the grid at lower levels. This is as the curvature of the value function (with respect to those state variables) will be greater at lower realisations of these states. 45 discrete points are used for cash assets and 16 for each of pension wealth and pension income.

There are two choice variables in the model - consumption and the contribution to the DC pension fund. Consumption is not placed on a grid - households can choose any feasible consumption

value. To avoid the computational burdens associated with having two continuous control variables, the proportion of earnings that can be contributed to a DC pension fund is restricted to take on one of 8 values. That is, households can contribute 0%, 5%, 10%, 15%, 25%, 40%, 75% or 90% of their earnings to the pension fund.

Integration Evaluation of the expectations in the households' problem involves integration of the value function over four stochastic variables. These are earnings, survival and the return on DC funds. Realisations of survival and earnings take one of a number of discrete outcomes – the former as it is naturally discrete, the latter as the procedure we apply (Tauchen (1986)) allows earnings to take only a discrete subset of outcomes. Integration over the possible realisations of earnings and survival is therefore carried out by taking a weighted average the value function realised at each possible outcome with the weights equal to the probability of that outcome. Realisations of the return on the DC fund are not restricted to a discrete subset. Integration over the distribution of possible outcomes is carried out using Gauss-Hermite quadrature with 10 nodes.

Approximation It is required to evaluate the functions $V_t(\cdot)$ at points in the state space other than those in the discrete sub-set of points in the discretised state space. Linear interpolation (in multiple dimensions) is used.

Optimisation In retirement households face a single choice each year - how much to consume (with the rest of their resources saved in a safe asset). Each optimisation is carried out by Golden Section Search. This will successfully find a maximum as our approximated Value Function is quasiconcave. In working life households face two choices - how much to consume and how much to pay into a pension (again, with the rest of their resources saved in the safe asset). Here we solve (again using Golden Section Search) for optimal consumption at each of the permitted rates of contribution to the DC fund. The optimal rate of contribution to the DC fund is that which, of these, maximises utility.

If the pension contribution was allowed as a continuous choice, the approximated value function would be quasiconcave. The discretisation of the pension contribution choice implies that the approximated value function may not be quasiconcave and therefore a local optimisation routine, like the Golden Section Search, finds the global optimum. The reasons for this are discussed in Appendix A of Low et al. (2009) where a similar issue arises. Those authors suggest that in problems where there is a lot of uncertainty (as there is in ours) local optimisers - such as the Golden Section Search routine - are likely to obtain the global solution. Their approach is to use the local estimator while estimating their parameters (which involves many solutions of the value function), and then at the set of parameters to check their result using a global optimiser. Our approach is similar. The results presented in the paper (on both the model set out and the sensitivity analyses) use the local optimiser. However, for our baseline estimates we check that the predicted level of wealth does not materially change when we take a different approach to optimisation - one that is robust to departures from quasiconcavity of the value function. This involves restricting consumption to be on a grid of 100 values (so that, in each period, households can choose to consume 1% of their resources, 2%, 3% etc.) and selecting (from the discrete subset of permissible selections) the levels of consumption and pension contribution that maximise utility. The results from this check support the use of the local optimiser - using the global optimiser, we get mean optimal wealth levels within

1% of the value found using the local optimiser and the same proportion of households to one decimal place (7.9%) are found to be undersaving.

b) Simulation

Once decision rules for pension saving ($\hat{d}c_t(\mathbf{X}_t; \boldsymbol{\theta}_i)$) and consumption ($\hat{c}_t(\mathbf{X}_t; \boldsymbol{\theta}_i)$) are obtained we can simulate the behaviour that a household member would exhibit if they followed those rules. The procedure is as follows:

1. Set initial values for state variables at the beginning of working life (age 20). The state variables that are relevant at the start of working life are cash, pension fund value, earnings and household composition. We set cash and pension fund value to zero. We set earnings to the value on the grid that is closest to actual observed earnings at the age of 20. For household composition we assume both members of the couple are alive and in a couple at that age.
2. Using these values for the state variables and our knowledge of the household's type ($\boldsymbol{\theta}_i$), and the decision rules ($\hat{c}_t(\mathbf{X}_t; \boldsymbol{\theta}_i)$ and $\hat{d}c_t(\mathbf{X}_t; \boldsymbol{\theta}_i)$) we obtain optimal consumption and optimal payments into the pension fund in period 20 (c_{20}^i, dc_{20}^i).
3. Obtain the new state variables for period 21. These are obtained as follows:
4. Cash assets in period 21 will follow from the consumption and saving decisions in period 20 along with equations (14) and (15) - two of the constraints on the optimisation problem faced by a working age household.

$$y_{20}^i = \tau(e_{20}^i + s(h_{20}^i, e_{20}^i, a_{20}^i, DC_{20}^i), 20, dc_{20}^i) \quad (14)$$

$$x_{21}^i = (1 + r)(x_{20}^i + y_{20}^i - c_{20}^i - dc_{20}^i) \quad (15)$$

- (a) Pension wealth in period 21 will be the sum of the stock of pension wealth in period 20, the flow into the pension wealth and the assumed growth rate of pension funds between ages 20 and 21 (equation (16)). That growth rate is assumed to be equal to the growth rate (from our time series of pension fund growth rates) in the year that this household turns 21.

$$DC_{21}^i = (1 + \phi_{21})(DC_{20}^i + dc_{20}^i) \quad (16)$$

- (c) Earnings in period 21 will be that point on the earnings grid that is closest to actual earnings observed at the age of 21.
 - (d) Household composition will remain set equal $h_{t+1} = 1$, that is both members of the couple are still alive. This is as we only retain sample members where nobody has died by the time they are observed in the data.
5. Repeat steps 2 and 3 to obtain optimal consumption and pension saving at each age up to the age at which the (male in the household) is observed in the data in 2002 (we call this age τ).

None of these men will have reached their state pension age before this period and therefore the decision rules of retired households are not needed in the simulations.¹⁷ This will allow a time series of the value held in both assets from the age of 20 to age τ : $\{(x_t^i, DC_t^i)\}_{t=20}^\tau$. Our central results involve comparing simulated optimal wealth at age τ (x_τ^i, DC_τ^i) with that observed in the data at that age.

¹⁷Though of course the decision rules for working age households could not have been calculated without first solving the retired households' problem.

D Derivation of moments of earnings process

This section derives the moment conditions in equations (9) to (10) that are used to estimate the parameters of the earnings process. Let us first derive the variances and autocovariances associated with the differences in the persistent stochastic innovation to earnings (u). Taking differences of both sides of equation 7 gives:

$$\Delta u_{it} = \rho \Delta u_{it-1} + \Delta \xi_{it} \quad (17)$$

The variance of this can be obtained as follows (suppressing i subscripts and letting *c.p.* refers to cross-products that have expectation zero):

$$\begin{aligned} \text{var}(\Delta u_t) &= \text{var}(\rho \Delta u_{t-1} + \Delta \xi_t) \\ &= \rho^2 \text{var}(\Delta u_{t-1}) + \text{var}(\Delta \xi_t) + 2\text{cov}(\rho \Delta u_{t-1}, \Delta \xi_t) \\ &= \rho^2 \text{var}(\Delta u_{t-1}) + 2\sigma_\xi^2 + 2E(\rho \Delta u_{t-1} \Delta \xi_t) \\ &= \rho^2 \text{var}(\Delta u_{t-1}) + 2\sigma_\xi^2 + 2E(\rho(u_{t-1} - u_{t-2})(\xi_t - \xi_{t-1})) \\ &= \rho^2 \text{var}(\Delta u_{t-1}) + 2\sigma_\xi^2 + 2E(\rho(\rho u_{t-2} + \xi_{t-1} - \rho u_{t-3} + \xi_{t-2})(\xi_t - \xi_{t-1})) \\ &= \rho^2 \text{var}(\Delta u_{t-1}) + 2\sigma_\xi^2 + 2E(\rho(\xi_{t-1})(-\xi_{t-1}) + \text{c.p.}) \\ &= \rho^2 \text{var}(\Delta u_{t-1}) + 2\sigma_\xi^2 - 2\rho\sigma_\xi^2 \\ &= \rho^2 \text{var}(\Delta u_{t-1}) + 2(1 - \rho)\sigma_\xi^2 \end{aligned} \quad (18)$$

The autocovariances at lead 1, and lead greater than 1 respectively are obtained as follows:

$$\begin{aligned} \text{cov}[\Delta u_t \Delta u_{t+1}] &= E[\Delta u_t \Delta u_{t+1}] \\ &= E[\Delta u_t (\rho \Delta u_t + \Delta \xi_{t+1})] \\ &= E[\rho \Delta u_t \Delta u_t + \Delta u_t \Delta \xi_{t+1}] \\ &= \rho \text{var}(\Delta u_t) + E[(u_t - u_{t-1})(\xi_{t+1} - \xi_t)] \\ &= \rho \text{var}(\Delta u_t) + E[(\rho u_{t-1} + \xi_t - \rho u_{t-2} - \xi_{t-1})(\xi_{t+1} - \xi_t)] \\ &= \rho \text{var}(\Delta u_t) + E[(\xi_t)(-\xi_t) + \text{c.p.}] \\ &= \rho \text{var}(\Delta u_t) - \sigma_\xi^2 \end{aligned} \quad (19)$$

$$\begin{aligned} \text{cov}[\Delta u_t \Delta u_{t+j}] &= E[\Delta u_t \Delta u_{t+j}] \quad \forall j \geq 2 \\ &= E[\Delta u_t (\rho \Delta u_{t+j-1} + \Delta \xi_{t+j})] \\ &= E[\rho \Delta u_t \Delta u_{t+j-1} + \Delta u_t \Delta \xi_{t+j}] \\ &= \rho \text{cov}(\Delta u_t \Delta u_{t+j-1}) + E[(u_t - u_{t-1})(\xi_{t+j} - \xi_{t+j-1})] \\ &= \rho \text{cov}(\Delta u_t \Delta u_{t+j-1}) + E[\text{c.p.}] \\ &= \rho \text{cov}(\Delta u_t \Delta u_{t+j-1}) \end{aligned} \quad (20)$$

The variances and autocovariances of the differences in the *iid* measurement error component are simpler to derive:

$$\text{var}(\Delta m_t) = 2\sigma_m^2 \quad (21)$$

$$\begin{aligned} \text{cov}[\Delta m_t \Delta m_{t+1}] &= E[\Delta m_t \Delta m_{t+1}] \\ &= E[(m_t - m_{t-1})(m_{t+1} - m_t)] \\ &= E[(m_t)(-m_t) + c.p.] \\ &= -\sigma_m^2 \end{aligned} \quad (22)$$

$$\begin{aligned} \text{cov}[\Delta m_t \Delta m_{t+j+1}] &= E[\Delta m_t \Delta m_{t+j+1}] \quad \forall j \geq 2 \\ &= E[(m_t - m_{t-1})(m_{t+j} - m_{t+j-1})] \\ &= E[c.p.] \\ &= 0 \end{aligned} \quad (23)$$

We can use these 6 expressions to obtain the moments that we bring to the data. These are the variances and autocovariances associated with the difference in the total deviation of recorded earnings from the deterministic component of earnings. This difference can be obtained from equation (6):

$$\Delta v_{it} = \Delta u_{it} + \Delta m_{it} \quad (24)$$

The variance of Δv_{it} , suppressing i subscripts and using equations (18) and (21) is:

$$\begin{aligned} \text{var}(\Delta v_t) &= \text{var}(\Delta u_t + \Delta m_t) \\ &= \text{var}(\Delta u_t) + \text{var}(\Delta m_t) \\ &= \rho^2 \text{var}(\Delta u_{t-1}) + 2(1 - \rho)\sigma_\xi^2 + 2\sigma_m^2 \end{aligned} \quad (25)$$

The autocovariance at lead 1 (using equations (19) and (22) is given by:

$$\begin{aligned} \text{cov}(\Delta v_t, \Delta v_{t+1}) &= E[\Delta v_t \Delta v_{t+1}] \\ &= E[(\Delta u_t + \Delta m_t)(\Delta u_{t+1} + \Delta m_{t+1})] \\ &= E[\Delta u_t \Delta u_{t+1} + \Delta m_t \Delta m_{t+1} + c.p.] \\ &= \text{cov}(\Delta u_t, \Delta u_{t+1}) + \text{cov}(\Delta m_t \Delta m_{t+1}) \\ &= \rho \text{var}(\Delta u_t) - \sigma_\xi^2 - \sigma_m^2 \end{aligned} \quad (26)$$

The autocovariances at leads greater than 1 (using equations (20) and (23)) are:

$$\begin{aligned}
cov(\Delta v_t, \Delta v_{t+j}) &= E[\Delta v_t \Delta v_{t+j}] \quad \forall j \geq 2 \\
&= E[(\Delta u_t + \Delta m_t)(\Delta u_{t+j} + \Delta m_{t+j})] \\
&= E[\Delta u_t \Delta u_{t+j} + \Delta m_t \Delta m_{t+j} + c.p.] \\
&= cov(\Delta u_t, \Delta u_{t+j}) + cov(\Delta m_t, \Delta m_{t+j}) \\
&= \rho cov(\Delta u_t, \Delta u_{t+j-1})
\end{aligned} \tag{27}$$

Equations (25) to (27) are the covariances given in equations (9) to (11) in Section 4.

E Estimating DC fund return

The DCisions index is an index of total fund return that reflects the asset allocation decisions made by leading DC pension plans in their default investment strategies. Over the period 1994 - 2010 the DCisions index exhibited slightly greater growth than that of the FTSE all-share index (an index representing the performance of the majority of companies listed on the London Stock Exchange). Across financial years where the FTSE all-share index grew in nominal terms, the median ratio of the growth in the DCisions index to the growth in the FTSE all-share index was 1.17, while across financial years where the FTSE all-share index fell in nominal terms, the median ratio was 0.89. This is the result of including re-investment of dividends (the DCisions index is a total return index while the FTSE all-share is an asset price index), slightly offset by the average DC pension plan being diversified into a portfolio with slightly lower return (but also lower risk) than the equities included in the FTSE all-share.

For years 1994 to 2010, therefore, ϕ_t (the model's rate of growth of funds in pension wealth) is assumed to be the real growth in the annualised DCisions index. For years prior to 1994 in which the FTSE all-share index increased in nominal terms, ϕ_t is assumed to be 1.17 times the growth in the FTSE all-share index; for years prior to 1994 in which the FTSE all-share index fell in nominal terms, ϕ_t is assumed to be 0.89 times the decline in the FTSE all-share index. The FTSE index is assumed to have grown by 4% per year in nominal terms in years before data on the FTSE all-share index are available.

F Calculating DB pension wealth

Projected DB pension *income* is estimated using survey responses (see Banks et al. (2005) for more details). The calculation of DB *wealth* needs to account for the fact that when ELSA respondents are interviewed, their actual pension income will depend on labour market outcomes that have not yet been realised. The assumption that underlies the estimates used here is to assume that individuals who are not working when interviewed do not return to the labour market, whereas those that are working when interviewed continue to work until their current DB pension's normal retirement age (NRA). The following calculation converts this projected DB stream of income at the pension's NRA into current wealth at age t - the age at which households are surveyed.

$$DB_t = \frac{db_{nra}}{\underbrace{q_{nra}}_A} \frac{yw_t}{\underbrace{yw_{nra}}_B} \underbrace{(1+r)^{t-nra}}_C$$

The first term (A) is projected DB income at the age of the NRA divided by the annuity rate at the NRA - this is a measure of DB wealth at the NRA. The second term (B) accounts for the fact that, by age t , individuals will not have accrued all the DB pension wealth in A and scales it down by a factor equal to the ratio of the number of years worked in period t to the number of years assumed to be worked at the NRA yw_{nra} . Finally, the product of A and B is wealth that will not be realised until the age of NRA, so to obtain its value in period t we must discount it back using the risk-free interest rate (term C).