Private Pensions and Public Pension Design*

Cormac O'Dea†

October 28, 2016

Abstract

Government pension spending in developed countries can be divided into three types: (1) Social Security-style benefits that depend on earnings during working life, (2) subsidies of private pension saving and (3) means-tested transfers to the elderly. Using an estimated lifecycle model that accounts for each of these, as well as endogenous labour supply, private savings and realistic uncertainty, this paper investigates the optimal combination of the three approaches. I show, using an optimal taxation framework, that for countries (such as the US and the UK) that currently provide public pensions that depend on career-average earnings, large welfare increases can be obtained by increasing means-tested old-age transfers that are funded by any of reducing public pensions, increasing taxes or (especially) reducing private pension subsidies. While means-tested transfers cause costly distortions, I show that these are more than offset by the value of the insurance they provide against low lifetime earnings potential, poor investment returns and longevity. The optimality of greater means-tested support is specific to older individuals: I find that such support to younger households should be at a lower level than that to the elderly. These results imply that governments should provide strong work incentives for the young, but provide pensions with good insurance properties for the old.

JEL Classification: D91, E21, D14

Keywords: Pensions; Lifecycle; Savings; Social Security; Household Finance

1 Introduction

The provision of support for the elderly in retirement is one of the most costly activities carried out by governments. These costs are growing as life expectancies increase. Spending on public pensions is

*Thanks to Richard Blundell and Hamish Low for comments and encouragement and to James Banks, Rowena Crawford, Thomas Crossley, Mariacristina De Nardi, Eric French, Andrew Hood, Guy Laroque, Peter Levell, Rory McGee, Barra Roantree and Ananth Seshadri for very helpful discussions. Thanks to the Economic and Social Research Council (Centre for Microeconomic Analsyis of Public Policy at the Institute for Fiscal Studies- ref: RES-544-28-50001 and Analysis of Secondary Data Grant - ref: ES/N011872/1) for funding this work. Correspondence to cormac.odea@ifs.org.uk. Any errors are my own.

†University College London and the Institute for Fiscal Studies

forecast to reach an average of 12% of GDP across the OECD by 2050 - up from 6% in 1990. Financial support to the elderly in the developed world is typically provided through one or more of three types of payments: i) Social Security-style benefits that depend on earnings during working life, ii) subsidies of private pension saving and iii) means-tested transfers to retirees. Countries differ in the extent to which they emphasise each of these. Many continental European countries provide public pensions that are proportional to career-average earnings (up to some cap); Social Security in the US replaces career-average earnings according to a progressive schedule while the UK's state pension system gives a payment that also varies with lifetime earnings, but to a lesser extent. These countries also have means-tested transfers that provide an effective consumption floor to those with little other income. In most countries, spending on these means-tested programs is substantially less than that on the core earnings-related public pension. Australia is an exception. There, the public pension – the 'Age Pension' – is entirely means-tested, and high income pensioners receive no public pension. The extent of and types of subsidies through the tax system for private pension saving also differ substantially across countries.

The aim of this paper is to assess the optimal balance of these different types of support for pensioners. Means-tested retirement consumption floors provide help to households when they need it most, providing valuable insurance against low lifetime earnings, poor investment returns and longevity, but they may provide strong disincentives to work and save. On the other hand, tax-favoured private pension schemes provide strong incentives to work and save, but do little to reduce inequality in lifetime outcomes. Progressive public pensions deliver a mix of these incentives and insurance properties. This paper develops a rich estimated lifecycle model that can be used to value the insurance provided by different types of support for pensioners and to quantify the costs that they impose. The model contains public pensions (e.g. Social Security), heterogeneous private pensions and old-age means-tested support. Labour supply and private savings are endogenous and households are exposed to realistic uncertainty over their employment, earnings, pension fund returns and mortality. The key contributions of this paper are to study the design of both public and private pensions within the same framework and to do so with a model in which preference parameters are estimated.

The estimated model is used within an optimal taxation framework to obtain the ex-ante optimal level of a means-tested consumption floor for retirees. Starting from the current UK old-age consumption floor (which is higher than that provided by Supplemental Security Income in the US), revenue-neutral reforms are considered (with a variety of instruments used to balance the budget). The current level offered in the UK is too low: increases funded by any of reducing the generosity of the public pension, increasing tax rates or (especially) reducing the extent of private pension subsidies deliver large welfare increases. The increase in welfare comes from a reduction in the variance of lifetime consumption and increases in leisure time, partially offset by lower average lifetime consumption (due to the now-larger distortionary means-tested payment).

The same optimal taxation framework does not, however, suggest extending such consumption floors to younger individuals - in fact, ex-ante welfare analysis suggests a *lowering* of current levels of means-tested support to those in the UK who don't accept a job offer. The negative effects on labour supply induced by providing a means-tested consumption floor are much greater for working-age than for older individuals. This makes the cost of rolling out a generous all-age consumption floor would cause distortions that substantially outweigh the benefits from insurance against poor earnings outcomes. These results point to the value of providing strong incentives to work to the young alongside pensions with good insurance properties to the old.

The first branch of the literature to which this paper relates concerns the design of public pensions. Following Auerbach and Kotlikoff (1987), this literature[1] has focussed on how Social Security might be made more affordable in light of the pressures imposed on it by changing demographics. Solutions that have been heavily studied include raising payroll taxes, delaying eligibility ages and reducing the generosity of benefits. The literature has, however, generally neglected to study the possibilities afforded by the means-testing of benefits, in part as such transfers are known to reduce labour supply and crowd out private saving.[2] An exception is Kitao (2014) who studies four options to make Social Security sustainable - one of which involves an extreme form of means-testing (whereby all Social Security benefits, after a small disregard, are withdrawn at an effective tax rate of 100%). This reform is rejected as "[due to] the large negative effects on economic activities and fiscal burden, it is unlikely to be a viable option for the social security reform".

However, less extreme increases in means-tested benefits have been shown using calibrated models to have the potential to be welfare-increasing. Sefton and Van De Ven (2009) find that enhancing the generosity of the means-tested component of the UK public pension system would be welfare-improving. Braun et al. (2016) show that a 33% increase in the generosity of means-tested social insurance programs in the US (Medicaid, Supplemental Social Security Income, food stamps and a number of smaller programs), would be welfare increasing if funded through the payroll tax. Huggett and Parra (2010) and Golosov et al. (2013) both find that making Social Security more progressive (albeit not by means-testing them) would be welfare-enhancing. The mechanism behind these results is the same as that laid out above: lower average consumption is outweighted by a smaller variance of consumption in the calculation of

[1] A partial list of contributions is De Nardi et al. (1999), Conesa and Krueger (1999), Kotlikoff et al. (1999), Huggett and Ventura (1999), Nishiyama and Smetters (2007) and Conesa et al. (2009).

[2] Neumark and Powers (1998, 2000) show that SSI reduces pre-retirement saving and labour. Hubbard et al. (1995) first showed the role that consumption floors play in reducing wealth accumulation in a lifecycle model. De Nardi et al. (2010) show that, in the presence of the large wealth shocks (such as those that come from healthcare needs at old ages), such floors can affect the behaviour even of those who have very high permanent income.

social welfare. The current paper adds to that literature by solving, in an estimated model, for an ex-ante optimal means-tested consumption floor, rather than exploring the welfare-implications of ad-hoc changes.

The second literature to which this paper relates is one which considers private pensions and their effect on behaviour. This literature is smaller. Nishiyama (2011) investigates the budgetary and welfare properties of tax-deferred savings vehicles (e.g. a UK Defined Contribution pension or US 401k plan). Chetty et al. (2014) finds that most savers who respond to subsidies for retirement accounts do so by shifting saving from other forms of saving into retirement accounts rather than by doing additional saving. Blau (2016) looks at the extent to which different types of private pension crowd out non-pension saving. None of these papers considers how the tax treatment of private pensions interacts with design issues around public pensions, as this paper does.

The interactions between private pension subsidies and the design of the public pension system are important. In the presence of private pension subsidies, increases in the consumption floor are welfare-enhancing for two reasons. The first is standard: consumption floors provide insurance which households are prepared to pay for through the tax system. The second is due to the fact that reforms that increase consumption floors reduce household private pension saving which (given that this form of saving attracts subsidies) defrays some of the cost of the more generous floor. While government spending on private pension subsidies and the generosity of means-tested consumption floors are therefore substitutes in the government's budget constraint, the latter have much better insurance properties and households prefer them – driving up the optimal consumption floor. When private pension subsidies are removed, consumption floors should still be increased, but to a much lesser extent.

The key questions that this paper seeks to address are how to design a public pension system and whether such systems should be complemented by subsidising private pensions. The analysis points to a greater role for means-tested consumption floors which provide, at an acceptable cost, valuable insurance to households against low working life earnings, poor investment returns and longevity. Pensions which amplify lifetime earnings risks (such as earnings-related public pensions or private pension subsidies) are substantially less preferred.

The paper proceeds as follows. To motivate the modelling choices which will come later, Section 2 very briefly discusses some typical features of public and private pensions. Section 3 outlines the model used in the paper before Section 4 details the estimation procedures, gives parameters estimates and discusses model fit. Section 5 uses the model to compare the welfare properties of international systems and to find the optimal level of a means-tested retirement consumption floor. Section 6 concludes.

2 Pensions

This section briefly describes some typical features found in international public and private pension systems - first describing the system of public pensions, and then private pensions. The aim of this section is to introduce some terminology that will be important throughout the rest of the paper and motivate some of the modelling decisions.

2.1 Public pensions

Public pensions can be either 'contributory' (they depend on earnings during working life) or can be 'means-tested' (they depend on income and assets in retirement). Examples of the former are Social Security in the US and the State Pension in the UK. Examples of the latter are Supplemental Security Income (SSI) in the US, Pension Credit (PC) in the UK and the Australian Age Pension (AP). Figure 1 illustrates these. Figure 1(a) shows how Social Security (US) and State Pension (UK) pension entitlements vary with average working life earnings for a sample of individuals born between 1935 and 1950 (the data used here will be discussed further in Section 4).[3] Figure 1(b) shows how post-means-tested income varies with pre means-tested-pension income in each of the US, UK and Australia. For the poorest pensioners both SSI and PC top pension income up to a minimum level. This is initially withdrawn at an effective tax rate of 100%, with benefits over a certain quantity taxed away at a lower effective tax rate of 40% in the UK. In Australia a small amount of income is disregarded in applying the income test, after which the Age pension is withdrawn at an effective tax rate of 50%.[4]

The conceptual difference between these two types of pension is whether earnings during working life (on the horizontal axis in the left-hand graph) or income in retirement (on the right-hand graph) determine the level of entitlement.

2.2 Private pensions

Private pensions can be grouped into two broad types - Defined Benefit (DB) pensions and Defined Contribution (DC) pensions. DB pensions pay a fraction of some function of earnings – for example, career average earnings or final earnings. DC (401k-style) pensions are investment accounts owned by the individual that can be used to purchase an annuity or otherwise provide an income in retirement. DB

[3] In both countries, this illustrative analysis is carried out at an individual level and does not take into account benefits that are earned on the basis of a spouse's contributions.

[4] These figures use the values for the 2013 systems. US and Australian dollars are converted to pounds using exchange rates for that year. The US figures are an average of the total (federal and state) entitlements for the states whose payments are delivered by the federal Social Security Administration. The x-axis here represents *pension* income - in the US and UK a certain level of employment income can be earned without affecting entitlement to the benefit. There are also differing rules regarding how assets holdings affect entitlement.

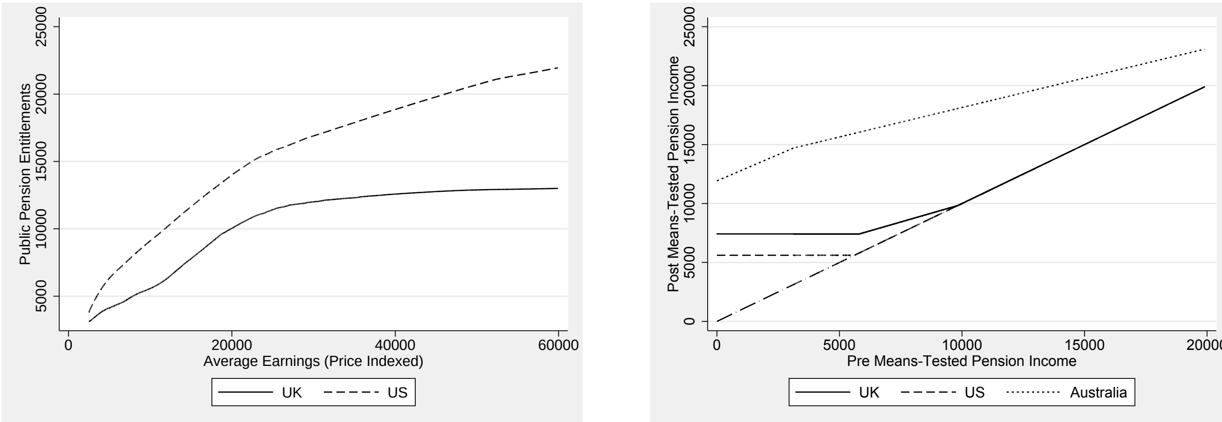

(a) Contributory pensions

(b) Means-tested pensions

Figure 1: Public Pensions in the UK and US

pensions can be thought of as a deterministic function of earnings, where DC pensions are a stochastic function of contributions into a pension fund.

Private pensions are subsidised by the tax system in many countries. The form of these vary internationally but typically involves some form of tax-deductibility (that is payments into a pension fund can be made out of gross earnings). If pension income is subjected to lower rates of tax than earned income, this tax deferral can be thought of as a subsidy (and will incentivise households to save in a pension - either by substituting consumption from during working life to retirement, or substituting towards pension saving from non-pension saving). Figure 2 illustrates that such favourable taxation is commonplace. It shows, for a selection of OECD countries, the average tax rate (black bar) on a worker earning average earnings and the average tax rate (grey bar) on a pensioner with a pension equal to average earnings. In most countries, the latter tax rate is lower (and often substantially so). The favourable treatment of pension saving, relative to non-pension saving, costs 1.1% of GDP in the UK and 0.9% of GDP in the US relative to a benchmark where pension income is taxed similarly to earnings (see Appendix A for details on calculation of these figures).

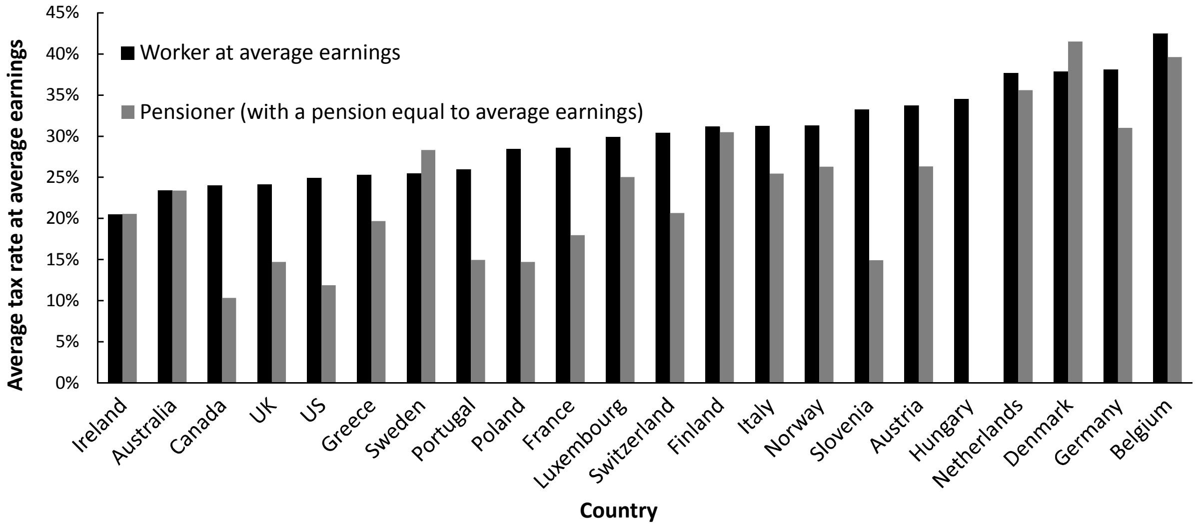

Figure 2: Average tax rates on earnings and pensions

Source: OECD (2016), Figure 6.6

3 Model

Before detailing the model I provide a short summary of its key features. The decision-making unit is a household that maximises an intertemporal utility function by choosing each of labour supply, consumption, pension saving and non-pension saving. Households are exposed to multiple sources of uncertainty: i) over whether they get an employment offer, ii) over productivity (their earnings if in work), iii) the investment returns they earn on any DC pension wealth and iv) their mortality. These households are heterogenous in their access to private pensions: some are offered a Defined Benefit pension through their employer, while all can save in a (risky) Defined Contribution pension should they wish to. Government partially insures households against risk by levying a progressive income tax, providing unemployment insurance and giving two types of payment to pensioners: a contributory pension and a means-tested pension.

The rest of this section discusses in turn the utility function, modelled pension and non-pension assets, the role of government and the household's maximisation problem. Table 9 in Appendix B gives a summary of all the variables introduced in this section.

3.1 Household composition, utility and decision

Household composition All households contain a married couple of age t who start their working life at age 20. Mortality is stochastic. Household composition (h) takes a value of 1, 2, 3 or 4 indicating, respectively, that both spouses are still alive, only the male is alive, only the female is alive or that both spouses are dead.

Heterogeneity Households are one of four types (indexed by j). These types are each pairwise combination of low/high education and having access to a DB pension or not. Household types are determined before the start of working life and are fixed for life. Different types have different productivity processes and have different risks over employment (allowing for the fact those who typically have DB pensions (e.g government employees) could have less labour market risk than those working in the private sector). Preference parameters differ across types.

Utility Households get utility from consuming, from leisure and from leaving bequests. The period utility function, given in equation (1) is non-separable in consumption (c) and leisure (l) with a weight on consumption ν and a coefficient of relative risk aversion on utility of γ.

$$u(c,l) = \frac{(c^{\nu} l^{1-\nu})^{1-\gamma}}{1-\gamma} \tag{1}$$

Households value bequests through a warm-glow bequest function, of a form used by Nardi (2004) and French (2005) and given in equation (2). a^b are assets bequeathed, θ determines the importance of bequest motives to households and K is a constant that ensures that the marginal utility of leaving a zero bequest is finite.

$$b(a^b) = \theta \frac{(a^b + K)^{(1-\gamma)\nu}}{1-\gamma}. \tag{2}$$

Decisions In each period households make four decisions. They decide i) employment at the extensive margin ii) non-housing consumption (c^{nh}), iii) how much, if anything, to contribute to their DC pension (dc) and iv) how much non-pension saving to do.

Employment and earnings The labour supply behaviour of only one household member - the 'main earner' - is modelled. The labor supply of the second earner is exogenous. This main earner chooses whether or not to supply labour if offered a job. The probability of not getting a job offer ($ue_t = 1$) evolves according to a conditional Markov process where the probability of unemployment, $\pi_1(\tilde{e})$, is conditional on current productivity ($\tilde{e}$).

When employed, earnings are equal to productivity, whose log (equation (3)) is the sum of a deterministic component (a quadratic in age) and a stochastic component (u).

$$\ln \tilde{e}_{it} = \delta_0 + \delta_1 t + \delta_2 t^2 + u_{it} \tag{3}$$

The evolution of the stochastic component in periods following an employment offer (equation (4)) follows an AR(1) process with innovations distributed normally. The variance of these innovations differs in the first and in subsequent periods.

$$\begin{aligned} u_t &= \rho u_{t-1} + \xi_t \;\; \forall \;\; t \;\; \text{with} \;\; ue_{t-1} = 0 \\ \xi_1 &\sim N\left(0, \sigma_\zeta^2\right) \\ \xi_t &\sim N\left(0, \sigma_\xi^2\right) \;\; \forall \;\; t > 1 \end{aligned} \tag{4}$$

In periods following a period of unemployment, productivity is drawn from a distribution E.

The labour supply behaviour of the second earners is exogenous. Households receive a fixed payment (e^s) up to a retirement age for the second earner (t^{ret_s}).

Consumption Consumption in utility function is the sum of non-housing (c^{nh}) and housing consumption (c^h). The former is a choice, the latter is the product of ($r^{houscon}$) – a preference parameter representing the rental value of housing – and $gh()$, gross housing wealth (which is an exogenous function of non-pension wealth and age):

$$c_t^h = r^{houscon} gh(a, t). \tag{5}$$

3.2 Assets

Households have access to up to three assets. These are, a Defined Benefit pension (for those types eligible), a Defined Contribution pension and non-pension wealth. These assets are now discussed in turn.

Defined Benefit pensions Two of the four household types accrue entitlements to DB pensions while working. They must make pension contribution (db_t) from their earnings at each age up to 65. This is set at a fixed proportion (ϑ) of pre-tax earnings. Once they reach the age of 65, they receive a taxable pension that is a function of career-average earnings at the age of 64. ($pp_t^{db} = db(ae_{64}, j)$). This function differs by education type (j).

Defined Contribution pensions Households can, each period, pay into a Defined Contribution (i.e. 401k-style) pension. The evolution of the stock of wealth in the DC fund depends on flows into the fund (dc -which is tax-deductible and so can be made out of *gross* income) and the return on the fund in each year (ϕ).

$$DC_{t+1} = (1 + \phi_{t+1})\big(DC_t + dc_t\big) \tag{6}$$

The return on DC funds is assumed to be *iid* and normally distributed with a mean of $\bar{\phi}$ and a variance of σ^2_{ϕ}.

DC wealth is decumulated from the age of 65. At this age a quarter of the fund is taken as a (tax-free) cash lump sum - this conversion of never-taxed pension wealth into non-pension wealth is one feature that makes saving in private pensions incentivised by the tax system.[5] The remaining three-quarters of the stock of DC wealth must be used to purchase a (taxable) life annuity. The lump sum ls^{dc}_{65} is given by $(0.25)DC_{65}$ and the pension income stream for each age after 65 is given by:

$$pp^{dc}_t = q(0.75)DC_{65}$$

where q is an annuity rate that is actuarially fair up to the deduction of a fixed proportion z to account for the administrative costs and profits of the annuity-providers.

Non-pension assets Households can save and accrue non-pension wealth (a) which accumulates according to the following inter-temporal budget constraint:

$$a_{t+1} = (1 + r_t)(a_t + y_t - c^{nh}_t - dc_t - db_t) \tag{7}$$

where (r_t) is the return on non-pension wealth and y is household income (the sum of gross earnings, unemployment insurance payments, public pension payments, private pension payments and interest less taxes). The tax function is discussed in the next subsection.

3.3 Government

The government levies taxes and provides unemployment benefits and pensions.

Taxes The progressive household tax function is modelled on the prevailing UK system and is fully detailed in Appendix D.7, the discussion here focusses on the tax treatment of private pensions. Private

[5]Such tax-free lump sums are permissible in the UK.

pensions are treated favourably through a combination of full tax-deductibility of payments into pension funds as well as two features of how pension income is taxed. The first is the tax-free lump sum, noted above. The second is that after the age of 65 tax rates become substantially lower - effectively the two main rates of tax fall from 32% and 42% to 20% and 40%.[6]

Unemployment benefits Unemployment shocks are assumed to be verifiable by the government. Affected households receive an unemployment payment (ui) irrespective of their accumulated assets. Those who get an offer but who simply choose not to work can receive an asset-tested payment (ui^{mt}) if they are sufficiently poor.[7]

Public pensions Those aged over 65 are entitled to two payments. The first is a Social-Security style public pension, payable from the age of 65 until death, and is modelled as a function of career-average earnings at the age of 64 and household composition $ss_t = ss(ae_{64}, h)$). The government also provides a means-tested consumption-floor in retirement ($mtcf(a_t, y_t, t, h)$) that depends on assets, income, age and household composition. The form of these functions was illustrated in Figure 1.

3.4 State variables and the household maximisation problem

This section gives the household's maximisation problem, making explicit the state variables of the Dynamic Programming problem.

3.4.1 State variables

The state variables are household type (j), age (t), non-pension wealth (a), whether unemployed in the current period (ue), productivity ($\tilde{e}$), DC pension wealth (DC), income from the DC pension (pp^{dc}), household composition (h) and average earnings (ae)[8]. The set of state variables is: $\mathbf{X}_t = \{j, t, a_t, ue_t, \tilde{e}_t, DC_t, pp_t^{dc}, h_t, ae_t\}$. There is uncertainty over the investment return (ϕ) earned on the DC fund, over whether an employment offer is received (ue), over productivity (e), and, due to stochastic mortality, over household composition (h). Below, the joint distribution of the first three of these will be denoted as $F(\phi, ue, e)$. The stochastic processes underlying each of these are mutually independent

[6] Formally, this is due to National Insurance Contributions – the UK's equivalent to the US payroll tax - not being levied after the age of 65. The treatment of private pensions and such social insurance contributions differs in the UK and US. In neither country are social contributions levied on private pension income. However, in the US, payroll taxes are levied on earnings paid into a private pension whereas in the UK they can be made exempt from NICs.

[7] These two payments reflect the UK's 'contribution-based job-seekers allowance' and 'income-based job-seekers allowance respectively'. The latter is payable at a particularly low level and is designed to protect against destitution (it can be thought of as the UK's equivalent of food stamps).

[8] DC wealth (DC) is a state variable only up to the age of 65, DC income (pp^{dc}) is only a state variable after the age of 65. Up to and including the age of 65, the state variable ae represents average earnings up to the previous year. From the age of 66 onwards, it represents average earnings at ages up to and including 64.

though the distributions of ue and e in period $t+1$ depend on their values in period t. s^m_{t+1} and s^f_{t+1} give, respectively, the probability that a man and a woman will survive to period $t+1$ conditional on them having survived to period t.

3.4.2 Household maximisation problem and value functions

Household's problem post-annuitisation Equation (8) gives the maximisation problem and associated value function of a household with both spouses still alive and which has already annuitised its DC wealth (that is for periods $t \geq 65$). After annuitisation, households choose their non-housing consumption (c^{nh}) that whether the main household earner works (which determines leisure (l)).

$$
\begin{aligned}
V_t(\mathbf{X}_t|h_t=1) = \max_{c^{nh}_t, l_t} \Bigg(\quad & u(c_t, l_t) \quad +\beta s^m_{t+1} s^f_{t+1} \int V_{t+1}(\mathbf{X}_{t+1}|1) dF(\phi_{t+1}, ue_{t+1}, e_{t+1}|ue_t, e_t) \\
& \qquad +\beta s^m_{t+1}(1-s^f_{t+1}) \int V_{t+1}(\mathbf{X}_{t+1}|2) dF(\phi_{t+1}, ue_{t+1}, e_{t+1}|ue_t, e_t) \\
& \qquad +\beta (1-s^m_{t+1})(s^f_{t+1}) \int V_{t+1}(\mathbf{X}_{t+1}|3) dF(\phi_{t+1}, ue_{t+1}, e_{t+1}|ue_t, e_t) \\
& \qquad +(1-s^m_{t+1})(1-s^f_{t+1}) b(a^b_{t+1}) \Bigg) \\
\text{s.t.} \quad & c_t = c^{nh}_t + c^h_t \\
& \text{and the intertemporal budget constraint in equation (7)}
\end{aligned} \tag{8}
$$

Household's problem pre-annuitisation The maximisation problem and associated value function faced by a household (again with both spouses still alive) which has not annuitised its DC wealth is given in (9). The problem differs from that of the post-annuitisation problem as there is now one additional choice variable – how much to contribute to the DC pension (dc) – and there are now two intertemporal budget constraints (equations (6) and (7) - that relating to the DC fund and that relating to non-pension wealth).

$$
\begin{aligned}
V_t(\mathbf{X}_t|h_t=1) = \max_{c_t^{nh},dc_t,l_t} \Bigg(\quad u(c_t,l_t) \quad & +\beta s^m_{t+1}s^f_{t+1}\int V_{t+1}(\mathbf{X}_{t+1}|1)dF(\phi_{t+1},ue_{t+1},e_{t+1}|ue_t,e_t) \\
& +\beta s^m_{t+1}(1-s^f_{t+1})\int V_{t+1}(\mathbf{X}_{t+1}|2)dF(\phi_{t+1},ue_{t+1},e_{t+1}|ue_t,e_t) \\
& +\beta (1-s^m_{t+1})(s^f_{t+1})\int V_{t+1}(\mathbf{X}_{t+1}|3)dF(\phi_{t+1},ue_{t+1},e_{t+1}|ue_t,e_t) \\
& +\beta (1-s^m_{t+1})(1-s^f_{t+1})b(a^b_{t+1})\Bigg) \\
\text{s.t.} \quad & c_t = c_t^{nh}+c_t^h \\
& \text{and the intertemporal budget constraints in equations (6) and (7)}
\end{aligned}
\tag{9}
$$

There are no analytical solutions to the problems outlined in (8) and (9). Solutions are obtained numerically - using methods discussed in Appendix G.

4 Estimation and results

4.1 Estimation

Estimation of the model parameters follows a two-step procedure.[9] In the first step, some parameters are estimated outside the model, or are set with reference to the literature. In the second step, preference parameters and the earnings processes are estimated using the method of simulated moments. Both these steps will be described below. Before that, the next subsection briefly introduces the main data source, defines the sample used and describes how household 'types' are characterised.

4.1.1 Data, sample and definition of types

The main data used in this paper come from linked survey and administrative data. The survey data is the English Longitudinal Study of Ageing (ELSA) - a biennial longitudinal survey that contains a representative sample of the English private household population aged 50 and over. ELSA is one of a number of international 'ageing surveys' - modelled on the Health and Retirement Study (HRS) in the US. ELSA contains detailed data on demographics, labour market circumstances, earnings and the level and composition of wealth holdings. The first wave of ELSA covered 2002/03 and data from the first five waves are used in this paper.

ELSA respondents were asked for their National Insurance number (equivalent to Social Security

[9]This two-step procedure is widely applied in papers that develop and estimate structural lifecycle models. See, amongst others, Gourinchas and Parker (2002), Cagetti (2003), French (2005), De Nardi et al. (2010), Low and Pistaferri (2015) and Blundell et al. (2016).

Table 1: Proportion of each types in sample

	No DB Pension	DB Pension
Low Education	16%	32%
High Education	11%	41%

number in the US) and permission to link to their history of National Insurance contributions. Data on these contributions allows a panel of earnings in each year of working life for ELSA respondents to be obtained. 80% of individuals consented to the link. Details on how I convert these data into a panel of earnings is given in Appendix E.1. These earnings data are used, in a manner described below, to estimate earnings processes, while the survey data yields moments of assets and employment which are used to estimate preference parameters.

A sample of couples is selected in which the primary earner in the couple (the member who has the highest lifetime earnings) was born between 1935 and 1950. There are 2,364 households in this cohort. Those who never married and those who are divorced are not included. Only those couples with linked National Insurance data and where National Insurance contributions were made in at least 5 years are included in the sample. A number of additional sample restrictions are imposed. Households where the education of primary earner is not recorded are excluded, as are those where either member of the couple didn't fully complete the survey, and those where the sum of years of self-employment carried out by either member of the couple is greater than or equal to 5.[10] After applying these exclusions, 1,121 couples (47.4%) of the original sample remain.

The model contains households of four types who differ in their education (low or high) and whether they have access to a DB pension. Households are characterised as low (high) education if the primary earner left school at (remained after) the age of 15 which was the compulsory schooling age for this cohort. The split of the sample into DB/non-DB types is complicated by the fact that many households have small amounts of DB wealth (because, for example, they worked for a year or two in a job that provided a DB pension before moving on). A household is defined as being a DB pension type household if the primary earner spent at least one-third of working years contributing to a DB scheme (years accruing DB pension rights is recorded in the administrative data). Table 1 gives the proportion of the sample in each of the four types.[11]

[10]The administrative data can be used to calculate the number of years spent in self-employment, but not the income from that employment.

[11]Some caution should be exercised in interpreting these as population shares. The sample is not fully representative of the original population-representative sample. Those with less education were less likely to give permission to link to the administrative records and so are under-represented. It is not possible to assess representativeness on access to the DB pension as information on the latter is not available for those who did not give permission to link to the administrative data. More detail on differences between the sample used here and the full ELSA sample is given in Appendix E.4.

4.1.2 Parameters estimated/set outside the model

This subsection gives details of the parameters set or estimated outside the model.

Stochastic component earnings process The data generating process for u_{it} - the stochastic component of earnings given in equation (4) - contains three parameters: the coefficient of autocorrelation and the variance of the first and subsequent innovations. These are estimated using a standard approach (see, for example, Guvenen (2009) or Low et al. (2010)). The first step is to run a regression of observed wages on a quadratic in age and an indicator variable (pt) for working less than 20 hours per week[12]:

$$\ln \tilde{e}_{it}^{data} = \bar{\delta}_0 + \bar{\delta}_1 t + \bar{\delta}_2 t^2 + \bar{\delta}_3 pt + \underbrace{u_{it} + m_{it}}_{\eta_{it}} \tag{10}$$

This differs from the model's earnings process given in equation (3) in Section 3 in two ways. First, the error term is η_{it} – the sum of the stochastic component of earnings (u_{it}) and measurement error ($m_{it} \sim N(0, \sigma_m^2)$). The dependent variable is denoted $\tilde{e}^{data}$ rather than ($\tilde{e}$) to indicate that earnings are measured with error. Second, the coefficients ($\bar{\delta}_0, \bar{\delta}_1, \bar{\delta}_2$) are not the true coefficients of the earnings process ($\delta_0, \delta_1, \delta_2$) as running the simple regression shown in (10) will yield coefficients that are biased due to non-random selection into employment. These biased coefficients will be used within the model's estimation procedure to estimate the true parameters of the earnings process (discussed in Section 4.1.3). The parameters of the data generating process for η_{it} ($\rho, \sigma_\zeta^2, \sigma_\xi^2, \sigma_m^2$) are obtained by choosing those values that minimise the distance between the empirical covariance matrix of estimated residuals ($\hat{\eta}_{it}$) for ages up to 50 and the theoretical variance covariance matrix of $\eta_{it} = u_{it} + m_{it}$.

Table 2 gives the estimates of these parameters for each type. Earnings processes differ by type in two principle ways. First, highly educated households have shocks to earnings that have a variance almost twice the level of those with low education. Second, while the two high education types have very similar earnings process estimates, within the low education group, those without a DB pension have shocks to their earnings process that have higher variance and lower persistence than those with such a pension.

Table 2 also shows the unemployment rate for each type (a household is considered unemployed in the data if the primary earner is recorded as having earnings of less than that provided by the UK's unemployment insurance level). Modelled unemployment shocks evolve according to a Markov process - where the probability of entering unemployment depends on decile of current productivity. The final row in Table 2 shows average unemployment rates (the full Markov transition matrices and the probability

[12]The objective here is to recover the process for full-time equivalent earnings (as there is no option to work part-time in the model). Including the part-time dummy controls for earnings observations where an individual is observed working part-time. More details on the construction of this variable is given in Appendix E.1.

distribution over productivity after an unemployment shock are shown in Appendix D). There are significant differences between groups in rates of unemployment – those in non-DB pension careers (who are drawn disproportionately from the private sector) have substantially higher rates of unemployment than those in DB careers (who are drawn disproportionately from those who work for the government).

Table 2: Earnings process estimates

	Type			
	Low Ed		High Ed	
	No DB	DB	No DB	DB
ρ	0.970	0.993	0.995	0.990
σ^2_ξ	0.007	0.005	0.010	0.012
σ^2_ζ	0.073	0.053	0.110	0.112
σ^2_m	0.053	0.027	0.057	0.023
Unemployment rate	0.13	0.05	0.15	0.03

Coefficient of relative risk aversion I show results for a range of values for γ, the coefficient of relative risk aversion on utility[13]. In related papers that also use a utility function that is non-separable in consumption and leisure, the choice for γ has usually been between 2 and 4. My main results are for $\gamma = 3$, but results are also given for $\gamma = 2$ and $\gamma = 4$.[14]

Other parameters set outside the model The other preference parameters set outside the model are K, L, and hrs. K is the parameter that determines the curvature of the bequest function - which is set at £650,000.[15] Finally, L is the hours endowment - it is set equal to an (annual) value of 5,824 (16 hours a day to be divided between work and leisure). hrs is the number of hours worked when employed and is set equal to 1,840 (40 hours a week for 46 weeks a year).

The other parameters set outside the model relate to the model's assets. Details are given in Appendix D: the distribution of returns on Defined Contribution pensions are discussed in Appendix D.1, the function relating career-average earnings to defined benefit pension income and public pension income are given in Appendix D.3 and D.4. Appendix D.5 discusses the return on non-pension wealth and the

[13]The coefficient of relative risk aversion on consumption is

$$-\frac{\partial^2 U}{\partial c^2}\frac{c}{\frac{\partial U}{\partial c}} = -(\nu(1-\gamma)-1)$$

[14]In the literature that values social insurance over the lifecycle, a choice of $\gamma = 4$ is the most common (e.g. Auerbach and Kotlikoff (1987), Kotlikoff et al. (1999), Conesa et al. (2009), Nishiyama (2011).) The results in this paper (which suggest a greater role for mean-testing) are strengthened when this level of risk aversion is used - that is the choice of $\gamma = 3$ is a conservative one given the results.

[15] French (2005) sets this at \$500,000 in 1987 prices which is approximately £650,000 in 2012 prices and converted to pounds sterling using the average exchange rate in that year.

function that returns gross housing wealth from non-pension wealth and age (in equation (5)).

4.1.3 Method of simulated moments estimation

Method of simulated moments estimation of preference parameters Four preference parameters, each of which differ by type (j), are estimated using data on employment, wealth and portfolio composition and the method of simulated moments. These parameters are the discount factor (β_j), the weight on consumption in the utility function (υ_j), the consumption flow value of housing ($r_j^{houscon}$) and the weight on bequests (θ_j). The moments used are the proportion of men in work at each age between 52 and 75, mean non-pension holdings between the ages of 52 and 90[16] and (for those types without DB wealth) DC pension wealth between the ages of 52 and 75.[17] Wealth moments are top-coded in both data and simulations at the 95th percentile to mitigate the impact of the very wealthy. The parameters are estimated using standard GMM techniques.

It is worth noting which aspects of variation in the data will bear most heavily on which estimated parameter. Total wealth (the sum of pension and non-pension wealth) will contribute to the identification of the discount factor (β) - the greater the holdings of wealth, the more patient are households and the higher will be the estimate of β. The trajectory of wealth late in life (the extent to which it is retained rather than consumed) contribute to the identification of the strength of the bequest motive (θ) - the greater the extent to which wealth is retained rather than consumed, the more households value the leaving of bequests and the higher will be the estimate of θ. The split in wealth between non-pension and pension wealth will contribute to the identification of the consumption flow value of housing ($r_j^{houscon}$). The higher is non-pension wealth relative to pension wealth, the more value households place on the housing consumption flow, and the higher will be the estimate of the consumption value of housing wealth. Finally, labour supply profiles will pin down the relative weight of consumption in the utility function (υ). The greater the extent to which older individuals remain at work, the higher will be their preference for leisure and the lower will be (υ) – their (relative) preference for consumption.

Deterministic component of earnings process To estimate the deterministic component of the earnings process of the primary earner[18] (the parameters of equation (3)) the first step is to obtain

[16]While the cohort born between 1935 and 1950 form the basis for all other moments used in estimation, calculating these moments involves using data from individuals born before 1935. Data from older cohorts is used here as moments from the phase of life where wealth is being (or not being) decumulated is important to help identify the strength of the bequest motive. Age, period and cohort effects are estimated using the method of Deaton and Paxson (1994).

[17]The reason that moments on work after the age of 75 are not used is that the numbers are very low and don't change much - and so the additional moments do not provide much additional information. Similarly, after the age of 75, pension wealth is being decumulated in a mechanical manner (as the annuity stream becomes less valuable as fewer years of receipt are left in expectation), and so additional moments representing the pension wealth of the very old are not used.

[18]The exogenous earnings of the secondary earner are given in Appendix D.6.

the parameters in equation (10) - the regression of observed earnings on a quadratic in age. Due to non-random selection into the labour market, these parameters $\{\bar{\delta}_0, \bar{\delta}_1, \bar{\delta}_2\}$ are not those of the true productivity process. These biased parameters will be used to estimate the true parameters during the implementation of the method of simulated moments procedure using a method introduced by French (2005). Briefly, the approach involves (i) first solving the model and simulating behaviour using this (biased) profile. (ii) With the simulated data (where both accepted and rejected wage offers are observed), the bias is calculated at each age. (iii) These biases are used to 'correct' the earnings process fed into the model in (i). (iv) This corrected earnings process can then be fed back into the model which is solved and behaviour is simulated once again. Steps (ii) to (iv) are then repeated until convergence.[19]

4.2 Estimates and model fit

Table 3 gives the estimates of the preference parameters. Those with more education are found to be more patient (a frequently-found result - see Dohmen et al. (2010) and Alan and Browning (2010)). This result reflects the fact that in the data those with more education accumulate more wealth as a proportion of earnings than those with less. Figure 16 in Appendix E.3 shows the ratio of mean total wealth (including public pension wealth), at its lifecycle peak – the ratio of this measure of wealth to average lifetime earnings for the four types is 19 and 22 for the two low education types, but is substantially higher at 25 and 27 for the two high education types.

The utility function is non-separable in consumption and leisure and so every value of υ (the consumption weight in the utility function) implies a certain ratio of consumption out of work to consumption in work necessary to keep marginal utility constant. These replacement rates necessary to keep marginal utility constant when moving out of work of 69.0% and 68.7% for the two low education types but of 75.5% and 74.3% for the two high education types. These estimates of υ imply that those with more education place a greater weight on consumption relative to leisure (perhaps driven by their work tasks being, on average, less onerous, especially at older ages, and so leisure time is less important).

The results on $r^{houscon}$ show that those with a DB pension get a greater consumption value from their housing wealth than those without. The parameter on the strength of bequests (θ) can be given an intuitive interpretation by calculating its implication for the average propensity to consume in the last period of life (when death by next period is certain). The estimated parameter implies an average propensity to consume out of final period wealth of, for example, £200,000, respectively, for the four types of 0.13, 0.12, 0.21 and 0.18 respectively.

[19] French (2005) points out that if the value function were concave, it would be possible to prove that this iterative procedure is a contraction and so a unique fixed point would exist. The value function here (as in French's paper) is not concave - however, using a number of starting values, it appears that unique fixed points for each type has been found.

Table 3: Preference parameter estimates

Type	Low education		High education	
	No DB	DB	No DB	DB
β	0.978	0.970	0.999	0.986
	(0.002)	(0.002)	(0.003)	(0.002)
υ	0.422	0.417	0.512	0.499
	(0.004)	(0.002)	(0.013)	(0.004)
$r^{houscon}$	0.022	0.031	0.027	0.036
	(0.001)	(0.002)	(0.013)	(0.002)
θ	0.020	0.008	0.029	0.008
	(0.002)	(0.001)	(0.013)	(0.002)
$\chi^2 stat.$	785	1171	230	653
(df)	83	59	83	59

The degrees of freedom in the χ^2 test differ across types as DC wealth moments are not used in the estimation procedure.

While the model is formally rejected by tests of overidentification for each type (shown, with associated degrees of freedom in the bottom two rows of the table), Figures 3 (labour supply), 4 (non-pension wealth) and 5 (DC wealth) show that the model closely fits the economically-important features of the data. These include the decline in labour supply at older ages, the fact that households tend not to decumulate their non-pension wealth in the UK, and the fact that those with Defined Benefit pensions accumulate very little Defined Contribution pension wealth.

Figure 3: Fit: Labour supply

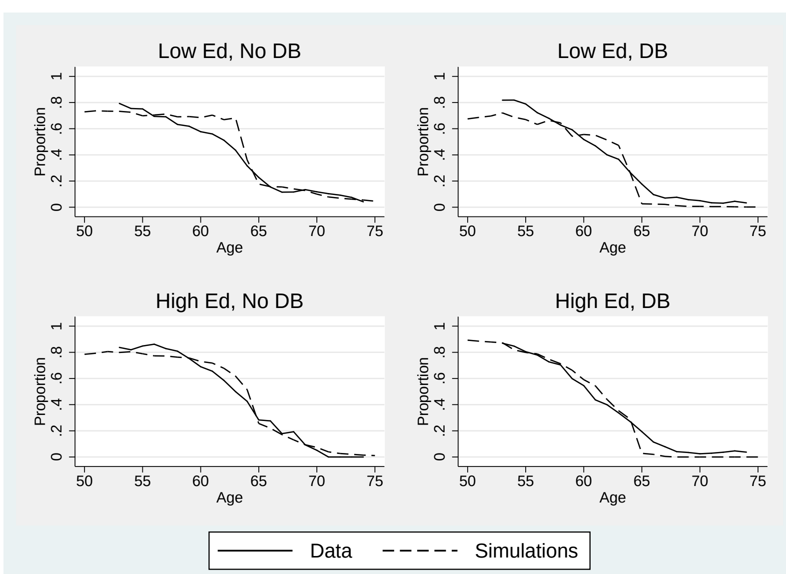

Figure 4: Fit: Non-pension wealth

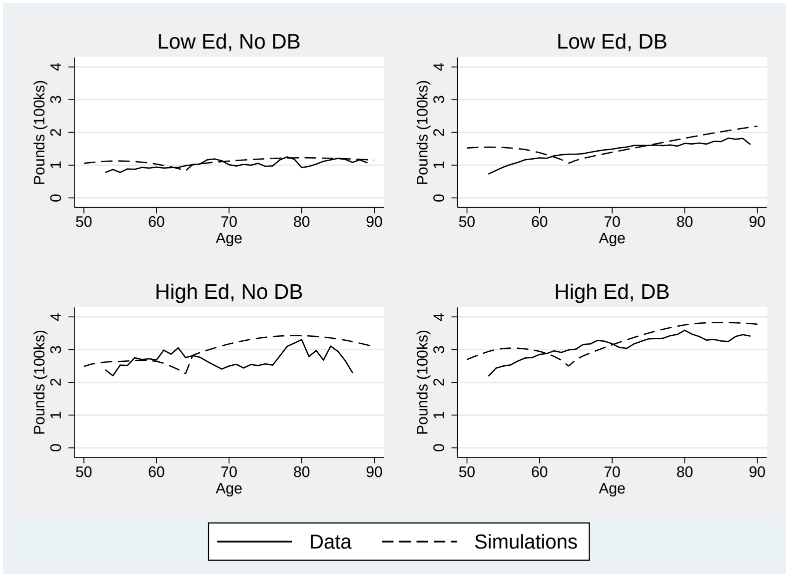

Figure 5: Fit: DC pension wealth

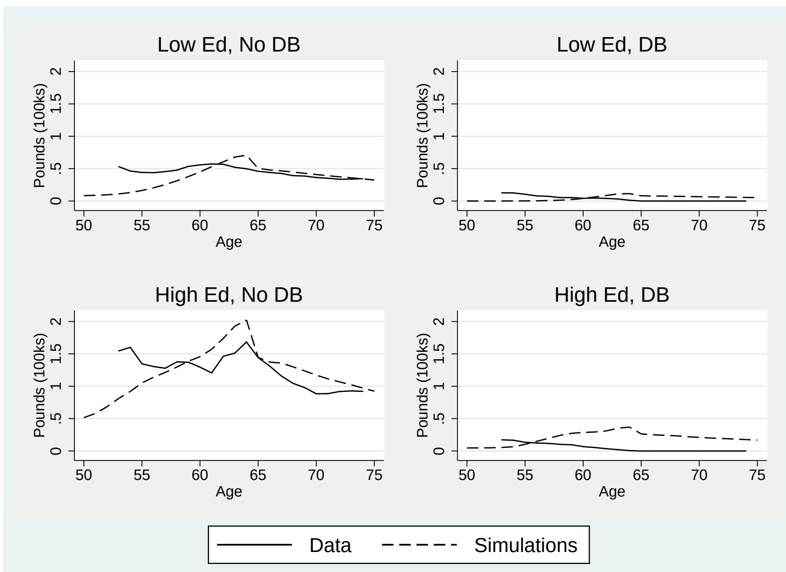

The data profiles in Figures 3 to 5 are the moments that have been used to estimate the preference parameters. Appendix H shows how model simulations compare to data not used in the estimation procedure. Figures 17 and 18 show the 25th and 75th percentiles of modelled wealth (the sum of Defined Contribution pension wealth and non-pension wealth) and Figure 19 shows consumption over the lifecycle. That the match is close in those figures also indicates that the model can capture the heterogeneity that exists in wealth accumulation and consumption growth.

5 Counterfactual analysis

This section uses the estimated model and an optimal taxation framework to obtain the ex-ante optimal means-tested old-age consumption floor. Before turning to that, the manner in which welfare changes are measured is briefly outlined.

5.1 Measuring Welfare Changes

The value function $V_t^j()$, outlined in equation (9) gives expected utility in time t for type j as a function of state variables in time t. $V_1^j()$ is therefore expected lifetime utility. Expected utility can also be expressed as an indirect utility function – a function of the optimal choices over those objects which give households utility (consumption, leisure and leaving bequests). Let $\mathfrak{s} = (\mathfrak{s}_1, \mathfrak{s}_2, \ldots, \mathfrak{s}_T)$ be the set of possible states of the world in every time period. Each element in $\mathfrak{s}_t$ gives the history of realisations of stochastic variables up to and including period t. Define household policy functions for each of consumption, leisure and bequeated assets as $c(\mathfrak{s}), l(\mathfrak{s}), a^b(\mathfrak{s})$. These functions, which give optimal behaviour as a function of the state of the world, are obtained in solving the Dynamic Progamming problem. The expected utility function can be expressed as:

$$\tilde{V}^j\big(c(\mathfrak{s}), l(\mathfrak{s}), a^b(\mathfrak{s})\big)$$

A policy change (for example a reform of the pension system) leaves that function unchanged but it will take different values of the arguments as households re-optimise in response to reform. Denoting the new policy functions as: $c^1(\mathfrak{s}), l^1(\mathfrak{s}), a^{b1}(\mathfrak{s})$, the new level of expected utility is:

$$\tilde{V}^j\big(c^1(\mathfrak{s}), l^1(\mathfrak{s}), a^{b1}(\mathfrak{s})\big)$$

The difference in welfare induced by the reform can be expressed by finding a 'consumption-leisure-bequest' equivalent variation – the proportionate change (Δ_j) in all *pre-reform* quantities that yield *post-reform* expected utility to households:

$$\tilde{V}^j\Bigg((1+\Delta_j)c(\mathfrak{s}), (1+\Delta_j)l(\mathfrak{s}), (1+\Delta_j)a^b(\mathfrak{s})\Bigg) = \tilde{V}^j\Bigg(c^1(\mathfrak{s}), l^1(\mathfrak{s}), a^{b1}(\mathfrak{s})\Bigg) \tag{11}$$

In evaluating any change from baseline policy, the social welfare function will be an equally weighted average of these Δs (so a utilitarian social welfare function is applied along with an assumption of a population containing equal proportions of these types[20]):[21]

[20] Recall the caveat around interpreting the shares in Table 1 as population shares.

[21] Common in the related literature (see for example Conesa et al. (2009) or Low et al. (2010)) is to express utility differences

$$W = \frac{1}{4}\Sigma_{j=1}^{4}\Delta_j \tag{12}$$

5.2 Optimal consumption floor

The estimated model can be used to solve for an optimal retirement consumption floor. To illustrate the types of changes that this experiment considers, Figure 6 illustrates how the mapping from pre-consumption floor retirement income to post-consumption floor income would change with each of a 25% decrease and a 25% increase in the consumption floor.[22]

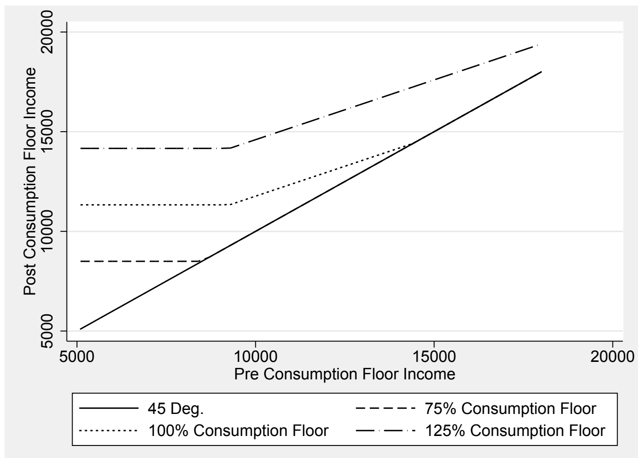

Figure 6: Illustration of changes to consumption floor

Government's problem The government's problem, outlined formally in Appendix I, is to maximise the social welfare function (equation 12) by choosing the level of the consumption floor (as depicted in Figure 6). This must be done subject to no change in the government budget balance over the lifecycle of this cohort. Different instruments are used (in different experiments) to balance the budget - these are i) changes in the basic rate of tax[23], ii) equal proportionate changes in both the UK's main income tax rates, iii) a proportional tax/subsidy on all income and iv) changes in the contributory public pension.

Table 4 reports the results of government's problem when the budget is balanced using the basic rate of tax (results using other instruments balancing the budget are shown later). To maximise the social welfare function, it is optimal to increase the consumption floor by 60.5% (bringing if from £11,300 for

as a consumption equivalent variation (CEV) - the proportionate increase in (only) consumption in each state of world that would obtain the expected utility post-reform. This is not a sensible measure when a social welfare function averages utilities across households with different preferences for consumption. For a given utility difference, the CEV will tend to be larger the lower is the consumption weight on utility (as the less valuable is consumption to an agent the more additional consumption that will be needed to obtain a particular level of expected utility). Those household types with the lowest weights on consumption would bear most heavily on social welfare function. To avoid the government's objective function having this characteristic, a consumption/leisure/bequest equivalent variable (Δ) is used instead.

[22] In all cases, the point at which the effective marginal tax rate falls from 100% to 40% is left unchanged.

[23] This is a tax rate levied at 20% on earnings between the 15th and 85th percentile of earnings.

Table 4: Optimal consumption floor - budget balance using basic rate of tax

	Type			
	Low Ed		High Ed	
	No DB	DB	No DB	DB
Change in consumption floor	60.5%			
Change in basic rate of tax	20% $\rightarrow$ 23.4%			
$100\bar{\Delta}$	0.16			
$100\Delta_j$	0.86	-0.04	0.36	-0.54

couples to just over £18,000 and from £7,420 for singles to £11,870). This requires an increase in the basic rate of tax from 20% to 23.4% to balance the budget. Expected lifetime utility increases by 0.16%. This is the average across all four types – the increase in expected utility before even their type is realised. The change in expected utility once a household's type (but no other uncertainty) is realised is given in the final row of the table. There are welfare gains for each of the household types who don't have a DB pension. Among those who have a DB pension, the low educated households experience a very small fall in expected utility, with the high educated households experiencing utility losses of just over half a percent of lifetime expected utility.

These welfare effects are generated by falls in consumption inequality and increases in leisure (the 'good news'), partly offset by falls in average consumption and in bequests left (the 'bad news'). Table 5 shows this.[24] Taking the 'good news' first: household expected utility increases due to additional leisure time (row (1)), and falls in the variance of consumption (row (3)). These increases in utility are offset by the fact average consumption for the two highly educated types falls (row (2)) and bequests fall for each type as wealth at older ages falls (row (4)).

Grouping households of all types together, Figure 7(a) illustrates the change in the distribution of modelled wealth[25] induced by the move to the optimal system. It shows the 10th percentile, median and 90th percentile of wealth before the reform (solid line) and after it (dashed line). The 10th percentile falls to zero by the age of 65 and is close to zero for most of the lifecycle – the new consumption floor almost completely eliminates the incentive for those at the bottom to save. Wealth at the median is lower post-reform and is decumulated faster as households have a greater incentive to run down their wealth and rely on the more generous consumption floor. Wealth at the 90th percentile does not fall faster -

[24]Appendix J formally decomposes the total welfare effect into contributions coming from changes in the distributions of each of consumption, leisure, bequests.

[25]This is the wealth in the model that is accumulated endogenously: the sum of non-pension wealth and DC wealth; DB pension wealth is not included.

Table 5: Welfare effect – channels

		Low Ed				High Ed			
		No DB		DB		No DB		DB	
	Quantity	Pre	Post	Pre	Post	Pre	Post	Pre	Post
(1)	Average years of work	35.9	31.8	34.3	31.4	37.5	36.1	39.3	38.1
(2)	Average consumption (£1000s)	16.4	16.7	17.8	17.9	23.0	22.6	27.4	26.9
(3)	Var(log consumption)	0.11	0.07	0.10	0.07	0.14	0.12	0.12	0.12
(4)	Bequests (£1000s)	110	47	219	153	283	204	362	310

Row (3) shows the average (over years) of the annual variance of the log of consumption.

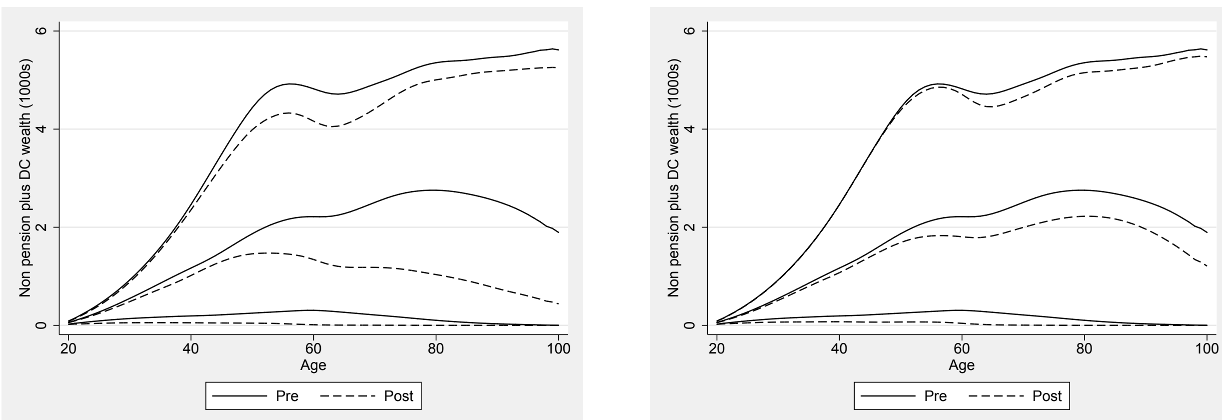

(a) Current tax treatment of private pensions

(b) Reformed tax treatment of private pensions

Figure 7: Modelled asset profiles under optimal systems

households who accumulated very high wealth stocks pre-reform are not attracted by the prospect of relying even on the more generous consumption floor. Wealth at these upper percentiles is lower over the whole of the lifecycle due to higher tax rates of income tax.

The effects on the wealth distribution induced by this reform are large (as were the effects of labour supply shown in Table 5). The next section shows how implementing an optimal consumption floor alongside changes in the treatment of private pensions can induce smaller distortions and yield substantially larger increases in welfare.

Changing tax treatment of private pensions The counterfactual experiment reported above kept unchanged the tax treatment of private pension saving. The tax treatment in the UK has two particular features - i) tax-deductibility of contributions from earnings and ii) lower taxation of pension income relative to earnings. These features are expensive – they cost over 1.1% of GDP. Table 6 shows the optimal consumption floor when the second of these two features is removed. This would move the UK

Table 6: Optimal consumption floor - reformed tax relief budget balance using basic rate of tax

	Type			
	Low Ed		High Ed	
	No DB	DB	No DB	DB
Change in consumption floor	34.1%			
Change in basic rate of tax	20% → 18.6%			
$100\bar{\Delta}$	0.61			
$100\Delta_j$	0.96	0.59	0.58	0.33

to a system which retains tax-deductibility of private pension contributions, but taxes pension income and earnings equivalently.[26] The additional tax revenue raised can be divided between a more generous retirement consumption floor and changes in the basic rate of tax. The optimal system increases the consumption floor by 34.1% and reduces the basic rate of tax to 18.6%. The welfare gains here are equal to 0.61% of lifetime utility, substantially more than the gains of 0.16% of lifetime utility reported in Table 4 when moving to the optimal consumption without reforming the taxation of private pension saving. Furthermore, in this case each of the four types gains from the reform. Figure 7(b) shows that the change in the distribution of modelled wealth induced by the move to this system is more modest than those when the optimal consumption floor, retaining private pension subsidies, is imposed.

It is worth exploring further the reason for the substantially higher increase in the consumption floor when private pension subsidies are retained (60.5%) compared to when they are removed (34.1%). Some intuition for what is driving this result can be obtained by looking directly at how the government budget balance responds to changes in the level of consumption floor. Figure 8 shows the change in government balance per household (on the vertical axis) for different proportional changes in the consumption floor (horizontal axis), keeping all other features of the tax system unchanged. Changes both allowing for and not allowing for behavioural responses are shown. In spite of the fact that removing the consumption floor leads to greater labour supply and tax revenue, the improvement in the government balance when the consumption floor is abolished is more modest when the effect of behavioural responses are incorporated. The explanation for this lies in the tax advantages associated with private pension saving. Reductions in government-provided insurance against low consumption in retirement lead households to self-insure by increasing their (subsidised) private pension saving. The additional government spending on subsidising

[26] In particular, the changes made are to remove the tax-free lump sum, to levy National Insurance Contributions on pension income, and to remove the fact that tax-free allowances (that portion of income which is exempt from income tax) are higher for those over the age of 65. This would bring the treatment of private pension saving in the UK close to that in the US.

this self-insurance partially offsets the revenue saved by removing the consumption floor.

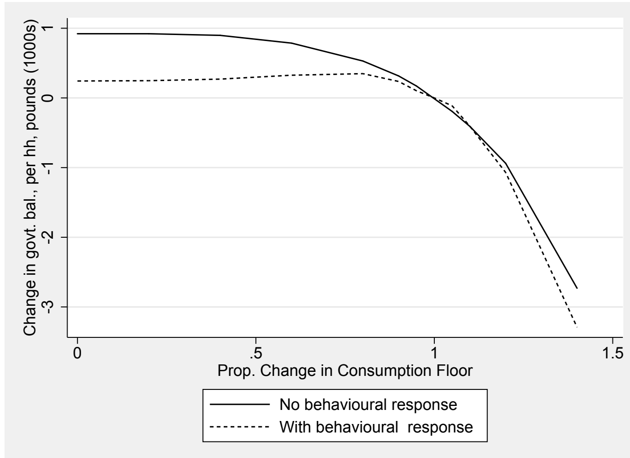

Figure 8: Fiscal implications of changing the consumption floor

Retirement consumption floors and private pension subsidies are therefore substitutes in the government budget constraint: when the generosity of the former is increased (reduced), government spending on the latter falls (rises). Therefore, there are two distinct reasons for the very large (60.5%) optimal increase in the consumption floor when private pension subsidies are left in place. The first is straightforward - it is that consumption floors provide valuable insurance to households at a cost that households are prepared to pay. The second is that increasing consumption floors moves government spending on pensions from private pension subsidies towards consumption floors, which households prefer.

Balancing the budget through other means The results so far have been for the case when the government budget is balanced using changes in the basic rate of tax. Table 7 shows equivalent results when the budget is balanced through reductions in the contributory public pension (panel (1)), through equal proportionate increases in each of the basic and higher rates of tax (panel (2)), and through a proportional tax/subsidy on all net income (panel (3)). In all cases, the optimal system under the current tax treatment of private pensions and the reformed system that treats earnings and pension income equivalently is shown.

When the budget is balanced using the social-security style contributory public pension, the result is an extreme one - the optimal system involves abolishing the pension and using the resulting funds to substantially increase the consumption floor. This is how public pensions are delivered in Australia: a means-tested pension is provided but there is no contributory public pension.

When both higher and basic rates of tax are used (increasing both in equal proportional terms) to balance the budget (panel (2)), the optimal increase in the consumption floor is smaller than in the base case. This is due to the fact that the labour supply effects of higher taxation are larger when higher earnings are more heavily taxed as they would be by this reform.

The case where the budget is balanced by taxing/subsiding all income by a fixed proportion (panel (3)) and private pension subsidies are removed is the one scenario examined where the optimal design problem suggests a *lowering* of the consumption floor. The reason for this is that increasing the consumption floor would involve reducing consumption possibilities in all states of the world, including, for example, those periods of working life when incomes are very low and the marginal utility of consumption is highest.[27] The utility cost of such taxes are large and the optimal policy is to reduce the consumption floor by 16.9%, allowing an increase in 1.8% of net income in each period. If private pension subsidies are left in place an increase in the consumption floor is optimal (in this case, consumption floors are less expensive as part of their cost is mitigated by reductions in household take-up of private pension subsidies).

Table 7: Balancing budget through other instruments

	(1)		(2)		(3)	
	Cont. Pub. Pen.		Both tax rates		Prop. tax (subsidy)	
Private pension treatment	Base	Reform	Base	Reform	Base	Reform
Change in consumption floor	67.9%	89.2%	40.3%	30.7%	38.9%	-16.9%
Change in contributory pension	-100.0%	-100.0%	-	-	-	-
New basic rate of tax	-	-	21.6%	18.3%	-	-
New higher rate of tax	-	-	43.2%	36.6%	-	-
Proportional tax (subsidy)	-	-	-	-	0.7%	-1.8%
Ave. Welf. Change ($100\bar{\Delta}$)	0.29	0.70	0.14	0.64	0.10	0.67
Low Ed, No DB	1.32	1.77	0.66	0.93	0.59	0.31
Low Ed, DB	0.18	0.45	0.07	0.61	0.03	0.77
High Ed, No DB	0.24	0.98	0.13	0.59	0.08	0.71
High Ed, DB	-0.60	-0.40	-0.30	0.42	-0.30	0.87

Sensitivity to risk aversion assumption and other private pension arrangements To investigate the sensitivity of the results suggested here, Table 8 shows the results from some alternative design problems. Column (1) reproduces the results from Table 6, that is, it shows the optimal consumption floor when a) private pension income is taxed equivalently to earnings and b) the budget is balanced using the basic rate of tax. Columns (2) and (3) show the optimal consumption floor when alternative values for the coefficient of relative risk aversion ($\gamma = 2$ and $\gamma = 4$ respectively) are used.[28] These can be compared to the baseline results in column (1) which are for $\gamma = 3$. The optimal level of the consumption

[27]Households in this situation are not affected by the basic rate of tax changes which exempts the first £8,000 of households from tax.

[28]The other preference parameters are re-estimated with the new levels of risk aversion imposed.

floor increases with risk aversion but even with a coefficient of relative risk aversion of 2[29], the current retirement consumption floor is found to be too low.

Column (4) assesses the relevance of the results for future cohorts - for whom Defined Benefit pensions will be less prevalent than for the cohort whose behaviour is studied here. The exercise conducted removes DB pension entitlement from those types who held them and also (as well as removing their obligation to pay for DB pension). These households are allowed to make their own arrangements for retirement by saving in a DC pension. It is not clear *a priori* whether the optimal consumption floor will be higher or lower in a post-DB world. Households who previously had a DB pension are now exposed to investment risk on their DC funds (a risk against which they were previously insured by their employer). This will make consumption floors (which provide insurance against these risks) more valuable. However, their pension saving will now be more elastic (previously their assumed DB pension was assumed not to change with changes in the policy environment), which will increase the distortions induced by and taxes needed to fund changes in the consumption floor. The latter effect dominates and the optimal increase in the retirement consumption floor is 22.1% (compared with 34.1% in the DB world).

The final column removes the tax deductibility of private pension savings - this makes the taxation of savings in a private pension equivalent to the taxation of non-pension savings. Tax must be paid on all income in the period that it is earned. The results here are very similar to those given in column (1) – the increase in the optimal consumption floor is slightly higher, the basic rate of tax is slightly lower, and the change in welfare is of a similar magnitude.

Table 8: Additional Results

	(1)	(2)	(3)	(4)	(5)
	Reform	$\gamma = 2$	$\gamma = 4$	Post DB	No tax deferral
Change in consumption floor	34.1%	7.8%	67.5%	22.1%	36.0%
Basic tax rate	18.6%	18.3%	21.6%	18.7%	17.7%
Ave. Welf. Change ($100\bar{\Delta}$)	0.61	0.32	0.63	0.41	0.69
Low Ed, No DB	0.96	0.32	1.49	0.62	1.17
Low Ed, DB	0.59	0.38	-0.05	0.42	0.41
High Ed, No DB	0.58	0.34	1.11	0.29	0.82
High Ed, DB	0.33	0.24	-0.05	0.31	0.38

[29]This coefficient of relative risk aversion on utility implies a coefficient of relative risk aversion on consumption for the different types that varies between 1.4 and 1.5, which is at the lower end of estimated values.

An all-age consumption floor? The results presented in this section have suggested that there is scope for revenue-neutral welfare-increasing increases in the consumption floor paid to those in old age. It is worth investigating the extent to which these results suggest a role for a more generous consumption floor for all ages. The advantages of an all-age consumption floor are similar to those for the elderly - they transfer resources from good states of the world (high income/high consumption states) to bad (low income/low consumption states). Once again, the costs are a diminished incentive to supply labour and a crowding-out of household saving.

There currently exists an all-age consumption floor for those not working in the UK that is less generous than the retirement consumption floor[30] which was denoted ui^{mt} in Section 3.3. I use the framework that generated the results in Table 4 to assess the optimal level of this consumption floor. This suggests that the current all-age consumption floor (already substantially lower than the old-age consumption floor) is too high – the optimal level would be 16% lower, which would allow a reduction of the 0.5% in the basic rate of tax. The reason for this is that the negative effect on labour supply is much greater when providing an all-age consumption floor than it is when providing an old-age consumption floor. These results emphasise the value of providing good incentives to work for the young alongside pensions with good insurance properties for the old.

6 Conclusion

Providing public pensions is one of the costliest activities undertaken by governments in the developed world and a variety of different approaches are observed internationally. Many provide replacement rates that are either proportional to, or vary progressively with, career-average earnings while others emphasise means-testing by providing pensions to those who have income below a certain level. This paper investigates how governments should structure their public pension schemes and how they should tax private pensions.

The analysis suggests that private pension subsidies (such as tax rates on pensions that are lower than those on earnings) and pensions that are related to lifetime-earnings should be replaced with a combination of higher means-tested consumption floors in retirement and lower taxes on earnings. Private pension subsidies and earnings-related pensions project lifetime earnings risk into retirement. Means-tested consumption floors, on the other hand, provide valuable insurance against such earnings risks, as well as against investment risk and longevity. While there are distortions induced by providing an entirely means-tested public pension, these are more than offset by the value of this insurance. Old-age

[30] Known as 'income-based job-seekers allowance' it guarantees (in the 2012/13 tax and benefit system) an income of just under £5,800 to couples. One can think of this as roughly analogous to the Food Stamps in the US.

is therefore a good time to provide insurance against lifetime productivity risk – calling into question the rationale for policies that, by amplifying those risks, do the opposite of this.

References

Alan, S. and M. Browning (2010). Estimating intertemporal allocation parameters using synthetic residual estimation. *The Review of Economic Studies 77*(4), pp. 1231–1261.

Altonji, J. G. and L. M. Segal (1996). Small-sample bias in gmm estimation of covariance structures. *Journal of Business & Economic Statistics 14*(3), 353–366.

Auerbach, A. J. and L. J. Kotlikoff (1987). *Dynamic fiscal policy.* Cambridge University Press.

Bank of England (2013). Interest rate data. Accessed 19Nov2013.

Banks, J., C. Emmerson, Z. Oldfield, and G. Tetlow (2005). *Prepared for retirement? The adequacy and distribution of retirement resources in England.* Number R67. IFS Reports, Institute for Fiscal Studies.

Barclays Capital (2012, March). The equity gilt study 2012. Technical report, Barclays Capital, London.

Blau, D. M. (2016). Pensions, household saving, and welfare: A dynamic analysis of crowd out. *Quantitative Economics 7*(1), 193–224.

Blundell, R., M. Costa Dias, C. Meghir, and J. Shaw (2016). Female labor supply, human capital, and welfare reform. *Econometrica 84*(5), 1705–1753.

Bozio, A., R. Crawford, C. Emmerson, and G. Tetlow (2010). Retirement outcomes and lifetime earnings: descriptive evidence from linked elsa-ni data. DWP Working Papers 81, Department for Work and Pensions.

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.

Braun, R. A., K. A. Kopecky, and T. Koreshkova (2016). Old, sick, alone, and poor: a welfare analysis of old-age social insurance programs. *Review of Economic Studies*.

Cagetti, M. (2003). Wealth accumulation over the life cycle and precautionary savings. *Journal of Business & Economic Statistics 21*(3), 339–353.

Chetty, R., J. N. Friedman, S. Leth-Petersen, T. H. Nielsen, and T. Olsen (2014). Active vs. passive decisions and crowd-out in retirement savings accounts: Evidence from denmark. *The Quarterly Journal of Economics 129*(3), 1141–1219.

Conesa, J. C., S. Kitao, and D. Krueger (2009). Taxing capital? not a bad idea after all! *American Economic Review 99*(1), 25–48.

Conesa, J. C. and D. Krueger (1999). Social security reform with heterogeneous agents. *Review of Economic dynamics 2*(4), 757–795.

Crawford, R. and C. O'Dea (2014, Sep). Cash and pensions: Have the elderly in england saved optimally for retirement.

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*(1), 39–75.

De Nardi, M., S. Imrohoroğlu, and T. J. Sargent (1999). Projected us demographics and social security. *Review of Economic dynamics 2*(3), 575–615.

Deaton, A. S. and C. Paxson (1994). Saving, growth, and aging in taiwan. In *Studies in the Economics of Aging*, pp. 331–362. University of Chicago Press.

Dohmen, T., A. Falk, D. Huffman, and U. Sunde (2010). Are risk aversion and impatience related to cognitive ability? *The American Economic Review 100*(3), pp. 1238–1260.

French, E. (2005, 04). The effects of health, wealth, and wages on labour supply and retirement behaviour. *Review of Economic Studies 72*(2), 395–427.

Golosov, M., A. Shourideh, M. Troshkin, and A. Tsyvinski (2013). Optimal pension systems with simple instruments. *The American Economic Review 103*(3), 502–507.

Gourinchas, P.-O. and J. A. Parker (2002, January). Consumption over the life cycle. *Econometrica 70*(1), 47–89.

Greene, W. (2004, 06). The behaviour of the maximum likelihood estimator of limited dependent variable models in the presence of fixed effects. *Econometrics Journal 7*(1), 98–119.

Guvenen, F. (2009, January). An empirical investigation of labor income processes. *Review of Economic Dynamics 12*(1), 58–79.

HMRC (2015). Registered pension schemes: cost of tax relief. mimeo., Her Majesty's Revenue and Customs.

Hood, A. and L. Oakley (2014, Nov). A survey of the gb benefit system. Technical report.

Hubbard, R. G., J. Skinner, and S. P. Zeldes (1995, April). Precautionary saving and social insurance. *Journal of Political Economy 103*(2), 360–99.

Huggett, M. and J. C. Parra (2010). How well does the us social insurance system provide social insurance? *Journal of Political Economy 118*(1), 76–112.

Huggett, M. and G. Ventura (1999). On the distributional effects of social security reform. *Review of Economic Dynamics 2*(3), 498–531.

Kitao, S. (2014). Sustainable social security: Four options. *Review of Economic Dynamics 17*(4), 756–779.

Kotlikoff, L. J., K. Smetters, and J. Walliser (1999). Privatizing social security in the united states - comparing the options. *Review of Economic Dynamics 2*(3), 532–574.

Low, H., C. Meghir, and L. Pistaferri (2010, September). Wage Risk and Employment Risk over the Life Cycle. *American Economic Review 100*(4), 1432–67.

Low, H. and L. Pistaferri (2015, October). Disability insurance and the dynamics of the incentive insurance trade-off. *American Economic Review 105*(10), 2986–3029.

Nardi, M. D. (2004). Wealth inequality and intergenerational links. *The Review of Economic Studies 71*(3), pp. 743–768.

Nationwide Building Society (2014). House price index data. Accessed 1Feb2014.

Nelder, J. A. and R. Mead (1965). A simplex method for function minimization. *The computer journal 7*(4), 308–313.

Neumark, D. and E. Powers (1998, May). The effect of means-tested income support for the elderly on pre-retirement saving: evidence from the ssi program in the u.s. *Journal of Public Economics 68*(2), 181–206.

Neumark, D. and E. Powers (2000, October). Welfare for the elderly: the effects of ssi on pre-retirement labor supply. *Journal of Public Economics 78*(1-2), 51–80.

Neyman, J. and E. L. Scott (1948). Consistent estimates based on partially consistent observations. *Econometrica 16*(1), pp. 1–32.

Nishiyama, S. (2011). The budgetary and welfare effects of tax-deferred retirement saving accounts. *Journal of Public Economics 95*(11), 1561–1578.

Nishiyama, S. and K. Smetters (2007). Does social security privatization produce efficiency gains? *The Quarterly Journal of Economics*, 1677–1719.

Pope, T. and B. Roantree (2014, Jan). A survey of the uk tax system. Institute for Fiscal Studies Briefing Note 9, Institute for Fiscal Studies.

Sefton, J. and J. Van De Ven (2009). Optimal design of means tested retirement benefits*. *The Economic Journal 119*(541), F461–F481.

Tauchen, G. (1986). Finite state markov-chain approximations to univariate and vector autoregressions. *Economics Letters 20*(2), 177–181.

US Department of the Treasury (2013). Fiscal year 2014 tax expenditures. mimeo., US Department of the Treasury.

A Appendix to section 2

United Kingdom The cost of tax relief on private pension saving less the tax revenue raised on private pensions was £21bn in 2013/14 (HMRC (2015)) or 1.1% of GDP

United States The cost of tax expenditures on private pensions was $152bn in 2014 (or 0.9% of GDP). This is the sum (from Table 3 of US Department of the Treasury (2013)) of tax expenditures on Defined Contribution Plans, Defined Benefit Plans and Individual Retirement Accounts.

B Parameter definitions

Table 9 summarises the parameters that enter the model.

Table 9: Parameter definitions

Preference Parameters	
β	Discount factor
ν	Consumption weight in util. func.
$r^{h_{houscon}}$	Value of housing in util. func.
θ	Weight on bequest
K	Determinant of curvature of bequest
γ	Coeff. of Rel. Risk. Aversion (util.)
Labour market	
π_0	Prob. remaining in unemployment
$\pi_1(\tilde{e})$	Prob. entering unemployment
$\{\delta_i\}_{i=0}^{2}$	Params. of det. component of earn. proc.
$\{\bar{\delta}_i\}_{i=0}^{2}$	Biased (due to selection) $\{\delta_i\}_{i=0}^{2}$
u	Stoch. component of earn. proc.
ρ	Autoregressive param. in stoch. earn.
ξ	Innovation to autoregressive component
σ_ζ	Variance of first innovation stoch. earn
σ_ξ	Variance of subsequent innovations
m	Measurement error in earnings
σ_m^2	Variance of measurement error
η	Stoch. earnings plus meas. err. $(u+m)$
e	Primary earner earnings
e^s	Secondary earner earnings
t^{ret_s}	Retirement age of secondary earner
e^{data}	Primary earnings in data (inc. meas. error)
$E()$	Post-unemployment productivity dist.
Assets	
a	Non pension wealth
r_t	Return on pension wealth
DC	DC wealth stock
ϕ	DC wealth return realisation
$\bar{\phi}$	DC wealth mean return
σ_ϕ^2	DC wealth return variance
ls^{dc}	DC tax free lump sum
q	Annuity rate (net of admin load)
z	Administrative load
db	DB pension contribution
ς	DB contribution earnings proportion
pp^{db}	DB pension income
ls^{db}	DB lump sum
ϑ	DB lump sum multiplier on (pension income)
$db(ae_{64})$	DB income function (on average earnings at 64)

State variables	
$\mathbf{X}$	Vector of all state variables
j	Household type
t	Age
a	Non-pension Wealth
ue	Unemployment shock
$\tilde{e}$	Productivity
DC	DC pension wealth
pp^{dc}	DC pension income
h	Household composition
ae	Average Earnings
Household choices	
c^{nh}	Non-housing consumption
l	Leisure
dc	DC pension contributions
Utility function and arguments	
$u()$	Single period utility function
c^h	Housing consumption
c	Consumption $(= c^h + c^{nh})$
L	Endowment of hours
hrs	Hours of work when employed
$b()$	Bequest function
a^b	Assets bequeathed
$V()$	Value Function
Taxes and transfers	
y	Net income
ss	Public pension income
$mtcf$	Means-tested consumption floor
ui	Unemp. insurance to those with no offer
ui^{mt}	Means-tested unemp. insurance.
Other	
s_{t+1}	Surv. prob. to $t+1$, (cond. on surv to t)
$F()$	Dist. over $ue, \tilde{e}, \phi$
Counterfactual Analysis	
$\mathfrak{s}_t$	History of stochastic realisations to period t
Δ	Cons, Leis, Bequest equiv. variation (CLBEV)
$\tilde{V}()$	Val. Function in indirect utility form

The remaining appendices are incomplete. Please email me (at cormac.odea@ifs.org.uk) if you would like a draft of the appendices.