THE UNIVERSITY OF MELBOURNE

FACULTY OF BUSINESS & ECONOMICS

Melbourne Institute Working Paper Series

Working Paper No. 19/13

The Influence of Design Costs on Investments in Individual Savings Accounts

Justin van de Ven

MELBOURNE INSTITUTE®
of Applied Economic and Social Research

The Influence of Decision Costs on Investments in Individual Savings Accounts*

Justin van de Ven
National Institute of Economic and Social Research, London; and Melbourne Institute of Applied Economic and Social Research, The University of Melbourne

Melbourne Institute Working Paper No. 19/13

ISSN 1328-4991 (Print)

ISSN 1447-5863 (Online)

ISBN 978-0-7340-4309-2

June 2013

* This study was made possible by funding from the Economic and Social Research Council, UK Treasury, and UK Revenue and Customs. The usual disclaimer applies.

Melbourne Institute of Applied Economic and Social Research
The University of Melbourne
Victoria 3010 Australia
***Telephone* (03) 8344 2100**
***Fax* (03) 8344 2111**
***Email* melb-inst@unimelb.edu.au**
***WWW Address* http://www.melbourneinstitute.com**

Abstract

This study considers the efficacy of a tax incentivised savings scheme in context of decision making rigidities. Analysis is based on a classical life-cycle model of savings and investment decisions, augmented with a salience cost over participation in Individual Savings Accounts (ISAs) currently run in the UK. Calibration results indicate that salience costs help to match the model to observed rates of participation in ISAs. The calibrated model suggests that the price effects of ISAs are insufficient to generate appreciable increases in private sector savings, with or without salience costs. In this context, salience costs have an important influence on the distribution of welfare benefits that are delivered by the ISAs scheme.

JEL classification: D12, D14, D81, H31, G02

Keywords: Decision costs, saving, uncertainty, pensions

1 Introduction

Concerns regarding the medium term sustainability of welfare provisions in context of aging populations have motivated interest in policy alternatives designed to raise household sector savings rates. In this context, there is a strong political economy appeal to policies that encourage households to save more voluntarily, relative to compulsory savings schemes. This appeal is, however, off-set by attendant uncertainty regarding the efficacy of voluntary savings schemes, particularly where such schemes rely for their success upon active public engagement. In this study we explore the issues involved with reference to a tax incentivised savings scheme currently operated in the UK.

Individual Savings Accounts (ISAs) are innovative tax-advantaged savings vehicles operated in the UK.[1] In common with many schemes of its kind, investment returns on assets held in an ISA (capital gains and income) are shielded from income tax, both when they are earned and upon withdrawal. The innovative aspect of the policy design for ISAs concerns the limits to which the scheme is subject. Annual contributions to an ISA – made out of post tax income – are subject to modest upper limits, whereas there are no limits imposed on either the aggregate value of the investments that can be held within an ISA, or when the balance can be withdrawn.[2] These terms have advantageous properties in context of a rational population: the upper bounds on annual contributions and the omission of a limit on aggregate account size encourage long-term investment horizons and exaggerate savings motives early in the life course, whereas allowing the full balance of an ISA to be withdrawn at any time omits the disincentives to participate that are associated with pension fund illiquidity.

ISAs have grown steadily in popularity and importance since the scheme's introduction on 6 April 1999, supporting the conjecture that they have been effective in raising household savings rates. The number of ISAs subscribed to in each year has increased from 9.3 million in 1999/00 to 14.2 million in 2011/12, with annual contributions rising from £28.4 billion to £53.5 billion. The total market value of investments held in ISAs was £391 billion in 2012, accounting for 14% of net financial assets held by the household sector in 2011 (the most recent data available at the time of writing).[3] Nevertheless, a question remains over how far the popularity of ISAs may have been limited by behavioural distortions that recent studies suggest have an important influence on the savings decisions that people make. The

[1] A number of alternative tax advantaged savings schemes operate in parallel to ISAs. The most prominent of these is the National Savings and Investments (NS&I), which was first established in 1861 to provide the government with an alternative channel of debt financing. ISAs are the largest such scheme in operation in the UK at the time of writing, with total funds under management exceeding those in the NS&I by a factor of four.

[2] Tax-Free Savings Accounts (TFSAs) administered in Canada are similar to ISAs, but with the distinction that TFSAs allow unused contribution limits from previous years to be carried forward without limit. In the US, the closest approximation to an ISA is a Roth Individual Retirement Arrangement (Roth IRA). The key differences between an ISA and a Roth IRA is that there are no penalties to early withdrawal imposed on ISA balances, and there are no upper income bounds on those who are eligible to contribute to ISAs.

[3] Based on a total market value of ISAs and PEPs in 2011 of £374 billion (*Individual Savings Account Statistics*, September 2012, Office for National Statistics, Table 9.6), and £2742.4 billion of net financial assets (National Accounts variable NZEA).

current study sheds some light on this issue.

There is extensive evidence on both experimental and field data that default options matter in a wide range of decision making contexts.[4] One example from this literature that is relevant to the current study concerns decisions over private savings for retirement. Madrian, Choi, Laibson and co-authors have shown that a broad range of decisions regarding 401(k) plans in the US – including plan participation, rates of contributions, and investment fund allocations – are strongly influenced by the respective default options that individuals face (e.g. Madrian and Shae, 2001, Choi *et al.*, 2002, Beshears *et al.*, 2008). These observations have been confirmed on data for other countries (see McKay, 2006, for the UK), and are now being taken into account in the design of pensions policy (e.g the US Pension Protection Act, 2006, the New Zealand Kiwisaver Act, 2006, and the UK Pensions Act, 2008). The observed influence of framing effects also calls into question how far behavioural models that omit such effects provide an appropriate basis for understanding responses to tax-advantaged savings schemes.

The potential importance of decision making rigidities has long been recognised, and has motivated calls for reform of the classical rational agent model (e.g. Simon, 1955). In some contexts, the observed importance of framing effects has focussed attention on the behavioural influences of transactions costs.[5] The trivial size of observable transaction costs that are commonly incurred when switching between alternative savings vehicles, however, suggests that something more must be at work in this context. O'Donoghue and Rabin, for example, have shown that it is not enough to assume that consumers are myopic to capture the framing effects observed for 401(k) plans given the financial costs that are involved – it is also necessary to assume that consumers are naïvely unaware of the time-inconsistency of their preferences.[6] In context of such a preference structure, procrastination can mean that policies which influence the default options that people face are more effective than those which are designed to alter incentives at the behavioural margin.

This last observation is recognised in a series of papers by Choi, Madrian, Laibson, and Metrick, which consider policy alternatives over default options for pension saving.[7] Choi *et al.* consider a stylised form of the model of procrastination explored by O'Donoghue and Rabin. Importantly, Choi et al. (2003) note that they assume "naïve beliefs because they increase the force of procrastination, but our qualitative results would not change if we instead assumed that agents hold rational expectations"

[4] A default option is the selection from the available alternatives that requires the minimum effort to enact on the part of the decision maker. As such, default options are often defined by the way in which decision alternatives are framed. Framing effects have, for example, been identified for decisions regarding insurance contracts (Johnson *et al.*, 1993), fitness club contracts (DellaVigna and Malmendier, 2004), email marketing (Johnson *et al.* 2002), and organ donation (Johnson and Glodstein, 2003).

[5] This tack has been particularly prominent in literature concerning stock accumulation at the firm level, and the demand for factors of production; see Khan & Thomas (2008) for a review. See also Flavin & Nakagawa (2008) on transactions costs in relation to decisions regarding owner-occupied housing.

[6] See O'Donoghue & Rabin (1999*b*), O'Donoghue & Rabin (1999*a*), and O'Donoghue & Rabin (2001), building on Akerlof (1991).

[7] See, e.g., Choi et al. (2003), and Carroll et al. (2009).

(p. 181). It seems reasonable to suppose that the same may be true for the distinction between the assumptions of myopic and time-consistent preferences: allowing for decision costs of sufficient scale might permit observed framing effects over savings schemes to be captured without the need to relax the classical assumptions of rational expectations *or* time consistent preferences. This alternative has not been explored in the literature that is referred to above, and forms the basis for the current study.

We consider whether a provision for framing effects helps to improve the ability of a rational agent model to reflect ISA participation rates described by contemporary survey data. The study follows O'Donoghue and Rabin in ascribing behavioural distortions associated with framing effects to "salience costs" that conceptually represent both explicit transactions costs and the opportunity cost of the effort required to make a decision. The salience cost is modelled as a reduction of utility, so that the approach reflects aspects of the literature on both adjustment costs (referred to above) and habit formation. In contrast to O'Donoghue and Rabin, however, the current study assumes a time-consistent preference relation and the analysis conforms to the assumption of perfect rationality.

The analysis reported here suggests that salience costs are important in capturing public engagement with the ISA scheme: in the absence of such costs, the rational agent model assumed for analysis suggests that participation rates would exceed those observed by a wide margin. The parameterised model is used to explore how framing effects influence the effectiveness of the ISA scheme in motivating increased household savings, building on the extensive literature that explores such effects on classical life-cycle assumptions (e.g. Gustman & Steinmeier (1986), Rust & Phelan (1997), French (2005), Sefton et al. (2008)).

The paper is organised as follows. Section 2 describes Individual Savings Accounts and the statistical evidence that motivates the study. A structural model of ISA savings decisions is described in Section 3. Section 4 reports calibration of the model to survey data, focussing upon the unobserved preference parameters that are necessary to match to rates of ISA take-up. The practical relevance of the decision costs that are identified in Section 4 are explored in Section 5, focussing upon counterfactual implications drawn from the structural model for welfare and saving. A summary and directions for further research are provided in the conclusion.

2 The UK System of ISAs

Individuals aged 16 or over and resident in the UK can invest in a wide range of alternative assets through an ISA.[8] Contributions to an ISA must be made from post-tax income, and are subject to strict upper bounds per year. Two principal types of ISA exist: cash ISAs, and stocks and shares ISAs.

[8] Junior ISAs were introduced in November 2011 for children under age 16, but these must be established and contributed to by a person with parental responsibility.

The defining difference is that stocks and shares ISAs require there to be a credible risk that at least five percent of an investment may be lost, whereas a cash ISA is primarily designed to cover money deposits.[9] Any investment returns realised on assets held within an ISA are tax exempt, and there are no bounds imposed on aggregate account value. Funds held within an ISA can be withdrawn at any time without penalty, and cannot be used as collateral for a loan.

Annual contributions to an ISA are subject to three upper thresholds; one for each of 'cash' and 'stocks and shares' ISAs, and one on the aggregate ISA contribution. When the ISA scheme was introduced in 1999, annual contributions to cash ISAs were limited to £3000, for stocks and shares ISAs to £4000, and aggregate contributions were limited to £7000. These limits were subject to somewhat haphazard alterations until the 2010 budget, when the contributions limits were set to £5100 for cash ISAs, £10200 for stocks and shares ISAs, and £10200 in aggregate. These limits have subsequently been increased in line with inflation, rounded to the nearest £120.[10]

Interest income earned on cash deposits and bonds held within an ISA is tax-free, as are capital gains on all transactions.[11] Dividend income on equity held within an ISA is also not assessable for individual income taxes, which delivers a benefit to higher rate tax payers. In contrast, the opportunity to claim a rebate in respect of company taxes paid was discontinued in 2004, so that basic rate tax payers (and below) now derive no advantage from receiving dividend income through an ISA. This has tended to skew the tax incentives of the ISA scheme in favour of equity investments that deliver capital growth. Competition for cash ISAs has typically seen these pay appreciably higher pre-tax rates of interest than other comparable accounts[12], which exaggerates differences in post-tax returns. Furthermore, trading company stocks within an ISA wrapper can reduce administrative burden, as there is no need to report associated activity to Revenue and Customs.[13]

In common with much of UK policy governing private sector savings, ISAs were introduced following a lengthy reform process. Importantly, this process involved the introduction of two separate schemes, which were superseded by ISAs. Personal Equity Plans (PEPs) were introduced in the 1986 budget to encourage people to purchase (UK) equities. This scheme was very similar to, and eventually rolled into, the stocks and shares ISA scheme. Tax-exempt Special Savings Accounts (TESSAs) were introduced in the 1990 budget, essentially to address concerns that the focus of PEPs on equities delivered little

[9] A stocks and shares ISA can hold cash deposits, which are denoted "awaiting investment" and subject to a 20% tax charge on interest income. Cash ISAs in principle can hold investments other than money deposits, if these do not qualify for the five percent rule.

[10] The inflation index used for this purpose will change from the RPI to the CPI in 2013/14.

[11] Capital losses within an ISA cannot be off-set for tax purposes against other gains.

[12] For example, variable interest ISA accounts are reported by the Bank of England to have attracted average annual interest rate of 2.4% p.a. from Jan 2011 to Feb 2013, relative to 1.4% p.a. for bank accounts over the same period (codes IUMB6VJ, IUMB6VL).

[13] Individuals must report details of all trades in equities whenever the total sales value within a given tax year exceeds a threshold defined by the Capital Gains Tax allowance.

advantage to people with a pronounced aversion to risk, particularly those on modest incomes. The TESSA scheme allowed an individual to invest a cash deposit in return for a tax free investment return. Unlike PEPs, however, the TESSA scheme imposed limits on contributions and account withdrawals that departed significantly from the ISA model. Specifically, consistent with its focus on people with modest financial circumstances, TESSAs imposed an upper limit on aggregate contributions of £9000, and withdrawal of capital within five years of the initial contribution incurred a claw-back penalty on interest income earned.

The policy uncertainty that accompanied introduction of the ISA scheme was recognised in the Acts that brought the scheme into being. Following industry petitioning, the government committed to a minimum life-span for ISAs of 10 years. This commitment was, however, something of a double-edged sword in the sense that it was perceived to be a *de facto* 'end date' for the scheme. ISAs were only considered to continue indefinitely following a government review in 2006.

That ISAs managed to attract 9.3 million subscriptions worth £28.4 billion in the scheme's first year of operation, despite the regulatory uncertainty noted above, is attributable in large part to the basic similarity between stocks and shares ISAs and PEPs. Nevertheless, growth of the scheme to the time of writing has been impressive. Aggregate assets held in ISAs were reported to be worth £391 billion in 2012, with the associated tax exemptions estimated to have cost the exchequer £1.75 billion in 2012/13. In this paper we are concerned with how far decision making rigidities may have dampened the influence of ISAs on household savings rates, and therefore the return that the government has achieved on its foregone tax revenue.

2.1 ISAs in the household portfolio allocation

Participation rates in ISAs of the adult population, distinguished by age and year, are reported in Figure 1. The age profile of ISA participation described in Panel A takes an intuitive shape. The proportion of the adult population reported to have an ISA account rises steeply from 10% at the beginning of adult life to around 30% at age 30, flattens out during peak child-rearing years, before continuing to rise between ages 45 and state pension age, peaking at around 50% of the population at age 65. The respective year specific series indicate a reasonable degree of stochastic variation between age groups, which complicates comparisons between them. Averaging over all ages, however, provides a clear indication of the steady rise in the proportion of the population reported to hold an ISA account, as reported in Panel B. Here we see that the incidence of ISA accounts rose by 6% over the entire adult population during the 8 year period covered by the available data, consistent with the expansion of the scheme discussed above.

Additional detail regarding the place of ISAs in the household portfolio allocation can be gained by

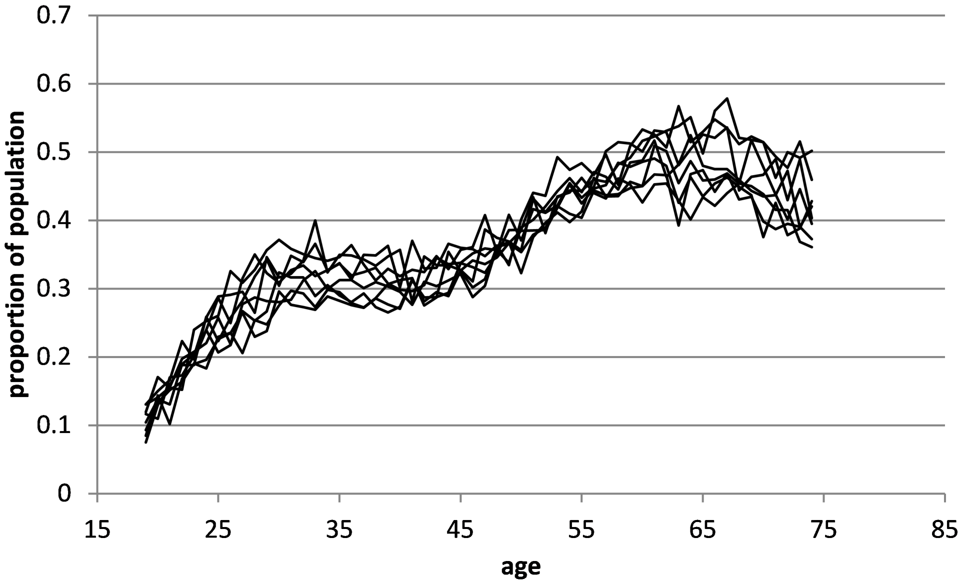

Panel A: participation rates by age and sample year

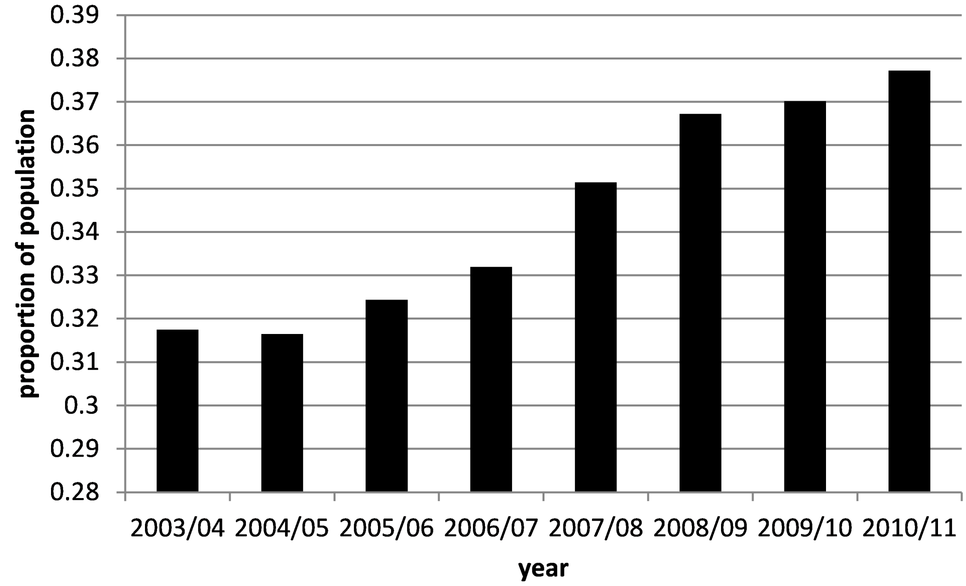

Panel B: average participation rates of full survey population

Source: Author's calculations on data from repeated waves of the Family Resources Survey: 2003/04 to 2009/10

Notes: Participation defined as membership of a benefit unit with aggregate ISA assets exceeding £0.10

Figure 1: Participation rates in ISAs by age and year

considering wealth data reported by the Wealth and Assets Survey. Associated statistics are displayed in Table 1. This table indicates that the vast majority of UK household wealth is held in either pensions or owner-occupied housing, which together account for over 80% of average net worth throughout the age distribution. Of the wealth that remains, investments held in ISAs account for 15% of the aggregate averaged over all households, with slightly higher weighting early and late in life. Just one in every five individuals between ages 20 and 29 are reported to hold any wealth in ISAs, rising to half the population between ages 60 and 69, before falling away slightly at higher ages. Averaging over the full population population aged 20 and over, more then three out of every five individuals report holding no assets in an ISA.

Of the six asset classes singled out in Table 1, the closest approximation to ISAs is given by 'other financial assets', which include cash, bonds and equities held outside of an ISA. It is notable that, within this limited subset of household assets, investments held is ISAs account for just 30% of value on average, falling with age to 25% for the population aged 70 and over.

Statistics that focus on the allocation of non-housing and non-pension wealth are reported in Table 2, distinguished by wealth quintile. These statistics indicate that ISAs were reported as a rarity amongst households in the bottom two wealth quintiles, but even amongst households in the top quintile, one out of every four reports no ISA assets. Focussing in on holdings of bonds, shares and cash – assets that could all presumably be held within an ISA – we see that the low rates of ISAs toward the bottom of the wealth distribution make sense in context of their wider portfolio allocations. The same is not true, however, for households in the top two wealth quintiles, who report substantial holdings of cash, bonds and stocks outside ISAs, indeed larger holdings than those held inside ISAs. Although ISAs have enjoyed a measure of success in attracting private sector savings, there therefore appears to be considerable scope for households to make more of the tax advantages that are offered through the ISA scheme. The remainder of the study considers the potential role of 'salience costs' in explaining observed decisions concerning ISA investments, and the influence of such costs on household savings rates.

3 A Structural Model of ISA Investments

Survey data regarding financial asset allocations suggest that decisions regarding ISAs are affected by considerations that extend beyond the simple risk/return trade-offs with which most portfolio theory is concerned. In this study, such 'considerations' are represented by a salience cost in an otherwise standard economic model of life-cycle savings.

ISAs are designed to complement private sector pensions in the UK. At a minimum, a behavioural model of ISA investments must consequently allow for allocations between three alternative asset classes:

Table 1: Household net asset allocations by age band (£2006)

	20-29	30-39	40-49	50-59	60-69	70+	all
ISAs	832	2282	4462	8735	13780	11879	6099
population with no ISA (%)	78.2	68.4	64.5	55.2	49.7	53.5	63.8
owner occupied housing	9262	55232	97014	141120	159534	151240	93287
other property	2196	8066	13658	17245	17527	10962	10459
private pensions	7651	51291	117474	223019	216549	134147	108995
other financial assets	418	8907	16386	33056	40453	37058	20242
total	20437	125846	249058	423231	447892	345341	239146

Source: Author's calculations on WAS data reported for the period between July 2006 to June 2007

Weighted averages reported in current prices

Sample omits households with own business assets

Table 2: Average net non-pension and non-housing assets, by wealth quintile (£2006)

	lowest	2nd	3rd	4th	highest	all
ISAs	68	27	647	4930	24838	6099
population with no ISA (%)	95.0	92.1	68.4	38.9	24.4	63.8
bonds	1	0	30	816	13561	2879
shares	10	6	159	1391	29961	6301
national savings products	4	7	86	628	5316	1208
insurance products	26	4	103	1454	9836	2283
cash	-3023	-1937	-846	4223	38218	7312
other property	-595	1	24	583	52351	10459
informal savings and loans	6	24	110	220	933	258
total	-3407	-1775	380	14285	175038	36904

Source: Author's calculations on WAS data reported for the period between July 2006 to June 2007

Quintile groups calculated on the aggregate value of non-housing and pension wealth, gross of financial loans

Weighted averages reported in current prices

Sample omits households with own business assets

ISAs, private sector pensions, and other liquid wealth. A key distinguishing feature between ISAs and pensions is that capital invested in an ISA can be accessed at any time, whereas it can only be accessed from age 55 in a pension. The influence of investment illiquidity on savings incentives depends crucially upon the uncertainty that is associated with evolving financial circumstances and consumption needs. We have therefore adopted the standard economic framework for exploring savings behaviour in the context of uncertainty, which assumes that decisions are made by well-informed agents to maximise expected lifetime utility.

Income sharing motivates our focus on the nuclear family (household) as the unit for analysis. The model follows a sample of 'reference adults' and their evolving families at annual intervals through time. It is assumed that, although the age of death of any reference adult is random and unknown ex ante, there is a maximum possible lifespan, which is taken to be 130 years. Households choose their consumption, labour supply, net ISA contribution, and pension participation to maximise expected lifetime utility, given their circumstances, preferences, and beliefs regarding the future. A household's circumstances are described by the age and year of birth of its reference person, their relationship status (single or couple), the wage potential of adult members, pension rights, ISA value, other net liquid wealth, and time of death. Preferences – that are the same for all households – are defined by a time separable iso-elastic utility function, with intra-temporal preferences over consumption and leisure taking a CES form. Beliefs are rational in that they are consistent with the processes that generate the intertemporal evolution of household circumstances.

In the interests of parsimony, the study does not include analysis of uncertain rates of return. Of the eight characteristics that define the circumstances of a household, three are stochastic (relationship status, wage rates, and time of death), and five deterministic (year of birth, age, pension rights, ISA value, and other net liquid wealth). Analysis proceeds by numerically solving the decision problem for any conceivable combination of household characteristics. These behavioural solutions are used to generate data for a simulated sample population, which are used as the basis for analysis.

This section describes key features of the assumed model. Details of the numerical procedures used to solve the utility maximisation problem and simulate households through time are provided in Appendix A.

3.1 Preference relation

A wide range of models have been proposed to take into account behavioural rigidities of the sort considered by this study.[14] The model that we assume retains the core elements of the traditional economic approach to inter-temporal decision making, augmented to include a 'salience cost' following

[14] See Gigerenzer & Selten (2002) for examples of some of the more radical departures from classical economic theory that have been advocated.

O'Donoghue and Rabin:

$$U_{i,a} = \frac{1}{1-\gamma} E_a \sum_{j=a}^{A} \delta^{j-a} \left\{ \phi^b_{j-a,a} \left[u\left(\frac{c_{i,j}}{\theta_{i,j}}, l_{i,j} \right)^{1-\gamma} + \Psi\left(d^{ISA}_{i,j}\right)^{1-\gamma} \right] + \left(1 - \phi^b_{j-a,a}\right)\left(\zeta_0 + \zeta_1 w^+_{i,j}\right)^{1-\gamma} \right\} \tag{1}$$

where $\gamma > 0$ is the (constant) coefficient of relative risk aversion; E_a is the expectations operator; A is the maximum potential age; δ is an exponential discount factor; $\phi^b_{j-a,a}$ is the probability of reference adult from birth year b surviving to age j given survival to age a; $u(.)$ is intra-temporal utility derived from equivalised consumption and leisure; $c_{i,a} \in R^+$ is discretionary composite (non-durable) consumption; $\theta_{i,a} \in R^+$ is adult equivalent size based on the "revised" or "modified" OECD scale; $l_{i,a} \in [0,1]$ is the proportion of family time spent in leisure; $\Psi(.)$ is the (intra-temporal) salience cost; $d^{ISA}_{i,a}$ is an indicator variable that equals 1 if reference adult i – with no pre-existing ISA wealth – chooses to make an ISA contribution at age a, and is zero otherwise (defined formally below); the parameters $\zeta_{0/1}$ reflect the "warm-glow" model of bequests[15]; and $w^+_{i,a} \in R^+$ is net non-pension wealth when this is positive and zero otherwise.

Intra-temporal utility is based upon a Constant Elasticity of Substitution function, augmented to include a salience cost:

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

$$\Psi\left(d^{ISA}_{i,a}\right) = \left(\psi_0 - \psi_1 d^{ISA}_{i,a}\right) \tag{3}$$

where $\varepsilon > 0$ is the (period specific) elasticity of substitution between equivalised consumption $(c_{i,a}/\theta_{i,a})$ and leisure $(l_{i,a})$; and $\alpha > 0$ is the utility price of leisure. Setting $\psi_0 > \psi_1 \geq 0$ imposes a (discontinuous) salience cost on the first contribution made to an ISA, and the additive separable form ensures that the salience cost does not influence the marginal rate of substitution between contemporaneous consumption and leisure.

Little is currently understood about the form or scale of salience costs, which might otherwise guide the specification of equation (2). In this study we focus on salience costs that depend on the extensive margin of ISA participation. This raises two key issues: (i) why constrain salience costs to focus exclusively on ISAs; and (ii) why assume the form represented in equation (3) for such costs.

On the former of these issues, we experimented with alternative specifications that extended salience costs to include participation in private pension schemes. This was omitted from the final specification, as we found extending salience costs to private pensions did not help to improve the match of the model to observed rates of pension participation described by survey data. Discussion of this issue is taken up in Section 4.1.

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

Regarding the functional form adopted for salience costs in relation to ISAs, the focus on the extensive margin of participation reflects the limitations of the publicly available micro-data, and the form for salience costs that we find most intuitively plausible. While data concerning participation in ISAs are readily available, none of the micro-data sources considered for analysis provide detail concerning the intensive margin of net ISA investment contributions. Furthermore, ISA providers typically charged few direct costs for investing through the scheme, and trading within an ISA typically involved a comparable or smaller administrative overhead to trading more generally.[16] There were, nevertheless, difficulties for prospective investors in familiarising themselves with the rules of the scheme, choosing between alternative ISA providers, and addressing the administrative overhead associated with account applications processes. The salience cost described in equation (2) is designed to reflect these difficulties, and we have adopted a form that allows the scale of costs to be altered while avoiding direct feedback on preferences regarding (contemporaneous) consumption and employment.

3.2 Individual Savings Accounts (ISAs)

As uncertain asset returns are not taken into account, the model does not distinguish between cash and stocks and shares ISAs, but rather focusses on ISA investments in aggregate. Otherwise, ISAs are represented in a way that captures the key features of the scheme. Each household can contribute to, or withdraw from, a household specific ISA account during any stage of the reference adult's life. Annual net contributions to each household's ISA account are made out of post-tax income, and cannot exceed an upper threshold per adult member. At the start of each period, all wealth held in an ISA is assumed to accrue the same rate of return, r^{ISA}, which is certain. Investment returns to ISA wealth are not taxed, no formal upper bound is imposed on aggregate wealth held in an ISA, and ISA wealth can be drawn down at any time without penalty.

In most periods, wealth held in an ISA, w^{ISA}, is assumed to vary inter-temporally as described by the equation:

$$w_{i,a}^{ISA} = r^{ISA} w_{i,a-1}^{ISA} + k_{i,a}^{ISA} \tag{4}$$

where $k_{i,a}^{ISA} \geq -r^{ISA} w_{i,a-1}^{ISA}$ denotes net contributions into the scheme (negative when there are net out-flows). The only departure from equation (4) is when the relationship status of a reference adult is identified as changing, in which case a fixed factor adjustment is used to alter ISA wealth to reflect the proportion of the account accruing to the reference adult's partner. Cohabitating partners are treated symmetrically throughout the analysis, so that marriage doubles ISA wealth, and marital dissolution halves it.

[16] Undisclosed charges for cash ISAs, for example, were implicitly administered through the interest rates on offer. In this regard it is notable that cash ISAs typically returned better rates of interest than comparable deposit accounts, suggesting that any such charges were more than off-set by competition amongst ISA providers.

We can now define the indicator variable $d_{i,a}^{ISA}$ that appears in equation (1) as:

$$d_{i,a}^{ISA} = \begin{cases} 1 & \text{if } w_{i,a-1}^{ISA} = 0 \text{ and } k_{i,a}^{ISA} > 0 \\ 0 & \text{otherwise} \end{cases} \tag{5}$$

3.3 Private pensions

The approach taken to model private pensions shares close similarities with that adopted for ISAs. Pensions are modelled at the household level, and are notionally defined contribution in the sense that the value of accumulated pension pots are a function of pension contributions and investment returns. Note that, in the absence of uncertainty over pension fund investment returns – as is assumed here – it is possible to express a defined contribution pension in terms of a (weighted) career average defined benefit pension.

In each year, a household with labour income exceeding a minimum threshold, g_l^P, can choose whether to make fresh contributions to its pension scheme. If a household chooses to contribute to its pension, then a fixed share of its total pre-tax labour income, π^P, is added to its accumulated pension fund. Contributing households also receive an employer contribution to their pension fund, which is specified as a fixed share of pre-tax labour income, π_{ec}^P. Eligible employer contributions to a household's pension are lost if the household chooses not to contribute to its scheme in the respective year. Wealth held in a private pension, $w_{i,a}^P$, is assumed to attract a fixed rate of return r^P, and in most periods evolves following the accounting identity:

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

where $d_{i,a}^P$ is an indicator variable, equal to one if household i at age a contributes to its pension, and zero otherwise. As with ISAs, the only departures from equation (6) are in respect of relationship formation (in which case it doubles), and relationship dissolution (in which case it halves).

A household can choose to access its accumulated pension fund at any time between ages 55 and 70, at which time 25% of the fund is taken as a tax-free lump-sum and the remainder used to purchase an inflation adjusted life annuity. The annuity income generated by a pension is assessable for income taxes, and is taken into account in the evaluation of means-tested welfare benefits. The annuity rates assumed for analysis are calculated with reference to the survival rates assumed for individual birth cohorts, an assumed return to capital, and an assumed transaction cost levied at the time of purchase. This specification of private pension opportunities broadly reflects the terms of occupational pension schemes administered in the UK.

3.4 Other net assets

The third and final investment asset that is included for analysis is denoted other net (liquid) assets, $w^o_{i,a}$. The model allows for the possibility that other net assets may be negative, representing unsecured debt. The level of unsecured debt permitted in the model is constrained by the requirement that all debts be repaid by age A', the minimum potential income stream, and the nature of the preference relation. Intertemporal variation of $w^o_{i,a}$ is, in most periods, described by the accounting identity:

$$w^o_{i,a} = w^o_{i,a-1} + k^o_{i,a-1} + \tau_{i,a-1} - c_{i,a-1} \tag{7}$$

$$k^o_{i,a-1} = -\left(k^{ISA}_{i,a-1} + \pi^P g_{i,a-1} d^P_{i,a-1}\right) \tag{8}$$

where k^o represents net investment flows from other asset classes, and τ denotes disposable income net of non-discretionary expenditure. The only potential departures from equation (7) occur when a reference adult is identified as getting married or incurring a marital dissolution, in which case wealth is doubled or halved to reflect the 'other wealth' attributed to the spouse.

Other liquid wealth, w^o, represents all marketable wealth other than private pensions and ISAs. As discussed in Section 2.1, much of this wealth is held in owner occupied housing. Housing is a peculiar asset class, in the sense that it delivers accommodation services in addition to standard investment returns. The current study ignores the behavioural influence of accommodation services associated with housing, due to the computational overhead involved, recognising that the salience costs for ISAs identified on this basis are likely to represent an upper-bound to those that apply in practice.

The tax function of the model takes a flexible form, represented by:

$$\tau_{i,a} = \tau(l_{i,a}, x_{i,a}, n_{i,a}, n^c_{i,a}, r^o_{i,a} w^o_{i,a}, b, a) \tag{9}$$

This function, which is designed to provide a realistic reflection of UK transfer policy, depends on labour supply, $l_{i,a}$; (taxable) private non-capital income, $x_{i,a}$; the number and age of adults, $n_{i,a}$, a; the number of dependent children, $n^c_{i,a}$; the return to other assets, $r^o w^o_{i,a}$ (which is negative when $w^o_{i,a} < 0$); and birth year, b. Private non-capital income $x_{i,a}$ is equal to labour income, $g_{i,a}$, less private pension contributions, $\pi^P g_{i,a-1} d^P_{i,a-1}$, plus pension annuity income. Note that investment returns on wealth held in both ISAs and private pensions do not influence τ, and nor do ISA contributions (though contributions to private pensions do).

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

$$r^o = \begin{cases} r^I & \text{if } w^o \geq 0 \\ r_l^D + \left(r_u^D - r_l^D\right) \min\left\{\frac{-w^s}{g^{ft}}, 1\right\}, r_l^D < r_u^D & \text{if } w^o < 0 \end{cases} \tag{10}$$

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

3.5 Labour income dynamics

The dynamic programming problem is simplified by modelling wages at the household level. Household labour income, $g_{i,a}$, depends upon the latent wage, $h_{i,a}$, the wage offer received, $o_{i,a}$, the household's labour supply decision, $l_{i,a}$, and whether the household has started to draw on its private pension wealth, $ret_{i,a}$:

$$\begin{aligned} g_{i,a} &= \lambda_{i,a} h_{i,a} \\ \lambda_{i,a} &= \lambda^{emp}\left(l_{i,a}\right) \lambda^{wo}\left(o_{i,a}\right) \lambda^{r}\left(ret_{i,a}, a\right) \end{aligned} \tag{11}$$

Latent wages

Latent wages are predominantly assumed to follow the stochastic process described by:

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

$$\omega_{i,a-1} \sim N\left(0, \sigma^2_{\omega, n_{i,a-1}}\right) \tag{12b}$$

where the parameters $m(.)$ account for wage growth, $\beta(.)$ accounts for time persistence in earnings, $\kappa(.)$ is the return to another year of experience, and $\omega_{i,a}$ is a family specific disturbance term. All model parameters vary by relationship status ($n_{i,a}$), and the wage growth parameters also depend upon the age and birth year of the reference adult. The only exceptions to this specification are when a reference adult enters or departs a cohabitating relationship, in which case we assume equal latent wages between spouses (so that h doubles following marriage formation and halves following marital dissolution). This parsimonious wage specification has been explored at length in the literature (e.g. Creedy (1985)), and requires the addition of just two state variables to the decision problem (h, ω).

Labour supply decisions

Each adult household member has three discrete labour supply alternatives, representing full-time, part-time, and non-employment. Although our central interest is in savings behaviour, we also include labour supply to reflect the endogeneity that exists between these two decision margins. Increased labour supply results in a lower fraction of time enjoyed as leisure, and a higher fraction of the family's latent wage received as labour income (assuming that a wage offer is received). In the case of couples, one adult is assumed to be the 'chief wage earner', and the other the 'second wage earner'. Full-time employment of the chief wage earner is assumed to return a larger fraction of the households's latent wage than that of the second wage earner (if one exists). The effects of alternative employment decisions on leisure and earnings are specified as fractions of full-time employment, π^l for leisure and π^g for earnings.

Wage offers

Wage offers are included in the model to allow for the possibility of (involuntary) unemployment. Following experimentation with alternative assumptions, the model considered here focuses on wage offers for the chief wage earner (defined above). Wage offers are modelled as uncertain between one period and the next, subject to age and relationship specific probabilities $p^o\left(n_{i,a}, a\right)$. If a wage offer is received by the chief wage earner, $o_{i,a} = 1$, then the household's labour income is an increasing function of their labour supply. If a wage offer is not received by the chief wage earner, $o_{i,a} = 0$, then the household's labour income depends only on the labour supply decision of the second wage earner (if one exists).

Accessing pension wealth

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

3.6 Household demographics

Three key demographic characteristics are taken into account by the model: relationship status, number of dependent children, and survival of the reference adult.

Modelling relationship status

The relationship status of each reference adult is modelled explicitly, and is considered to be uncertain from one year to the next. The transition probabilities that govern relationship transitions depend upon a reference adult's existing relationship status, their age, and birth year. These probabilities are stored in a series of 'transition matrices', each cell of which refers to a discrete relationship/age/birth year combination.

Modelling children

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

Modelling survival

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

4 Model Calibration and the Scale of Decision Costs

The current study follows the now-standard two-stage approach to parameterising models of the type described in Section 3. In the first stage, a subset of parameters were evaluated exogenously from the model structure using standard statistical techniques. Given the model parameters identified in the first stage, remaining model parameters were adjusted to match selected 'simulated moments' implied by the structural model to 'sample moments' estimated on survey data. This section focusses upon parameter values identified endogenously to the model structure, with which the study is principally concerned; parameters identified in the first-stage of the parameterisation are reported in Appendix B.

We begin with an outline of the identification strategy employed. Data considered for analysis are then described, before reporting the results obtained.

4.1 Identification strategy

The assumed preference relation includes eight parameters (see Section 3.1): relative risk aversion, γ; an exponential discount factor, δ; two parameters for the warm-glow model of bequests, $\zeta_{0/1}$; the intra-temporal elasticity, ε; the utility price of leisure, α; and two parameters for salience costs of ISA investments, $\psi_{0/1}$. As these parameters are unobservable, they were identified in the second stage of the model parameterisation. Experience effects and wage responses to pension take-up have an important bearing on employment incentives, and make wage potential endogenous to the structural model. Parameters governing wage growth, $m(.)$, experience $\kappa(.)$, factor effects of pension take-up λ^{r}, and earnings volatility $\sigma^2_{\omega,n_{i,a}}$ were consequently also identified in the second stage of the parameterisation.[17]

[17] The persistence of latent wages β and factor effects of alternative labour supply decisions, λ^{emp}, were identified exogenous of the model structure; these parameters are reported in Appendix B.

The full structural model takes just over 7 days to evaluate behavioural responses for a single parameter combination on a modern workstation (dual Intel Xeon processors with 24GB of RAM). This run-time makes econometric estimation of the parameters referred to above impractical, and so the second stage of the parameterisation was implemented by manual calibration. Following extensive experimentation, we settled upon the following step-wise procedure to identify associated model parameters.

We started with a variant of the model that omits explicit consideration of ISAs. Suppressing ISAs is advantageous because the resulting model solves in 8 minutes, rather than 7 days, which permits greater flexibility in testing over alternative parameter combinations.[18] The parameters of this limited model were then divided into two mutually exclusive and exhaustive sets: set A comprising the parameters governing wage growth and earnings volatility, and set B comprising all other calibrated parameters excluding $\psi_{0/1}$. We began by setting all wage growth parameters $m(.) = 1$, and made initial guesses for earnings volatility, $\sigma^2_{\omega,n_{i,a}}$. Given these assumptions for set A parameters, and the model parameters identified exogenously from the model structure in the first stage of the analysis (reported in Appendix B), we adjusted the parameters in set B to reflect observed household characteristics endogenous to the structural model. Having obtained first approximations for set B parameters, we then adjusted the parameters $m(.)$ and $\sigma^2_{\omega,n_{i,a}}$ to reflect historical earnings data. We then extended the model to allow for ISA investments, and adjusted the salience cost parameters $\psi_{0/1}$ to reflect rates of participation in ISAs described by survey data, given the calibrated parameter values identified in the preceding two steps. This procedure was repeated until convergence between the three sets of parameters A, B and $\psi_{0/1}$, was obtained.

Identification of wage growth and earnings volatility parameters; 'set A'

The drift parameters, $m(.)$, and the dispersion parameters, $\sigma^2_{\omega,n_{i,a-1}}$, were calibrated against historical data by projecting a reference population cross-section backward through time. We begin by describing adjustment of the drift parameters.

The drift parameters were adjusted to reflect age and year specific geometric means of employment income calculated separately for singles and couples from survey data. The model includes a separate drift parameter for each age, year, and relationship combination, so that a close match could be obtained to the associated sample moments. Given the large number of model parameters involved, this stage of the parameterisation was undertaken using an automated procedure. First, age, year, and relationship specific means of log employment income implied by the model were calculated from the simulated

[18] The additional computational burden associated with including ISAs may be surprising, particularly for non-specialists. It is attributable to three key factors: the increase in the domain of the state-space (increasing the computation burden by a factor of 22), the need to use cubic rather than linear interpolation methods (factor of 8), and the need to solve over two continuous decision variables (factor of 9).

panel data projected back in time for the reference population cross-section. These simulated moments were subtracted from associated sample moments estimated from survey data. Note here that the endogeneity of employment decisions in the structural model ensures that selection effects are properly taken into account. The differences between the simulated and sample moments were then multiplied by a 'dampening factor' of 0.4, and the exponent of the result was multiplied by the associated drift parameter to obtain a revised value for the parameter. Application of this procedure to repeated simulations was continued until the average absolute variation of drift parameters within any year fell below 5 percentage points.

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

Identification of preference parameters and experience effects; 'set B' and $\psi_{0/1}$

Unlike existing models of its type, the current structural specification is adapted to permit parameterisation of unobserved preference parameters on data observed for a single population cross-section. See Lucchino & van de Ven (2013) for extended discussion of this issue.

We started by adjusting preference parameters to match the model to age-specific moments of labour supply observed for a single population cross-section. Increasing the utility price of leisure α tends to decrease labour supply throughout the working lifetime. Increasing the effect of experience on earnings (κ in equation 12) increases the price of leisure early in the working lifetime, relative to later in life. Reducing the factor effect of pension take-up, λ^{ret}, tends to decrease labour supply late in the working lifetime. These three model parameters consequently provide a high degree of control over the employment profile throughout the life-course.

Having obtained an approximate match to moments of employment, we next considered household income. The income potential of each reference adult is exogenously given by the cross-sectional data upon which the specification of preference parameters was based. The distribution of observed wages in the reference cross-section is consequently a product of labour supply decisions. The distribution of employment implied by our structural model is influenced by the intra-temporal elasticity: increasing ε tends to exacerbate the utility cost of leisure for high wage potential households, and vice versa for low wage potential households. This parameter was consequently adjusted to match simulated to sample moments of household income.

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

As suggested in Section 3.1, the model was originally specified to include salience costs in relation to pension scheme participation, in the expectation that these would help to obtain an improved match to observed rates of pension take-up. The analysis undertaken here, however, revealed that the remaining model parameters could be adjusted to obtain a close match to sample moments of consumption and pension scheme participation at the same time, without the need to resort to salience costs over private pensions. The key in this regard was the parameter of relative risk aversion γ. Put another way, the addition of pension scheme participation to the margins considered for analysis is required to identify γ. Whereas raising δ tends to imply lower consumption and higher pension scheme participation, *ceteris paribus*, raising γ tends to imply lower consumption *and* lower pension scheme participation, by exaggerating the precautionary motives to save. We consequently started with a relatively high value of γ (=5), and adjusted δ and $\zeta_{0/1}$ to match the age profile of consumption. If the associated rates of pension scheme participation were then too low, we reduced γ and readjusted δ and $\zeta_{0/1}$. Repeating this process enabled a close match to be obtained to consumption and rates of pension scheme participation at the same time.

Finally, the salience cost parameters, $\psi_{0/1}$, were adjusted to match the model to the net increase in the proportion of households aged 20 to 60 identified as participating in ISAs. The associated adjustments were trivial to implement, due to a simulated monotonic relationship between ISA participation and scale of salience costs.

The identification strategy described above obtains a tight fit for γ, which is of note given the importance of this parameter in determining behavioural responses to savings incentives, and the disparate estimates that are reported in the associated literature (discussed at length below). Nevertheless, limitations of the model assumed for analysis may have an important influence on our ability to identify the salience costs with which the study is concerned. Two key issues may be singled out in this respect. First, owner-occupied housing is not treated as a separate state variable in the model. As noted in Section 3.4, the salience costs for ISAs identified on this basis are likely to represent an upper bound to those that apply in practice. Secondly, the identification strategy outlined above does not permit a separate set of salience costs to be identified for decisions regarding private pensions.

The four sets of preference parameters governing intertemporal savings decisions – γ, δ, $\zeta_{0/1}$, and $\psi_{0/1}$ – are adjusted to match the model to four broad behavioural margins – consumption early in life, consumption late in life, pension scheme participation, and ISA participation. Including salience costs over private pension decisions would consequently require an additional behavioural margin for identification, which is complicated by limitations of available survey data and computational capacity. If decisions regarding private pension scheme participation are influenced by salience costs, then our failure to account for such costs will result in a form of omitted variable bias. The form that such biases take will depend crucially on the form of the omitted salience costs, complicating any predication of the associated influences on $\psi_{0/1}$.

4.2 Data

The model was matched to data structured around the population cross-section observed between July 2006 and June 2007. As discussed in Section 2, the ISA scheme was introduced following a lengthy reform process. Such a process adds a layer of uncertainty to individual investment decisions; each scheme has its own rules that must be understood and complied with, and there is no certainty regarding how long each policy incarnation will persist. Such uncertainty is likely to exaggerate behavioural headwinds to scheme success, particularly at the point of introduction. This motivated a preference for focussing on data observed for a recent population cross-section. Set against this preference, however, is a desire to avoid behavioural distortions attributable to the global financial crisis of 2007/08. Our sample period extends beyond the date that the government first committed to retaining ISAs permanently, but is prior to the deterioration of consumer sentiment leading up to the financial crisis.[19]

Two sets of data were required to identify the unobserved model parameters: micro-data describing individual circumstances for a reference population cross-section, and data that describe variation of endogenous sample moments. These two sets of data are described in turn.

Micro-data for a population cross-section

The structural model described in Section 3 projects individual specific characteristics through time, starting from a cross-sectional data-set describing the household circumstances of a population of 'reference adults'. The cross-sectional data that were used as a starting point for the model are drawn from wave 1 of the Wealth and Assets Survey. The Wealth and Assets Survey is a rich micro-data source that provides detailed information regarding demographics, income, assets, and liabilities for a sample of households selected from the Postcode Address File to reflect the population accommodated in private households in Great Britain, excluding Scotland north of the Caledonian Canal, the Scottish

[19] In the two years to July 2007, the GfK Consumer Confidence Index was approximately stable at -4 (negative being an indication of pessimism on average). The index fell between August 2007 and July 2008, from -4 to -39. The Treasury Economic Secretary announced the retention of ISAs in November 2006.

Islands, and the Isles of Scilly. The survey was designed to over-sample from high wealth households, and information was solicited from all individuals aged 16 or over in each responding household (excluding full-time students between 16 and 18 years of age). Data were collected continually between July 2006 to June 2008, and the survey achieved a response rate of 55 percent, providing information for 71,268 individuals in 30,595 households.[20]

The data reported by wave 1 of the WAS were screened to omit those responding after June 2007. Any household with a member reported to be self-employed was also omitted, due to well-recognised difficulties in evaluating their financial circumstances. Households with a member who was recorded as having a non-contributory pension scheme (primarily consisting of public-sector employees) were also excluded, to limit the heterogeneity in savings incentives described by our sample. The sample population that remained following this screening process was then re-organised into the nuclear families that are the unit of analysis of the structural model. Most family-level statistics required for the model could be obtained by summing over the individual specific data reported by the WAS within each family unit. The notable exception was home ownership, where associated (net) equity was allocated to the family unit of the household reference person (as identified by the survey).

Gender neutrality is a guiding principal adopted for the analysis, and we consequently elected to represent each adult aged 20 or over reported by the WAS separately in our reference population cross-section. Model characteristics specific to the reference adult, including year of birth and age, were equated to the characteristics of the respective adult reported by the WAS. All other model characteristics were set equal to their values reported for each adult's family.[21] This approach meant that family characteristics of couples reported in the WAS were represented twice in our data – once for each partner. To avoid over-sampling couples in our empirical analysis, we divided the weighting factor attached to each couple by two. The total derived sample size considered for analysis was 22,143 family units.

Almost all of the household specific characteristics that are distinguished by our structural model are reported by the WAS. The most notable exception applies to adults who are not reported as working full-time in the WAS, in which case the survey does not report their full earnings potential. These data were imputed wherever necessary, using reduced-form regression equations that are reported in Appendix C.

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

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

Data used to estimate sample moments

The calibration strategy described above was implemented by matching the model to the following sample moments:

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

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

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

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

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

6. The net increase in the proportion of the population aged 20 to 60 reported to hold assets in an ISA, estimated on data observed between 2003 and 2011.

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

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

Table 3: Net rates of ISA take-up averaged between ages 20 and 60 by survey year

year	2004/05	2005/06	2006/07	2007/08	2008/09	2009/10	2010/11
mean	0.0103	0.0159	0.0081	0.0318	0.0222	0.0008	0.0150
std. dev.	0.0035	0.0035	0.0036	0.0038	0.0039	0.0039	0.0039

Source: Author's calculations on annual FRS data, 2003/04 to 2010/11

Notes: Arithmetic averages of age specific statistics calculated from 20 to 60

There is a paucity of publicly available micro-data concerning investment decisions in ISAs. While it is not possible to identify individual net contributions to ISAs or the establishment of new accounts, repeated waves of the FRS do permit calculation of the net increase in the proportion of the population reporting to hold ISA accounts by age and year. That is, net take-up rates of ISAs can be identified from the FRS by subtracting the proportion of individuals identified as holding an ISA at age a in year y from the proportion at age $a + 1$ in year $y + 1$. Unfortunately, as Figure 1 suggests, variation of the annual data series complicates identification of an age profile of net ISA take-up. Following a detailed analysis, the model was calibrated to rates of net ISA take-up averaged between ages 20 and 60. Year specific averages for this statistic are reported in Table 3.

The statistics reported in Table 3 indicate that the year-specific series exhibit substantial variation, even at the high degree of aggregation considered here. The model parameters were consequently adjusted to match to estimates calculated for 2006/07 and 2007/08, which approximately span the range of values implied by the available data.

4.3 Calibrated model parameters

The parameters calibrated endogenously to the structural model are reported in Table 4. Table 4 also reports the simulated moments for net take-up of ISAs under alternative assumptions for salience costs. The match obtained by the model to the remaining moments considered for the calibrations is reported in Appendix D.

The calibrated value for the parameter of relative risk aversion – at 2.25 – is within the broad range identified by the associated literature. Simulations undertaken by Auerbach & Kotlikoff (1987), for example, are based upon a coefficient of risk aversion of 4, while Cooley & Prescott (1995) consider a value of 1. Grossman & Shiller (1981) and Blundell et al. (1994) report estimates just over 1.0, while Hansen & Singleton (1983), Mankiw et al. (1985), and Ziliak & Kniesner (2005) report estimates of approximately 1. Values of the coefficient of risk aversion required to explain the equity premium puzzle (Mehra & Prescott (1985)) are high by comparison, supported by econometric estimates reported by Grossman & Shiller (1981), Mankiw (1985) and Hall (1988). Nevertheless, evidence from attitudinal surveys suggest that the value is unlikely to be greater than 5 (Barsky et al. (1997)).

The value obtained for the intra-temporal elasticity of substitution implies that consumption and leisure are direct substitutes.[22] The utility price of leisure is in the region of 1.0 by construction[23], and the discount factor implies a higher rate of time discounting than the return to ISAs and positive balances of net liquid wealth (1.5% p.a.), or pension assets (2.5% p.a.). The bequest motive increases as the bequest parameters decrease, so that the relatively small parameter values arrived at for ζ_0 and ζ_1 indicate the importance of bequests in matching the model to survey data. The experience effects, which are age independent, imply that earnings potential increases by 4.3% for each year of full-time employment by single adults, and by 1.3% for each year of full-time employment for adults in couples. The factor effects of pension take-up apply a 40% discount to earnings at age 55, and increase linearly to 100% from age 65.

The middle panel of Table 4 reports the calibrated parameter values for salience costs, alongside the associated simulation moments of average net ISA take-up. When salience costs are omitted from the analysis ($\psi_1 = 0.000$), the table indicates that the simulated net increase in ISA participation for the reference population cross-section would be 48% averaged over all ages between 20 and 60. Recall from Table 1 that 64% of the population was reported as holding no ISA in 2006/07. The model therefore suggests that three quarters of this population subgroup would benefit from investing through an ISA.[24] Given this observation, it is little surprise that a salience cost is required to match to the much lower rates of net take-up of ISAs against which the model was calibrated. The salience cost parameters denoted with two and three asterisks in Table 4 were adjusted to match the model to within a single standard deviation of the sample statistics reported for 2006/07 and 2007/08 (Table 3). Although it is difficult to gauge the scale of the calibrated salience costs on the basis of the parameter values reported in the table, an indication of their importance in matching to the sample moments is given by the proximity of ψ_1 to ψ_0 in each case, and the relationship between these two in the assumed preference relation: $\left(\psi_0 - \psi_1 d_{i,a}^{ISA}\right)^{1-\gamma}$. Detail regarding the practical implications of the salience costs identified here is provided by the analysis of behavioural responses to policy that are reported in Section 5.

5 Savings Responses to ISAs and the Role of Decision Making Rigidities

This section explores the effectiveness of the ISA scheme to encourage private sector savings, and how salience costs influence associated behavioural responses. The analysis is based upon simulated data generated by the structural model described in Section 3, subject to parameters derived as discussed in Section 4. Starting from the micro-data of a reference population cross-section (described in Section

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

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

[24] Although Table 1 reports average participation for the full population of 64%, this same participation rate corresponds with the average taken for the population subgroup aged 20 to 64 considered here.

Table 4: Calibrated model parameters adjusted to match behaviour reported for the British population cross-section in 2006

relative risk aversion (gamma)	2.2500
intratemporal elasticity (epsilon)	0.8580
utility price of leisure (alpha)	1.7700
discount factor (delta)	0.9660
bequest motive constant (zeta0)	8.167E-09
bequest motive slope (zeta1)	4.083E-08
salience cost constant (psi0)	1.345E-04
salience cost slope (psi1)*	0.000E+00
ISA net take-up*	0.4754
salience cost slope (psi1)**	1.344E-04
ISA net take-up**	0.0113
salience cost slope (psi1)***	1.341E-04
ISA net take-up***	0.0329
factor effects of pension take-up	
at age 55	0.6000
from age 65	0.0000
experience effect singles (kappa)	0.0425
experience effect couples (kappa)	0.0125

Notes: * no salience cost imposed
** adjusted to match average net ISA take-up in 2006/07
*** adjusted to match average net ISA take-up in 2007/08
measures of "ISA net take-up" calculated on age specific moments, averaged over ages 20 to 60

4.2), the structural model was used to project each reference household forward at annual intervals for 40 years (from 2006 to 2046). Four separate simulations of this type are considered here, which are identical to one another, with the exception of assumptions made regarding the existence of ISAs and associated salience costs:

1. ISAs not included for analysis
2. ISAs included for analysis on the assumption of zero salience costs
3. ISAs included for analysis on the assumption of salience costs reflecting net ISA take-up observed in 2006/07 (*high costs*)
4. ISAs included for analysis on the assumption of salience costs reflecting net ISA take-up observed in 2007/08 (*low costs*)

The only differences between simulations (2), (3), and (4) listed above are the assumed values for the salience cost parameter, ψ_1, which were set, respectively, to each of the three values reported in Table 4. Results describing household savings behaviour generated by the model for the first two simulation scenarios listed here are reported in Table 5.

The rows of Table 5 are ordered into three cascading levels: policy scenario, birth cohort, and lifetime income quintile. Birth cohorts are aggregated into 10 year bands, based upon the birth year of the simulated household reference person. Income quintiles are based on arithmetic averages of net income over the entire simulated life course, and are identified within each birth cohort, so that households in the second income quintile of the 1957-1966 birth cohort may have higher lifetime incomes than those in the third income quintile of the 1937-1946 birth cohort. The columns of the table distinguish between asset classes and simulation years. Note that three asset classes are considered for analysis: ISA investments, private pension wealth, and other liquid wealth, which aggregate up to total wealth – statistics for other wealth can be calculated from the table as residuals.

The basic message of the statistics reported in Table 5 is both simple and clear: the structural model considered for analysis, on the assumption of zero salience costs, implies that households will invest heavily through the ISA scheme, but that most of this saving will be a transfer out of other liquid wealth rather than new household saving; simulated wealth held in pensions and in total are little different with or without ISAs.

Consider, for example, the statistics reported for the highest income quintile of the group of birth cohorts born between 1977 and 1986. These households had reference adults who were aged between 20 and 29 in 2006, which is the starting cross-section from which simulated projections were made. By 2046, this population subgroup will be aged between 60 and 69, which is when household wealth tends to reach its peak. These households are projected by the model to each invest £323,000 on average through the ISA scheme. This same population subgroup is reported to reduce average private pension wealth by 2046 from £922,000 to £916,000, while average total net wealth is projected to increase from £1.38 million to £1.45 million. Hence, of the £323,000 that this population subgroup is projected to invest through ISAs on average, £259,000 is projected to be off-set against private pensions (£6,000) and other liquid wealth (£253,000).

Comparing the statistics for total wealth reported in the top half of Table 5 against the associated statistics reported in the bottom half of the table indicates that simulated total net wealth is as often lower in context of ISAs as it is higher. This is in part attributable to the countervailing price and wealth effects implied by the increase in effective (post-tax) investment returns that are delivered through the ISA scheme. It is also likely to be attributable to the marginal nature of the tax incentives offered through the ISA scheme, relative to the decision making environment more generally.

Table 6 reports statistics for the two simulations that include ISAs and calibrated salience costs. This table indicates that allowing for salience costs in relation to ISAs has a strong influence on the scale of investments made through the simulated scheme, but otherwise leaves unchanged the broad outcomes described above. The assumed salience costs are projected to have the most pronounced effects on the youngest population subgroups, for whom the average value of ISA investments simulated for 2046 is projected to fall from £130,000 in the absence of salience costs to £19,000-24,000 with salience costs. Even amongst households aged 60 and over in the reference population cross-section, however, the simulated salience costs are projected to approximately halve investments made through ISAs.

Focussing on the cohort born between 1977 and 1986, the largest declines in ISA savings after allowing for salience costs are observed for households in the top income quintile. This is not surprising, as these households are projected to invest most heavily in ISAs in general. Of greater interest is the marginal variation in ISA investments observed under simulation scenarios based on *high* and *low* salience costs. A notable feature of the statistics reported in Table 6 is that simulated ISA investments are projected to fall by a similar absolute value in response to a marginal increase in salience costs for households across the income distribution. This is important because it underscores the role that salience costs may play in determining who is able to derive benefit from tax incentivised savings schemes. On balance, results reported in Table 5 suggest that there is potential for a sizeable majority of households to derive some benefit from the ISA scheme, in the absence of salience costs. Results reported in Table 6 suggest that salience costs may limit these benefits to a relatively small population subset.

Table 5: Household wealth simulated with and without ISAs, by birth cohort, income quintile, and calendar year (£000, 2006)

year of	income	ISA investments				private pension wealth				total wealth			
birth	qunitile	2006	2011	2026	2046	2006	2011	2026	2046	2006	2011	2026	2046
						ISAs not included for analysis							
	lowest	0.0	0.0	0.0	0.0	2.2	4.1	8.4	5.8	8.9	21.6	52.9	57.5
1986	2	0.0	0.0	0.0	0.0	3.0	6.3	22.2	17.2	11.1	28.0	79.4	83.9
to	3	0.0	0.0	0.0	0.0	5.6	14.8	81.7	89.7	18.9	49.6	170.5	206.4
1977	4	0.0	0.0	0.0	0.0	8.6	25.2	233.2	411.3	23.2	68.3	372.2	655.1
	highest	0.0	0.0	0.0	0.0	10.8	46.0	527.7	921.6	30.5	106.2	787.2	1382.3
	lowest	0.0	0.0	0.0	0.0	23.6	29.1	19.6	8.8	87.7	98.5	86.6	139.6
1966	2	0.0	0.0	0.0	0.0	46.6	60.5	51.6	22.2	133.0	158.3	151.3	173.4
to	3	0.0	0.0	0.0	0.0	68.1	96.5	105.9	43.8	179.9	221.6	244.9	222.7
1957	4	0.0	0.0	0.0	0.0	133.1	190.0	245.9	98.8	271.5	348.3	463.8	344.1
	highest	0.0	0.0	0.0	0.0	214.9	324.9	565.1	258.1	405.0	548.3	911.1	669.0
	lowest	0.0	0.0	0.0	0.0	38.8	34.8	17.8	4.0	160.7	157.0	161.0	191.1
1946	2	0.0	0.0	0.0	0.0	56.1	47.9	23.0	6.9	206.7	196.6	184.6	205.8
to	3	0.0	0.0	0.0	0.0	93.3	80.4	41.8	13.5	265.1	251.0	226.9	230.8
1937	4	0.0	0.0	0.0	0.0	153.7	133.1	63.3	15.7	374.3	352.3	285.5	255.8
	highest	0.0	0.0	0.0	0.0	524.1	359.1	180.1	41.3	899.2	729.4	535.5	380.3
						ISAs included for analysis on the assumption of zero salience costs							
	lowest	0.7	9.4	44.2	52.8	2.3	4.4	9.8	6.9	9.7	22.5	57.0	63.0
1986	2	0.5	10.9	56.5	62.5	2.9	6.5	23.9	18.9	12.9	29.5	86.7	88.8
to	3	0.6	18.7	73.7	72.2	5.7	15.0	81.7	91.9	16.5	48.3	166.1	195.1
1977	4	0.9	26.2	117.2	141.3	8.3	25.2	226.8	402.4	23.4	68.8	360.1	640.9
	highest	1.1	34.6	201.0	323.1	10.9	45.1	521.2	915.7	30.3	104.7	781.9	1446.3
	lowest	1.7	27.9	54.3	68.4	27.1	33.0	21.7	9.6	94.9	102.9	88.2	135.1
1966	2	3.2	41.5	62.7	65.0	47.9	63.6	54.1	23.1	135.7	158.5	146.1	161.9
to	3	3.3	54.0	73.8	70.7	72.8	100.4	106.2	43.3	183.0	219.8	231.1	205.5
1957	4	5.0	65.2	108.3	93.8	130.1	185.6	241.0	97.2	271.9	341.6	450.4	330.1
	highest	7.8	75.3	192.0	143.8	208.5	317.1	559.6	256.7	393.5	533.8	909.1	662.1
	lowest	7.0	24.9	66.2	50.6	46.1	38.0	18.9	4.2	186.2	175.0	168.0	189.9
1946	2	8.9	25.7	68.8	70.0	67.1	59.0	27.0	7.6	229.5	218.1	188.5	207.2
to	3	10.0	26.8	74.4	94.3	96.5	83.3	41.3	13.6	266.7	251.3	219.1	231.0
1937	4	12.5	28.0	74.1	84.8	144.0	124.9	60.5	14.8	351.4	332.0	271.4	250.0
	highest	25.2	38.7	97.2	132.2	512.3	349.7	179.4	41.1	878.6	713.5	527.6	379.7

Notes: simulated weighted averages of asset values within population subgroups, defined in £000, at 2006 prices

quintile groups identifed within birth cohorts, and with respect to average net income earned over the entire simulated lifetime

Table 6: Household wealth simulated with ISAs, by assumed salience cost, birth cohort, income quintile, and calendar year (£000, 2006)

year of birth	income qunitile	ISA investments 2006	ISA investments 2011	ISA investments 2026	ISA investments 2046	private pension wealth 2006	private pension wealth 2011	private pension wealth 2026	private pension wealth 2046	total wealth 2006	total wealth 2011	total wealth 2026	total wealth 2046
ISAs included for analysis on the assumption of salience costs reflecting net ISA take-up observed in 2007/08 (low)													
	lowest	0.8	1.7	8.6	12.8	2.5	4.4	9.1	6.2	9.9	21.1	50.2	58.2
1986	2	0.5	1.9	10.6	14.3	2.9	6.3	21.5	16.6	10.5	26.4	78.0	83.7
to	3	0.6	3.8	15.6	18.0	5.6	14.5	81.0	88.8	19.0	48.7	170.4	200.6
1977	4	0.9	5.9	20.0	34.9	8.4	25.0	225.5	404.3	23.9	68.5	361.9	645.5
	highest	1.1	6.5	18.2	37.9	10.7	45.0	526.1	919.4	29.6	104.1	784.8	1383.7
	lowest	2.6	5.7	13.4	20.1	27.2	32.8	21.7	9.6	99.1	105.1	89.2	139.1
1966	2	2.9	9.3	17.5	21.7	47.3	62.6	53.9	23.0	131.3	156.3	148.6	170.0
to	3	3.7	15.1	25.5	29.9	73.6	101.7	108.8	43.9	193.6	232.1	246.0	216.8
1957	4	4.6	17.4	37.8	43.1	129.4	184.1	233.3	94.8	262.6	335.7	442.5	330.3
	highest	7.2	23.5	68.2	69.0	208.9	318.1	561.0	256.8	392.6	536.3	907.0	662.3
	lowest	8.5	9.4	25.0	21.8	46.7	40.3	19.4	4.7	188.5	179.7	169.1	193.1
1946	2	8.4	11.8	29.5	42.7	60.1	51.9	25.7	8.1	213.7	201.9	187.5	209.7
to	3	9.2	13.9	35.7	53.9	97.0	83.7	40.7	12.4	266.7	252.6	220.7	227.2
1937	4	13.1	19.1	42.5	40.7	148.4	128.6	62.2	15.1	363.5	343.8	279.2	251.7
	highest	24.5	32.5	72.8	64.9	513.7	350.7	178.0	41.0	880.1	713.2	523.9	378.5
ISAs included for analysis on the assumption of salience costs reflecting net ISA take-up observed in 2006/07 (high)													
	lowest	0.8	1.7	6.8	8.8	2.5	4.4	9.1	6.3	10.1	21.4	50.4	58.1
1986	2	0.5	1.8	7.8	9.1	2.9	6.3	21.8	17.0	10.4	26.7	78.2	83.7
to	3	0.6	3.7	13.2	14.6	5.6	14.4	81.2	88.9	19.0	48.8	170.6	201.3
1977	4	0.9	5.8	17.5	29.6	8.3	24.8	226.4	406.5	23.8	68.5	363.4	648.4
	highest	1.1	6.4	16.3	32.8	10.8	45.2	525.7	918.7	29.6	104.5	784.3	1382.0
	lowest	2.6	5.8	13.5	20.2	27.0	32.6	21.7	9.6	99.0	104.9	89.3	139.1
1966	2	2.8	9.3	17.6	21.7	47.5	62.8	53.9	23.0	131.4	156.5	148.6	169.9
to	3	3.7	15.1	25.6	29.9	74.2	101.9	109.2	43.9	194.3	232.3	246.5	216.8
1957	4	4.6	17.4	37.7	43.1	128.8	183.9	233.0	94.8	261.8	335.5	442.0	330.3
	highest	7.2	23.5	68.2	69.0	208.9	318.1	561.0	256.8	392.6	536.3	907.0	662.3
	lowest	8.5	9.4	25.0	21.6	46.7	40.3	19.4	4.7	188.5	179.7	169.1	193.1
1946	2	8.3	11.9	29.5	42.3	60.5	51.8	25.7	8.1	214.4	201.6	187.5	209.7
to	3	9.2	13.9	35.7	53.4	96.6	83.8	40.7	12.4	266.0	252.9	220.7	227.2
1937	4	13.1	19.1	42.5	40.6	148.4	128.6	62.2	15.1	363.5	343.8	279.2	251.7
	highest	24.5	32.5	72.8	64.9	513.7	350.7	178.0	41.0	880.1	713.2	523.9	378.5

Notes: simulated weighted averages of asset values within population subgroups, defined in £000, at 2006 prices

quintile groups identifed within birth cohorts, and with respect to average net income earned over the entire simulated lifetime

Further detail regarding the welfare stakes involved is provided by Table 7, which reports compensating variations to introducing the ISAs scheme. The statistics reported in Table 7 are based on the hypothetical policy counterfactual that the ISAs scheme was introduced in 2006. Table 7 reports the money value that the structural model implies would need to be given to (positive) or taken from (negative) different population subgroups on average, to make each household indifferent following introduction of the ISAs scheme. Hence, negative compensating variations indicate a welfare benefit derived from ISAs (as a household would need to have money taken away from it following the reform to be as well off as it was prior to the reform). For example, the model implies that introducing the ISAs scheme in 2006 would improve the welfare of individuals in the highest income quintile and born between 1977 and 1986 by an amount equivalent to £5910 on average, if no associated salience costs applied.

The simulated welfare benefits following introduction of ISAs in the absence of salience costs rise from an average of £4200 amongst households with reference adults aged 20 to 29 (in 2006), up to £10400 for reference adults aged 40 to 49, before falling to £2100 for those aged 60 to 69. Furthermore, the welfare effects are strongly positively related to (lifetime) income for younger households, and independent of income quintile for older households. These observations reflect the inverse relationship between age and prospective investment horizons.

As expected, the projected welfare benefits of ISAs fall as the assumed salience costs rise. The steepest falls are projected for households of reference adults aged 40 to 49, who are projected to enjoy the largest gains under ISAs in the absence of salience costs. The distributional effects of the assumed salience costs are of particular note. The salience costs that are required to match the structural model to observed rates of net ISA take-up are sufficient to reverse the positive relationship that is noted above between income and the welfare benefits derived under ISAs in the absence of salience costs. Allowing for salience costs is consequently sufficient in the current context to ensure that ISAs deliver a more substantial welfare benefit to those on modest lifetime incomes than those at the top of the distribution.

This striking result is a product of two key factors: the functional form that salience costs are assumed to take, and the scale of costs required to match to observed rates of ISA take-up. The preference structure assumes that the salience cost is independent of financial circumstances in general, and of the marginal utility of consumption in particular. As marginal utility is decreasing in consumption, rich households place a higher money value on the assumed salience cost than poor households. As noted in Section 3.4, the omission of housing from the model is likely to have increased underlying incentives to participate in ISAs, which may suggest the need for a punitive salience cost to match the model to observed rates of ISA take-up. These two factors in concert drive the distributional welfare effects that are reported here.

Table 7: Compensating Variations to introduction of the ISA scheme in 2006, by birth cohort, income quintile, and salience costs (£2006)

year of birth	income qunitile	salience cost of ISA investment no cost	low cost*	high cost*
	lowest	-2709	-604	-528
1986	2	-3075	-562	-489
to	3	-4193	-505	-433
1977	4	-5174	-590	-490
	highest	-5910	-528	-430
	lowest	-7707	-258	-257
1966	2	-9395	-122	-124
to	3	-11277	-32	-32
1957	4	-11491	-2	-1
	highest	-12227	-1	-1
	lowest	-1812	-20	-17
1946	2	-1776	-16	-15
to	3	-2106	-18	-18
1937	4	-1924	-7	-6
	highest	-2741	-3	-3

Notes: simulated weighted averages of compensating variations within
population subgroups, defined in £000, at 2006 prices
quintile groups identifed within birth cohorts, and with respect
to average net income earned over the entire simulated lifetime
*low cost matched to reflect net ISA take-up in 2007/08
*high cost matched to reflect net ISA take-up in 2006/07

The results reported in Table 7 indicate that salience costs can have the capacity to turn relative winners from a policy reform into relative losers, and vice versa. An example suggested by the current framework would be an investment scheme that delivers increased returns to investment – thereby benefiting cash-rich individuals – but requires a substantial time investment – thereby penalising time-poor individuals. Any welfare analysis of such a scheme that focussed only on the increased investment returns is likely to be highly mis-leading.

Nevertheless, the form of the salience costs that are considered here, and the associated welfare effects, suggest plausible alternatives for analysis that would likely alter the results obtained. It seems reasonable, for example, to assume that social networks and individual characteristics might imply an inverse relationship between salience costs of ISA participation and household financial circumstances. Alternatively, the existence of financial advisory services argues in favour of an upper bound to the financial equivalent of the salience cost incurred in relation to associated decision making. Each of these alternatives would work in favour of improved welfare effects for higher net worth households, relative to the responses reported in Table 7.

One of the features that distinguishes ISAs from other tax advantaged savings schemes is the modest upper bounds that the scheme has imposed on annual contributions. This feature is interesting because, unlike lifetime limits, annual allowances can have an important influence on the timing of associated investments. The influence of annual limits on the timing of investment contributions depends crucially on the strength and temporal variation of preferences for saving; modest contributions limits will be most effective in encouraging higher contributions early in life where preferences for saving start low and display strong growth during the life course. Table 8 reports statistics that indicate how these effects bear out in the structural model considered for this study.

Table 8 reports the net increases in ISA investments simulated during decade intervals for reference people aged 20 to 29 in 2006, distinguished by the annual contributions limits imposed and the assumed salience costs. These statistics, taken together, indicate that the annual contributions limits tend to have the most bite early in the simulated lifetime.

Focussing first upon the simulations that suppress salience costs, the lower contributions limit is reported to push net increases in ISA investments forward through time for households throughout the income distribution. It is notable, however, that this simulated shift in the timing of ISA investments does not appreciably alter the total investments simulated to be held in ISAs by state pension age; doubling the annual contributions limit from £7000 to £14000 increases average ISA investments at 2046 by just £2000 (£130000 c.f. £132000) where salience costs are suppressed.

Relaxing the annual contributions limit is also found to increase ISA investments early in life where investment decisions are assumed to be influenced by salience costs. Unlike the simulations that suppress salience costs, however, doubling the annual contributions limit and allowing for salience costs results in higher net increases in ISA investments on average throughout the simulated period 2006-16. This result is observed because the assumed salience costs work to screen out households that derive marginal benefit from ISAs, and for whom the constraints imposed by contributions limits are least important. In this context, salience costs consequently work to exaggerate the reduction in ISA investments attributable to contributions limits.

6 Conclusions

This study uses a structural model of savings and labour supply in the context of uncertainty to consider behavioural responses to Individual Savings Accounts; a tax incentivised savings scheme currently operated in the UK. The model assumed for analysis augments the standard life-cycle framework to include salience costs incurred on the first investment made through an ISA.

Matching the model to survey data indicates that salience costs are useful in reflecting observed rates of participation in the ISA scheme. Micro-data concerning investments made through the ISA scheme suggest that net annual rates of take-up averaged across a representative sample of the British population aged between 20 and 60 do not exceed 5%. In the absence of salience costs, the structural model assumed here suggests that 48% of the population aged 20 to 60 would choose to make their first

Table 8: Average net accretion of ISA investments simulated during preceding ten years by cohorts aged 20 to 29 in 2006 (č2006)

year	lowest quintile*	2	3	4	highest quintile	population average
£7000 upper bound on annual ISA contributions and zero salience costs						
2016	23654	28949	46550	63318	85337	49562
2026	19774	27079	26545	52948	114508	48171
2036	-2926	2677	1837	36734	104791	28623
2046	11582	3239	-3391	-12605	17306	3226
£14000 upper bound on annual ISA contributions and zero salience costs						
2016	23711	29828	48959	67309	101928	54347
2026	16271	24531	24959	56463	162338	56912
2036	3737	3739	4178	40297	84343	27259
2046	8565	4130	-4863	-17714	-25484	-7073
*£7000 upper bound on annual ISA contributions and high salience costs***						
2016	2992	2946	8347	11846	11929	7612
2026	3005	4374	4198	4685	3302	3913
2036	1096	1366	884	7821	6859	3605
2046	865	-18	505	4308	9675	3067
*£14000 upper bound on annual ISA contributions and high salience costs***						
2016	5079	5684	10990	14889	17523	10833
2026	1803	3286	2645	15632	21234	8920
2036	507	329	379	9054	11838	4421
2046	774	-323	59	5055	17268	4567

Notes: simulated weighted averages of asset values within population subgroups, defined in £2006
* quintile groups defined on average net income earned over the entire simulated lifetime
** "high" salience costs adjusted to match rates of ISA take-up reported for 2006/07

ISA investment in 2006/07, equivalent to three out of every four individuals who are reported to hold no ISA investments at that time. The calibrations reported here reveal that allowing for salience costs on the first investment in an ISA allows the model to reflect the rates of net ISA take-up described by survey data.

Policy analysis conducted with the model suggests that the price effects associated with the ISA scheme are insufficient to motivate appreciable increases in household saving, with or without the addition of salience costs. In the absence of a salience cost, households are projected to invest heavily through ISAs. By the time that cohorts aged 20 to 29 in 2006 are in their sixties, for example, they are projected on average to invest just under two thirds of their non-pension wealth – £130,400 in 2006 prices – through the ISA scheme. At the same time, these households are project to reduce their pension savings by £2,000 on average, and reduce savings through other liquid wealth by £118,600, so that aggregate net savings in context of ISAs increase by £9,800 on average.

The salience costs that match the model to net rates of ISA take-up described by survey data reduce the scale of investments that are projected to be made through the ISA scheme, but leave most other effects on population averages qualitatively unchanged. The cohorts aged 20 to 29 in 2006, for example, are projected to invest £19,000-23,600 on average through ISAs by the time they are in their sixties, off-set by a reduction of £2,000 in pension wealth and £19,700-24,200 in other liquid wealth, so that total net savings are projected to fall by approximately £2,000 on average.

The assumed salience costs are also found to have an important influence on the distributional effects of ISAs. Where salience costs are incurred, they are assumed to decrease expected lifetime utility by the same absolute amount, independent of a households's financial circumstances. This form, and the diminishing marginal utility of consumption that is a property of the assumed preference relation, implies that rich households attach a larger financial value to the assumed salience costs than poorer households. The assumed salience costs consequently act to depress the welfare benefits derived from ISAs by richer households, relative to poorer households, and to disproportionately affect the behaviour or richer households as a result. The results obtained in this respect highlight the capacity of salience costs to have a pronounced influence on determining relative winners and losers of policy reform. They also suggest interesting alternatives regarding the function form to assume for salience costs; exploring this issue remains a subject for future research.

References

Akerlof, G. A. (1991), 'Procrastination and obedience', *American Economic Review* **81**, 1–19.

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

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

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

Beshears, J., Choi, J. J., Laibson, D. & Madrian, B. C. (2008), The importance of default options for retirement saving outcomes: evidence from the United States, *in* S. J. Kay & T. Sinha, eds, 'Lessons from Pension Reform in the Americas', Oxford University Press, Oxford.

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

Carroll, G. D., Choi, J. J., Laibson, D., Madrian, B. C. & Metrick, A. (2009), 'Optimal defaults and active decisions', *Quarterly Journal of Economics* **124**, 1639–1674.

Choi, J. J., Laibson, D., Madrian, B. C. & Metrick, A. (2002), 'Defined contribution pensions: plan rules, participation choices, and the path of least resistance', *Tax Policy and the Economy* **16**, 67–113.

Choi, J. J., Laibson, D., Madrian, B. C. & Metrick, A. (2003), 'Optimal defaults', *American Economic Review* **93**, 180–185.

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

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

Flavin, M. & Nakagawa, S. (2008), 'A model of housing in the presence of adjustment costs: a structural interpretation of habit persistence', *American Economic Review* **98**, 474–495.

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

Gigerenzer, G. & Selten, R. (2002), Rethinking rationality, *in* G. Gigerenzer & R. Selten, eds, 'Bounded Rationality - the adaptive toolbox', MIT Press, London.

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

Gustman, A. L. & Steinmeier, T. L. (1986), 'A structural retirement model', *Econometrica* **54**, 555–584.

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

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

Johnson, E. J., Bellman, S. & Lohse, J. L. (2002), 'Defaults, framing, and privacy: why opting-in-opting-out?', *Marketing Letters* **13**, 5–15.

Johnson, E. J. & Goldstein, D. (2003), 'Do defaults save lives?', *Science* **302**, 1338–1339.

Johnson, E. J., Hershey, J., Msezaros, J. & Kuhnreuther, H. (1993), 'Framing, probability distortions, and insurance decisions', *Journal of Risk and Uncertainty* **7**, 35–51.

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

Khan, A. & Thomas, J. K. (2008), Adjustment costs, *in* S. N. Durlauf & L. E. Blume, eds, 'The New Palgrave Dictionary of Economics', Palgrave Macmillan, Basingstoke.

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

Lucchino, P. & van de Ven, J. (2013), Empirical analysis of household savings decisions in context of uncertainty; a cross-sectional approach. National Institute Discussion Paper 406.

Madrian, B. & Shea, D. (2001), 'The power of suggestion: inertia in 401(k) participation and savings behaviour', *Quarterly Journal of Economics* **141**, 1149–1187.

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

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

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

O'Donoghue, T. & Rabin, M. (1999*a*), 'Doing it now or later?', *American Economic Review* **89**, 103–124.

O'Donoghue, T. & Rabin, M. (1999*b*), Procrastination in preparing for retirement, *in* H. J. Aaron, ed., 'Behavioral Dimensions of Retirement Economics', Brookings Institute, Washington.

O'Donoghue, T. & Rabin, M. (2001), 'Choice and procrastination', *Quarterly Journal of Economics* **116**, 121–160.

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

Rust, J. & Phelan, C. (1997), 'How social security and medicare affect retirement behavior in a world of incomplete markets', *Econometrica* **65**, 781–831.

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

Simon, H. A. (1955), 'A behavioural model of rational choice', *Quarterly Journal of Economics* **69**, 99–118.

Vigna, S. D. & Malmendier, U. (2004), 'Contract design and self-control: theory and evidence', *Quarterly Journal of Economics* **119**, 353–402.

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

A Numerical Solution Routines

An analytical solution to the utility maximisation problem described in Section 3 does not exist, and numerical solution routines were consequently employed. These solution routines are structured around a 'grid' that over-lays all feasible combinations of individual specific characteristics (the state space). Log scales were used to divide the state space assumed for liquid wealth, earnings potential, private pension rights and ISAs. These log-scales provide additional detail at the low end of the income/wealth distribution, where the assumed tax function can exhibit substantive non-concavity. Following extensive experimentation, the following dimensions were assumed for the respective grid dimensions: 16 points for liquid wealth between ages 20 and 74, and 51 points between ages 75 and 130; 16 points for earnings potential between ages 20 and 74; 11 points for private pension rights from age 20 to 74, and 51 points between ages 75 and 130; 11 points for ISAs from age 20 to 130; 2 points for wage offers between ages 20 and 74; 2 points for pension receipt from age 55 to 75; 2 points for relationship status from age 20 to age 89. Hence, the grid considered for analysis comprised 11,381,403 individual cells.

As noted in Section 3, the model assumes that there is a maximum potential age to which any individual may survive, A. At this age, the decision problem is deterministic, and trivial to solve. The solution routine that we employed starts by solving for utility maximising decisions at all intersections of the grid that correspond to this final period of life, and stores both the maximising decisions and optimised measures of utility (the value function). These solutions at grid intersections for age A were used to approximate solutions at age A more generally, via the cubic interpolation routine that is described in Keys (1981). Cubic interpolation was necessary in the current context, because of the curvature of the preference relation and our focus on more than one intensive decision margin.

Given results for age A, the solution routine that we used then considers decisions at intersections corresponding to the penultimate age, $A-1$. Here, expected lifetime utility is comprised of the utility enjoyed at age $A-1$, and the impact that decisions taken at age $A-1$ have on circumstances – and therefore utility – at age A. Given any decision set at age $A-1$, Φ_{A-1}, the solution routine projects forward the set of individual specific characteristics at age A, z_A, that is implied by the processes assumed to govern intertemporal transitions (e.g. equation 12 for wage potential). If characteristics at age A are uncertain, then each potential characteristic vector z_A^p is projected forward with an assigned probability pr_A^p.

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

Families are assumed to decide over two continuous domains defining the split of the household budget into consumption, ISA assets, and (other) liquid wealth, and a series of decisions over discrete alternatives relating to labour supply, pension participation, and the take-up of pension benefits. The model uses repeated calls of the value function via Brent's method (described in Press et al. (1986)) to search over feasible decisions on the continuous domain for a local maximum to expected lifetime utility, for each potential combination of discrete decision alternatives. Of all feasible alternative solutions, the one associated with the maximum numerical approximation of expected lifetime utility was taken as the solution to the lifetime decision problem.

As the decision problem considered here is neither smooth nor concave, we performed a supplementary routine to test the robustness of solutions identified by the procedure described above. This involved adding an additional step to the procedure in which a limited grid test above and below the initial candidate solution obtained by Brent's method was explored to identify alternative candidate solutions. If an alternative local maximum was identified, then the solution routine extends the grid search domain above or below, depending on the location of the new candidate solution, relative to the initial candidate. This process was continued until no further local maxima were found. Of all feasible solutions, the one that maximises the value function was then selected.

Simulating data for a reference population through time

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

B Exogenously Identified Model Parameters

The parameters that were evaluated exogenously from the model structure were obtained are concerned with four key issues: private income, taxation, private pensions, and demographics.

B.1 Private income

The model requires three separate (real) interest rates: the return to liquid assets, and lower and upper bounds for interest charges on debt. The return to liquid assets – applied to both ISA investments and positive balances of other liquid wealth – was set equal to 1.5% per annum, which is the average yield on long-term Gilts between 1970 and 2010 (reported by the Bank of England). The lower bound interest charge on debt was set to 8.4% per annum, which is the average interest paid on personal loans between 1995 and 2010, and the upper bound was set to 15.4% per annum, which is the average interest paid on credit card debt between 1995 and 2010 (also reported by the Bank of England).[25]

As noted in Section 4.1, the persistence of latent wages, β, and the factor effects of alternative labour supply decisions, λ^{emp}, were identified exogenously (see Section 3.5 for formal definitions of these parameters). The specification of latent wages was defined as a random walk with drift, so that $\beta = 1.0$. Full-time employment of all adult family members was assumed to reduce family leisure time by 40%. We assume constant wage rates, so that part-time employment has the same proportional effect on time spent in labour as it does on aggregate wages earned, relative to full-time employment. Data reported by the 2006 Annual Survey of Hours and Earnings indicate that part-time employees worked on average 50% of the hours of full-time employees for 30% of the (pre-tax) pay. The model assumes the midpoint between these two statistics for adjusting labour hours and wages for part-time employment; $\lambda^{emp} = 0.4$. Finally, full-time employment of the 'chief wage earner' in an adult couple was assumed to return 60% of the household's full latent wage.

B.2 Tax and benefits policy

Taxes and benefits are formally modelled on the transfer system that applied in the UK in 2006/07. Simulated transfer policy distinguishes between two periods of the life-course, subject to an age threshold set equal to state pension age $t_{SPA}(b)$. The program code adopted to simulate taxes and benefits in the model can be obtained from the authors upon request.

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

Taxes and benefits during the working lifetime

Prior to t_{SPA}, taxes and benefits are based upon the *Tax Benefit Model Tables* (TBMT) produced by the Department for Work and Pensions, which are designed to capture the key elements of the transfer system that applied to the healthy working-aged population. These include income taxes, national insurance contributions, the working tax credit, the child tax credit, the child benefit, housing benefit, council tax benefit, Jobseeker's allowance, healthy start allowances, and free school meals.

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

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

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

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

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

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

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

Taxes and benefits from state pension age

A similar approach was taken to model taxes and benefits from state pension age as described above for the working lifetime. Unlike for the working lifetime, however, the specification of transfer payments from state pension age could not be based on the TBMT, as these do not cover retirement benefits.

Rather, we referred to official rates and thresholds of the transfer schemes that were applied in 2006/07 to specify this aspect of the transfer system.

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

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

Transfer policy through time

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

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

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

B.3 Private pensions

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

There is a great deal of diversity in private pension arrangements in the UK, and in the details of occupational pensions in particular. Panel A of Figure 2 reveals that – although not universal – a sizeable majority of employees were offered some form of contribution in respect of participation in an employer sponsored pension. Eligibility to an employer sponsored pension is reported to increase from a low of between 30 and 40 percent among individuals on less than half of median earnings (increasing by age group), to between 75 and 85 percent among individuals on more than one and a half times median earnings. Approximately 60% of employees on median earnings are reported as being eligible to some form of employer pension contribution, with this rising to just over 80% for employees with very high earnings. The figure also indicates that eligibility to an employer pension contribution exhibited a stronger relationship with employee earnings than it did with age. Following these observations, we set g_l^P equal to 75% of median earnings throughout the simulated lifetime.

Panel B of Figure 2 indicates that, for employees who received an employer pension contribution, the distribution of employer pension contributions was dominated by a single mode between 12.5 and 15 per cent of employee wages. Bearing in mind that the decision by an employee not to participate in their employer's sponsored pension plan would usually result in the forfeiture of any matching employer pension contributions on offer, the scale of the employer contributions reported in Panel B provides an indication of how important these contributions were in supporting the UK system of private sector pension provisions. Panel B of Figure 2 also reveals that there was very little difference between the distributions of employer pension contributions offered in low-pay industries and the wider labour market, with the principal disparity being that employer contributions in excess of the mode were less frequent among employees in low pay industries. We consequently set the rate of employer contributions to 14%; the rate of private contributions to pension wealth was set to the 'normal' contribution rate stated in the guidance to interviewers for the FRS, equal to 8%.

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

B.4 Demographics

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

Life expectancy

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

We assume a maximum potential age of 130 for the simulations.[26] The age specific mortality rates reported by the ONS were extended beyond age 100 using a smooth sigmoidal function to equal 1.0 (certain death) at age 130. Furthermore, the time series dimension of the age specific mortality rates reported by the ONS was extended to all age and year combinations feasible for any modelled birth cohort by assuming a constant exponential growth factor of 0.975 from the most approximate year described by the ONS data to exogenously assumed age and sex specific asymptotes for the distant past

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

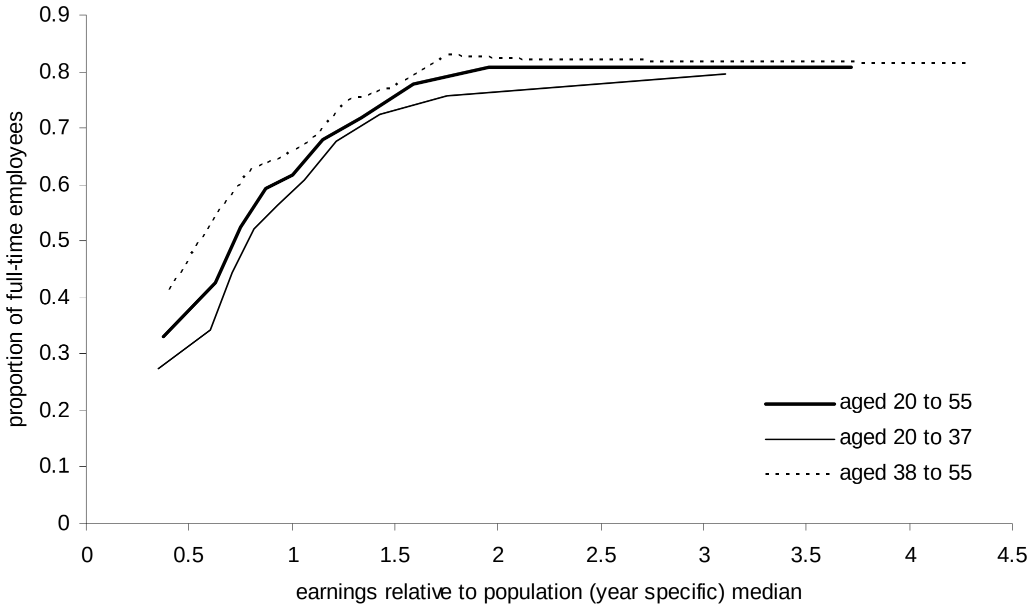

Panel A: eligibility to any employer pension contribution

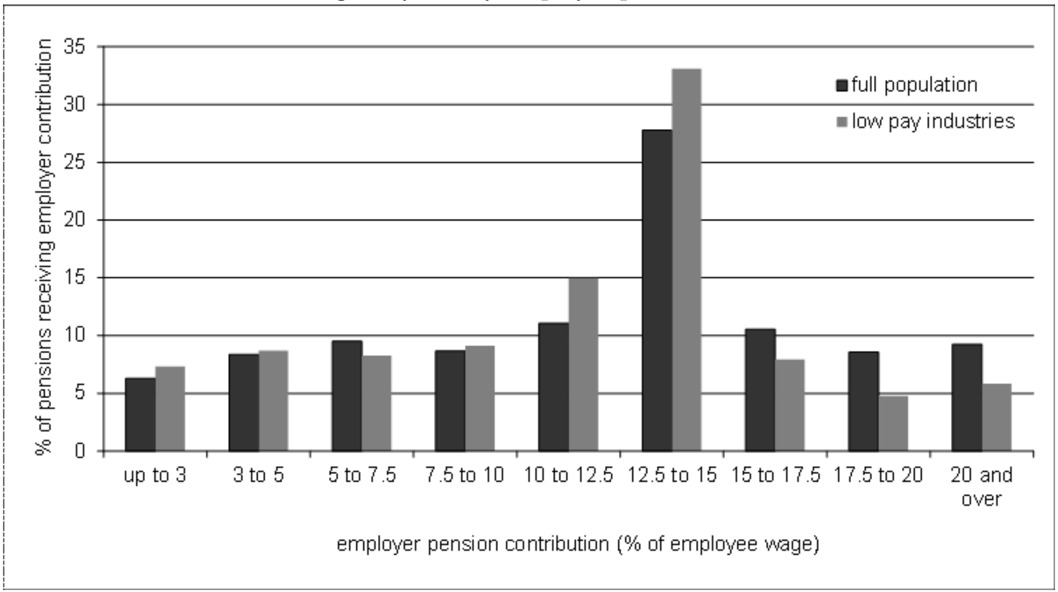

Panel B: employer contribution where some contribution was made

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

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

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

Excludes employees in the public sector and the self-employed

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

Figure 2: Eligibility rates of full-time employees to employer sponsored pensions by age and earnings

and future.

The model specification does not distinguish reference adults by their gender. The gender specific mortality rates that are reported by the ONS were consequently combined into a single series based on implied gender weights. Consider, for example, the cohort born in 1960. Assuming zero migration and equal numbers of males and females at age 16, the gender specific mortality rates reported for this birth cohort by the ONS can be used to project the ratio of men to women through time. This ratio was used to obtain a weighted average of the gender specific mortality rates reported by the ONS for each modelled birth cohort. To avoid imposing unwarranted structure on the parameters, the mortality rates were stored in the form of a transition matrix, comprised of 111 rows (representing ages 20 to 130), and 112 columns (representing years 1951 to 2060, with two additional rows to represent the distant past and future). The transition probabilities used can be obtained from the author upon request.

Relationship status

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

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

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

The first and second problems identified above were addressed by adjusting marriage rates to age

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

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

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

Number of dependent children

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

C Regression Models Used to Input Missing Earnings Data

Four regression equations were estimated on the cross-sectional WAS data, which were used to impute earnings for individuals who were not reported to be working full-time in the sample; separate equations for men and women, and separate equations for those aged under age 50 and those aged 50 years or over. The specifications adopted for this analysis were constrained only by the information reported by the WAS, which includes a high degree of financial detail. After experimenting with various alternatives, regression results for the assumed earnings equations are reported in Table 9.

The parameter values reported in Table 9 indicate that earnings are positively correlated with education, lower for students, and tend to vary positively with health and socio-economic status. Self-reported savers tend to earn more than non-savers, and earnings are positively related to aggregate wealth, home ownership, and mortgage value.

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

D Simulated vs Sample Moments

Table 9: Regression Estimates for log Earnings, Controlling For Sample Selection

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

Regression estimates for target equations controlling for sample selection

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

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

Robust standard errors reported

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

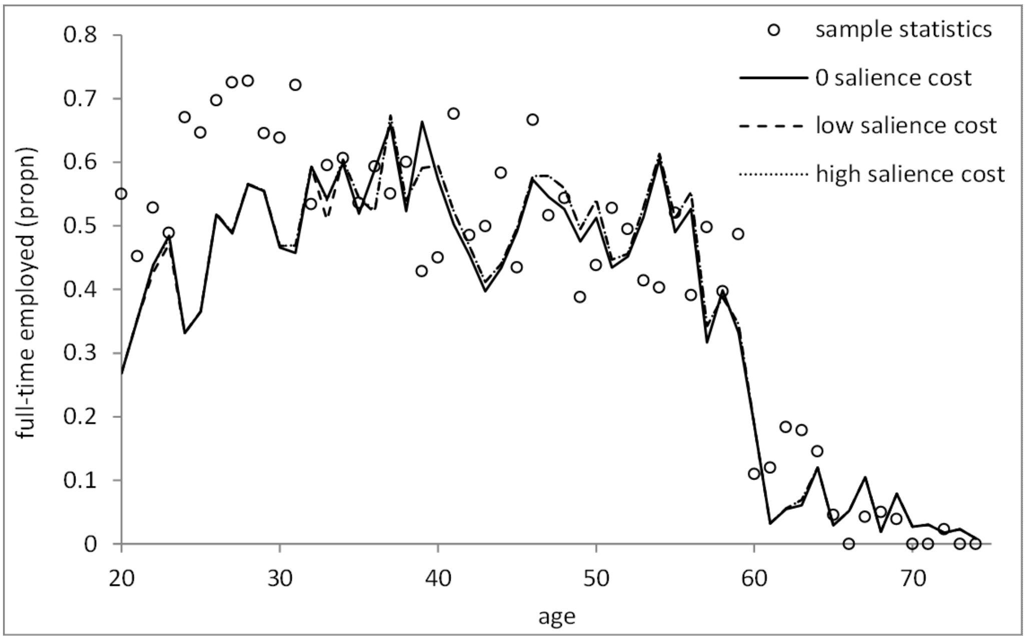

Panel A: rates of full-time employment

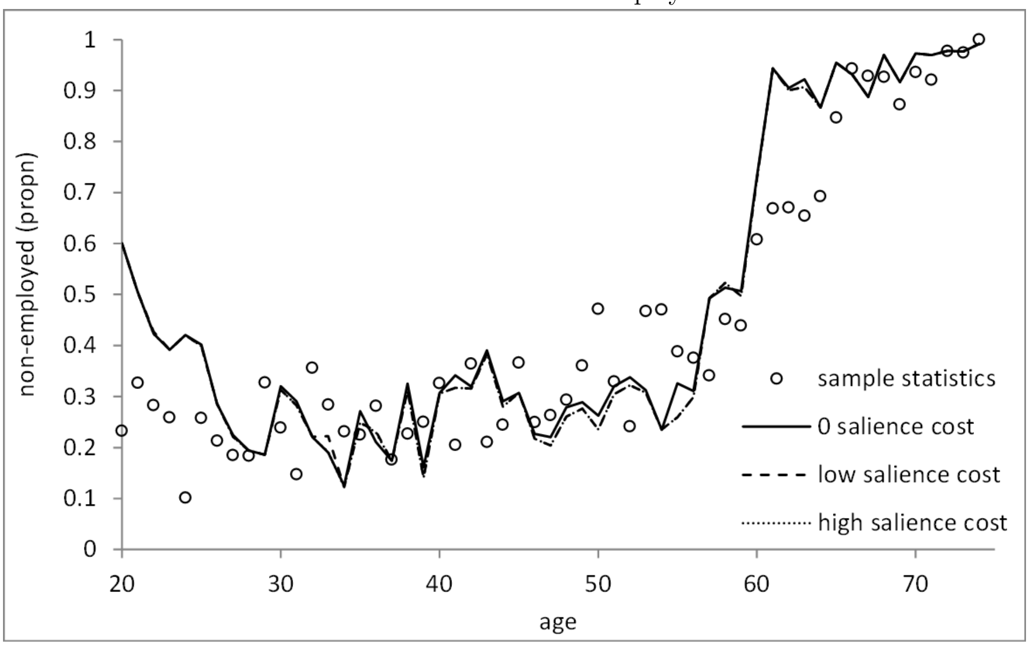

Panel B: rates of non-employment

Figure 3: Employment rates of single adults by age; simulated and sample statistics for 2006

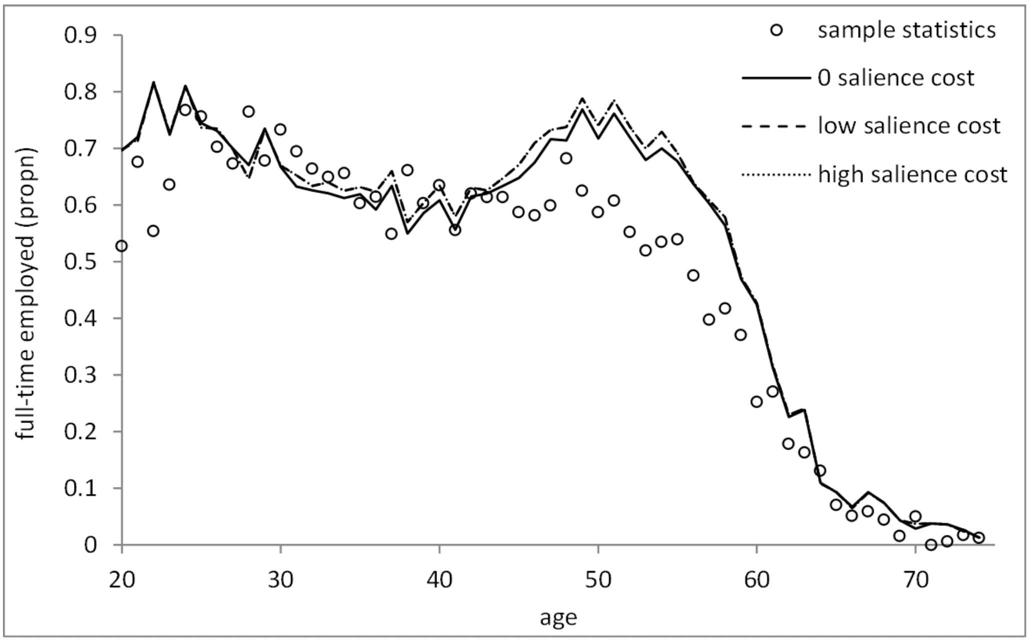

Panel A: rates of full-time employment

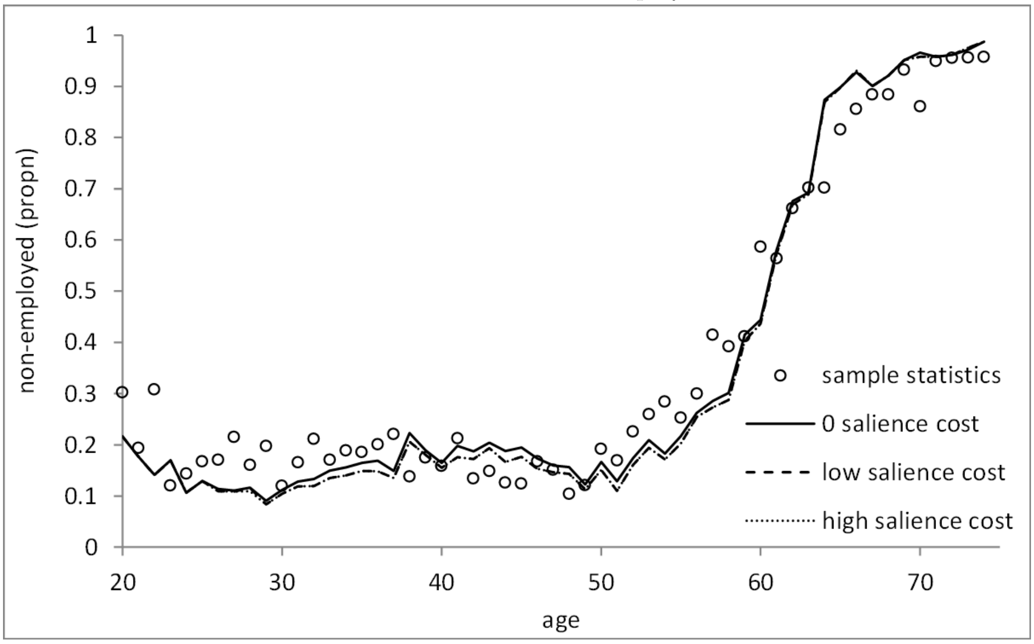

Panel B: rates of non-employment

Figure 4: Employment rates of adults in cohabitating couples by age; simulated and sample statistics for 2006

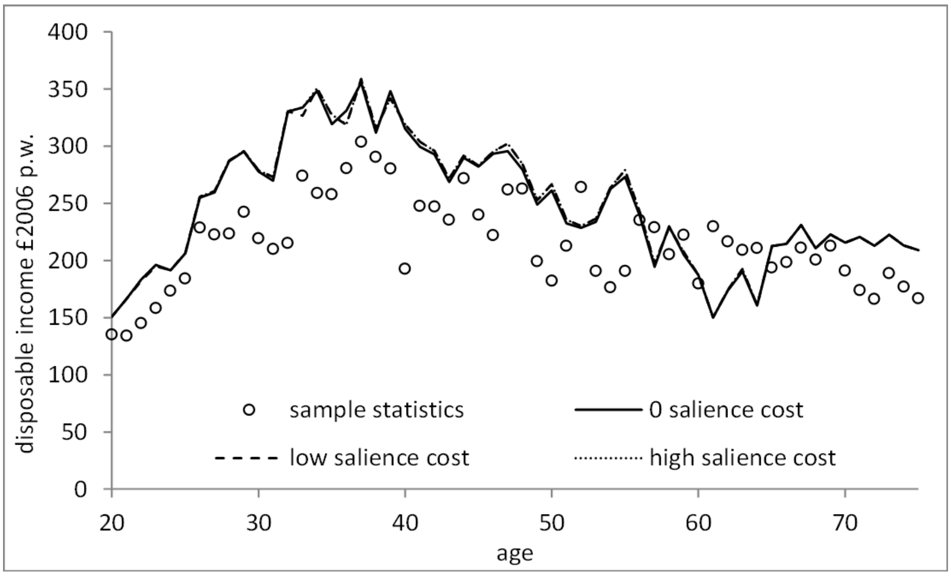

Panel A: single adults

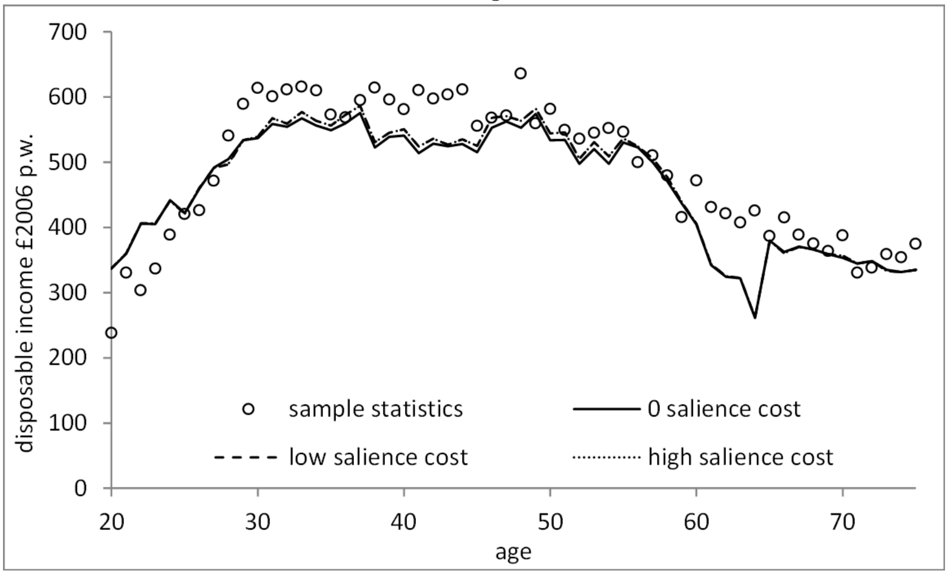

Panel B: adult couples

Figure 5: Geometric mean of disposable household income by age and relationship status; simulated and sample statistics for 2006

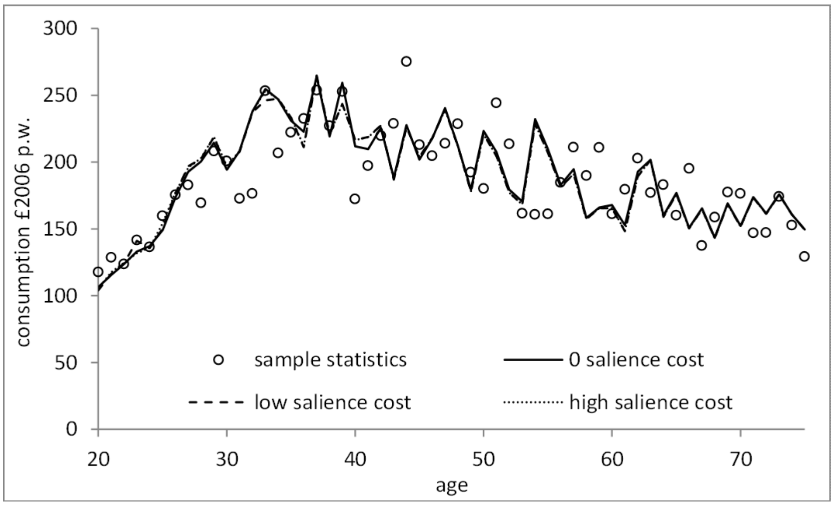

Panel A: single adults

Panel B: adult couples

Figure 6: Geometric mean of household consumption by age and relationship status; simulated and sample statistics for 2006

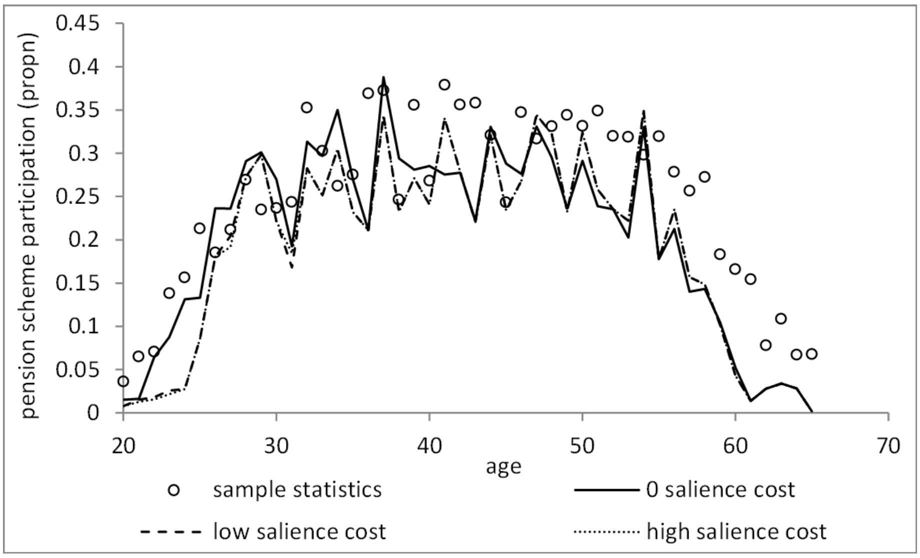

Panel A: single adults

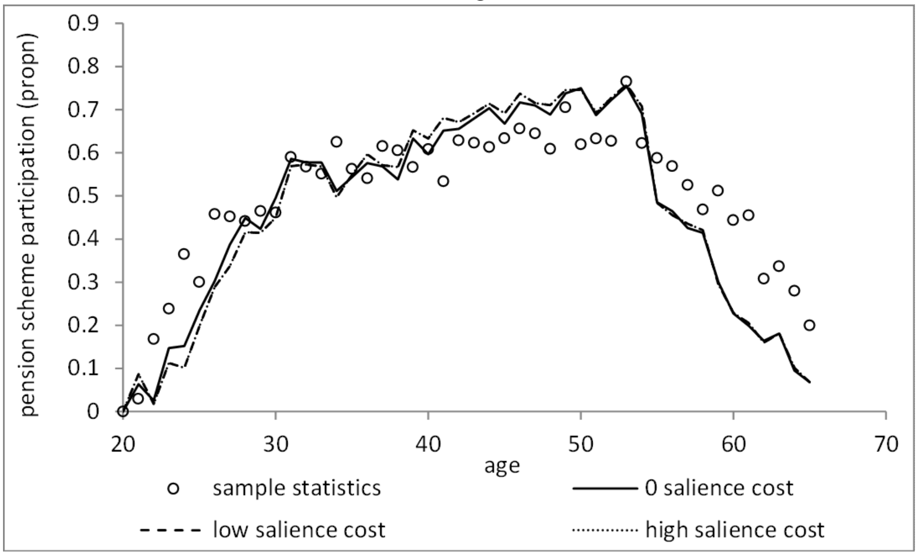

Panel B: adult couples

Figure 7: Rates of pension scheme participation by age and relationship status; simulated and sample statistics for 2006