Institute for Fiscal Studies

Retirement sorted? The adequacy and optimality of wealth among the near-retired

IFS Working Paper W14/23

Rowena Crawford
Cormac O'Dea

JRF JOSEPH ROWNTREE FOUNDATION

E·S·R·C ECONOMIC & SOCIAL RESEARCH COUNCIL

Retirement sorted? The adequacy and optimality of wealth among the near-retired

ROWENA CRAWFORD† AND CORMAC O'DEA§

†*Institute for Fiscal Studies*

§*Institute for Fiscal Studies and University College London*

Abstract

Much of the focus of the UK pensions policy debate over the past decade has been on the adequacy (or otherwise) of private retirement saving. In this paper, we present the first assessment of the *optimality* of the retirement resources of English couple households born in the 1940s. Here, 'optimal' wealth holdings are those that allow households to enjoy the same level of living standards in both working life and retirement. We use a life-cycle model of consumption and saving to calculate this level of wealth, and compare that with how much wealth households are observed to hold. We find that the majority of households hold more wealth than our model suggests is optimal and that this would still be true even if housing wealth were excluded from observed wealth holdings. A comparison of this approach with the replacement rate approach commonly used to assess the adequacy of households' retirement resources suggests that using a simple replacement rate benchmark could give a misleading picture of households' preparedness for retirement as it cannot capture the vast heterogeneity in households' circumstances.

Note: This paper contains a non-technical summary of much of the material contained in Crawford and O'Dea (2014).

The authors gratefully acknowledge funding from the Joseph Rowntree Foundation (project reference 1112004A) and the Economic and Social Research Council (ESRC) through both the Centre for the Microeconomic Analysis of Public Policy at IFS (grant reference RES-544-28-5001) and through the ESRC–NCRM Node 'Programme Evaluation for Policy Analysis' (PEPA) – reference ES/I02574X/1. The authors would like to thank Katy Heald and Gemma Tetlow for comments and suggestions. The views expressed and any errors are the authors' alone.

Data from the English Longitudinal Study of Ageing (ELSA) are made available through the UK Data Archive (UKDA). ELSA was developed by a team of researchers based at the National Centre for Social Research, University College London and the Institute for Fiscal Studies. The data were collected by the National Centre for Social Research. The funding is provided by the National Institute of Aging in the United States, and a consortium of UK government departments co-ordinated by the Office for National Statistics. National Insurance data for ELSA respondents were made available with permission from the ELSA Linked Data Access Committee.

Keywords: wealth; savings; life-cycle model

JEL classification numbers: D31, D91, E21, D12, H31, H55

1. Introduction

The UK pensions landscape has changed considerably over the past two decades. Life expectancies have been increasing rapidly and, in the absence of an equivalent increase in retirement ages, this has led to an increase in the resources individuals need in order to maintain their standard of living in retirement. At the same time, the availability of defined benefit (DB) pension schemes to private sector workers has declined, and government reforms have removed the earnings-related component of the UK state pension system. This means that, at a time when accumulating resources for retirement is more important than ever before, the onus has shifted more to individuals to be responsible for their own retirement saving.

It is against this backdrop that the focus of much policy debate in the UK over the past decade has been on the adequacy (or otherwise) of individuals' private pension saving. For example, the Department for Work and Pensions (DWP) estimated in 2012 that 11 million individuals in the UK were at risk of having an inadequate income in retirement (Department for Work and Pensions, 2012), and the DWP intends to continue using estimates of inadequacy as a way of assessing pension-related policy reforms (see, for example, Department for Work and Pensions (2013)).

There are two significant difficulties in assessing whether or not individuals are likely to have adequate resources for their retirement.

First, estimating the likely level and distribution of retirement resources is difficult. The introduction of household surveys that collect detailed information on wealth holdings, such as the English Longitudinal Study of Ageing and, more recently, the Wealth and Assets Survey, has considerably improved our knowledge of current levels of wealth. For households close to retirement, this measure gives a good indication of their likely retirement resources. However, for younger individuals, it is much harder to extrapolate their likely position at retirement without longer-run data on wealth holdings that give some indication as to the *trajectory* of their wealth in addition to its current level.

Second, even if individuals' future income and wealth holdings at retirement were known, assessing the adequacy of those wealth holdings poses a considerable challenge. It would be relatively simple to assess the 'adequacy' of resources in terms of whether they will be sufficient for an individual to avoid absolute poverty (or to avoid being reliant on the state benefit system to ensure their income was above some poverty threshold). However, assessing the 'adequacy' of resources in terms of whether they are sufficient to provide individuals with their desired standard of living is far more difficult, since it requires knowledge of individuals' desired standard of living.

To date, the UK literature that has addressed this second point has taken the approach of comparing individuals' retirement income with some standard benchmark income defined relative to current or previous levels of income or earnings (see, for example, Banks et al (2005), Crawford and O'Dea (2012), Department for Work and Pensions (2012) and Pensions Policy Institute (2012, 2013)). Retirement resources are then deemed to be ‘adequate’ if individuals are not likely to experience more than a certain decline in their income in retirement. This benchmarking approach is simple and transparent, but has the significant downside that the benchmark chosen is necessarily somewhat arbitrary.

An alternative approach is to use an economic model of consumption and saving over the life cycle to estimate for each household what their ‘optimal’ level of wealth accumulation is. In such models, the ‘optimal’ level of wealth will typically be that which facilitates the same standard of living in retirement as during working life. The key advantage of this approach is that one does not need to choose an arbitrary benchmark for what constitutes ‘adequate’ resources – instead, the model can be used to simulate how much each household will *want* to have saved for retirement given their individual circumstances (such as the number of children they have, their expectations over their lifetime earnings, the timing of those earnings and their entitlements to the state pension). However, the disadvantage of this approach is that, in order to construct the model, a number of assumptions need to be made about individuals’ expectations and preferences.

It is important to point out that the concepts of ‘optimality’ and ‘adequacy’ are subtly different; how they overlap will depend on how ‘adequate’ is defined and on the assumptions made in the model used to simulate optimal wealth. However, generally speaking, it is possible for an individual to have a level of wealth that is adequate but less than optimal (for example, if they replace the benchmark proportion of current or previous income but the model suggests that, given their circumstances, they should have saved more) or to have a level of wealth that is optimal but not adequate (for example, if they are a very low earner for their entire lives, they may not have the capacity to save much for retirement even if that results in their having a low income when they get there). Similarly, it is important to acknowledge that if all households saved optimally, this would not imply the absence of wealth (or retirement income) inequality and it would not necessarily imply the absence of poverty in retirement. Poverty could still exist, but would arise from low lifetime incomes rather than from the household undersaving.

In this paper, we calculate the first assessment of the *optimality* of the retirement resources of English households, focusing on couples born in the 1940s. Since these households are ‘near-retired’ when observed in 2002–03, we assume that their pension rights and wealth holdings at that point are indicative of the resources that will be available to them at the state pension age (SPA) and we therefore avoid the first problem described above of estimating decades’ worth of future saving behaviour. In order to place our results in the context of the existing UK literature on the adequacy of resources, we also consider the ‘adequacy’ of resources for the same sample of households using the common benchmarking approach. This builds on previous work (Banks et al. (2005) and Crawford and O’Dea (2012)),[1] although newly-available data mean we can make significant methodological improvements.

[1] Analyses in this spirit for the US have been carried out by Munnell et al. (2007, 2012).

We find that the majority of households hold more wealth than our model suggests is optimal and also look likely to have a replacement rate in retirement above the commonly-used thresholds for adequacy. We can also use the model (which predicts optimal wealth on retirement and optimal income in retirement) to calculate an implied optimal replacement rate. We find a wide distribution of optimal replacement rates, and that for many households the commonly-used thresholds for adequacy are a more demanding test than optimality. This suggests that considerable care should be taken when using simple benchmarks to assess the adequacy of households' resources, since they cannot capture the wide variety of household circumstances and therefore may result in a misleading impression of how prepared households are for retirement.

This paper proceeds as follows. Sections 2 and 3 discuss the methodology used to assess the adequacy and optimality of retirement resources respectively. The life-cycle model is only given a brief (and non-technical) treatment here – a full description is given in Crawford and O'Dea (2014). Section 4 describes the data we use and Section 5 our sample. Section 6 presents our results on the adequacy of resources, Section 7 presents our results on the optimality of resources and Section 8 compares the two approaches. Section 9 discusses the implications for the policy debate and concludes.

2. Method for assessing adequacy

We assess the adequacy of retirement resources using the 'benchmarking' approach, where we compare estimated retirement income with some standard replacement rate benchmarks. This has been the approach taken by almost all of the literature in the UK to date but, due to the availability of a new data set (discussed later), we are able to make significant methodological improvements over many existing studies that use this approach (including our own previous work).

2.1. Defining a replacement rate

Defining the benchmark against which retirement income is assessed in terms of a replacement rate (rather than an absolute level) means we are attempting to assess the adequacy of an individual's resources for maintaining their standard of living in retirement (rather than for avoiding poverty)

While the concept of a replacement rate is intuitive – broadly speaking, it is retirement income divided by pre-retirement income – in practice the exact definition of the replacement rate to be used requires careful consideration of a number of conceptual issues. In addition, data constraints often mean that the preferred definition of replacement rate may not be calculable in practice. An excellent discussion of the issues and choices involved can be found in MacDonald and Moore (2011). We will not repeat that discussion in full here, but limit ourselves to summarising some of the key decisions made in defining the replacement rate used in our analysis. These are: how to define income pre-retirement; how to define income post-retirement; over which age range pre-retirement income and post-retirement income should be measured; how to compare income in different periods; and how to take into account household composition (which changes over the life cycle).

What is pre-retirement 'income'?

The most inclusive measure of the total resources available to a household is total net income, and therefore conceptually we are most interested in the replacement of that by an equivalently inclusive measure of income in retirement. However, the data that we use in this paper (discussed in detail in Section 4) contain information on gross earnings, and therefore we necessarily focus on the replacement of pre-retirement gross earnings.

What is post-retirement 'income'?

We consider three definitions of income in retirement: gross pension income (the sum of state and private pensions), gross pension income plus the annuitised value of non-housing wealth, and gross pension income plus the annuitised value of total wealth. These last two measures of income allow us to take into account that many older households in England have wealth holdings that they could use to fund consumption in retirement in the absence of (or in addition to) specific pension resources. We do not take into account means-tested or other non-pension benefits that households may be entitled to after retirement so that our post-retirement definition of income is as comparable as possible to our pre-retirement definition of income (gross earnings).

Over what period is 'pre-retirement' income measured?

The choice of what period pre-retirement income is measured over reflects the period of a household's life with which one wants to compare standards of living in retirement. We focus on the pre-retirement 'lifetime', defined as between the ages of 20 and 50 inclusive, therefore choosing to compare standards of living in retirement with the mean over most of the working-age life. We include years in this age interval spent in unemployment (when earnings are zero) in the calculation of this average, so periods in unemployment will depress pre-retirement income. We truncate working-age life at age 50 for two reasons. First, all individuals in our data are observed until at least that age, and therefore we avoid drawing on assumptions about earnings in subsequent years. Second, we do not want years spent in early retirement (when earnings are zero) to depress our measure of average pre-retirement income.

Over what period is 'post-retirement' income measured?

The most comparable approach to our definition of pre-retirement income would be to measure post-retirement income as the average of income in all years between retirement and death. However, for simplicity, we focus only on income in a given year of retirement: at age 65.[2] The extent to which this will differ from the average will depend on the indexation arrangements of the relevant income streams and on our choice of how to compare income in different periods (see below). In England, state pension income is likely to be somewhere between price- and earnings-indexed (the basic state pension is earnings-indexed, while additional pension income from the state earnings-related pension and the state second pension is price-indexed) and private pension incomes are likely to be less than price-indexed on average. Therefore whether focusing only on income at age 65 understates or overstates the replacement rate relative to using average retirement income will vary between households according to the mix of public and private pension assets that they hold.

How is income in different periods compared?

A given (nominal) income level can purchase a different quantity of goods and services in different years, as prices typically rise year-on-year. Also, a given income level will place the recipient at a different point in the income distribution in different years, as earnings typically also rise each year (and, on average, at a rate faster than prices). It is necessary to revalue incomes in all periods in order to make them comparable. The obvious candidates to revalue incomes with are prices and earnings.

The choice of revaluation factor is not merely a technical matter but has implications for what the replacement rate is measuring. Uprating past income using inflation (i.e. uprating in line with prices) would mean that the replacement rate is measuring the replacement of the purchasing power enjoyed during working life. On the other hand, uprating past income in line with average earnings would mean that the replacement rate is measuring the individual's ability to maintain their living standards relative to the current working population. Since, on average, earnings tend to increase faster than prices, individuals need to save more and have a higher nominal income in retirement to maintain their position in the earnings distribution than to maintain their purchasing power.

For the analysis presented in this paper, we assume that individuals want a given level of replacement of their average lifetime consumption bundle, and therefore we convert income/earnings in all periods to 2002–03 prices using a price index. However, this is a crucial assumption, and so the sensitivity of our main findings to this are discussed in Section 6.4.

How is household size taken into account?

Household size is important when thinking about income replacement, since households change in size over time and larger households need a higher level of income for the same amount of consumption per person. This means that one reason why income needs might be lower in retirement than during working life could be that a household has dependent children during

[2] We refer to this throughout as the 'age at which the household reaches age 65'. To be more accurate, it is the year in which the man in the household reaches age 65. (Recall that we focus on couples.)

working life who no longer need to be financed by the time the parents reach retirement.

To take account of this, we 'equivalise' income for household size. For retirement income, we do this by dividing income in the year of interest by 1.5 if the household is a couple (and by 1 if it is a single individual). This is based on the modified OECD equivalence scale; the divisor is less than 2 for couples on the assumption that there are economies of scale in providing consumption within a household. We adjust pre-retirement income by dividing average lifetime earnings by an average equivalence factor:

$$S_{\bar{t}} = 1 + 0.5\ [\text{if have partner}] + \left(0.3 * \frac{14}{45} + 0.5 * \frac{5}{45}\right) * \text{Number of children} \qquad [1]$$

This is again based on the modified OECD equivalence scale (which weights additional adults and children aged between 14 and 18 by 0.5, and children aged between 0 and 13 by 0.3[3]) and it assumes that all children are born and reach age 18 while the household is aged between 20 and 65 (the male state pension age) and then leave home at 19. Since some of working-life income is saved to fund retirement, the average equivalence scale is arguably too large, as that portion of pre-retirement income that is saved for retirement should really be equivalised by the smaller, post-retirement, equivalence scale. If we took this into account, our equivalised pre-retirement income would be larger and our reported replacement rates would be lower. Because the magnitude of any errors will be small,[4] and to avoid making an arbitrary, ad hoc, adjustment, we proceed and use this equivalence scale.[5]

Summary

In summary, the replacement rate we focus on can be described as the ratio between household income in the year in which the household reaches age 65 and average lifetime earnings of the household, both adjusted for household size. More formally, this is defined as

$$\text{Replacement rate} = \frac{\delta_T (Y_T^A + Y_T^B)/S_T}{\frac{1}{S_{\bar{t}}}\left(\frac{1}{31}\sum_{t=\text{DOB}+20}^{\text{DOB}+50} \delta_t E_t^A + \frac{1}{31}\sum_{t=\text{DOB}+20}^{\text{DOB}+50} \delta_t E_t^B\right)} \qquad [2]$$

where

- t indexes a particular year in working life and T indexes the year of retirement;

[3] These age bands explain the two fractions in the equation. Every child spends 14 years (of its parents 45 in working life) getting a value of 0.3 in the equivalence scale and spends five years getting a value of 0.5 in the equivalence scale.

[4] For a couple with two children who save 15% of their income for retirement, the 'corrected' equivalence scale would be approximately 2.5% smaller than the one we use and therefore our reported replacement rates would need to be reduced by 2.5%.

[5] The model presented in the next section is able to take this fact into account in a more satisfactory manner.

- δ_t is the revaluation factor that converts income/earnings received in year *t* into 2002–03 terms using the change in prices between *t* and 2002–03;
- $(Y_T^A + Y_T^B)$ is household gross pension income (or pension income plus the annuitised value of wealth) in the year the man reaches age 65;
- S_T is the size of the household in the retirement year of interest;
- $S_{\bar{t}}$ is the average size of the household during working life;
- $\frac{1}{31}\sum_{t=DOB+20}^{DOB+50} \delta_t E_t^A$ is mean (real) gross annual earnings of individual A between the ages of 20 and 50 (inclusive);
- $\frac{1}{31}\sum_{t=DOB+20}^{DOB+50} \delta_t E_t^B$ is mean (real) gross annual earnings of individual B between the ages of 20 and 50 (inclusive).

2.2. Choosing an appropriate benchmark

Having defined our replacement rate, the important question is what level of replacement would typically constitute an 'adequate' retirement income.

It is not obvious whether expenditure needs increase or decrease in retirement. Some expenditure needs – such as work costs – will cease in retirement and retired individuals have more time to shop around and so can often achieve a given level of consumption at lower cost. On the other hand, increased leisure time may increase desired expenditure or declining health may increase healthcare-related expenditure.

Even if expenditure needed to be perfectly maintained in retirement for a household to have an adequate standard of living, in general this can be achieved with less than 100% replacement of gross income. There are two main reasons for this. First, retired households have much lower saving needs than working-age households (who are largely saving to accumulate housing wealth and pension resources for retirement). Second, households over their SPA in the UK on average face lower tax rates for three reasons. First, working-age individuals pay National Insurance contributions (NICs) on earnings, whereas individuals aged over their SPA do not pay NICs. Second, as incomes tend to fall on retirement and the income tax schedule is progressive, pensioners typically face lower average tax rates than they did in working life. Third, during the period we consider, the UK tax system gave a higher tax-free allowance to those aged 65 or over, so pensioners faced lower average tax rates than working-age individuals with the same level of income. Given these factors, it is likely to be the case that replacement rates of gross income of less than 100% will result in incomes net of saving and tax that are comparable to those experienced during working life.

However, it is not clear (without appealing to an economic model of the type we discuss later) how much less than 100% replacement rates can be while maintaining living standards. The existing literature that has defined a benchmark replacement rate for adequacy has therefore typically taken one of two main approaches to make such a choice seem less arbitrary.

The first has been to refer to economic studies that investigate actual income in retirement and how this compares with income pre-retirement (for example, Blundell and Tanner (1999), Bardasi et al. (2002) and Crawford and Tetlow (2012)). These studies find net replacement rates in the UK that average around 80%, being higher (lower) among those with lower (higher) pre-retirement incomes. However, there are two drawbacks to such an approach. First, different answers may be obtained with different definitions of replacement rate (for example, the UK's progressive personal tax system means that replacement of net income is likely to be greater than replacement of gross income for any given individual), and therefore care must be used when choosing an appropriate benchmark based on empirical evidence that uses a different definition. Second, and more importantly, this approach presupposes that these retired individuals have incomes that are exactly 'adequate', when in fact arguments could be made that these observed incomes are higher or lower than these individuals actually desired.

The second approach has therefore been to try to elicit from individuals what they would deem to be adequate (for example, Mayhew (2002), Binswanger and Schunk (2012) and Aegon (2013)). Most individuals seem to think that they do not need to replace all their income on retirement: Mayhew (2002) found that 71% of individuals suggested an adequate retirement income would be less than or equal to their current incomes (10% did not know) while Aegon (2013) found that 93% of individuals reported their required replacement rate to be 100% or less.[6]

There is little strong evidence in the UK that argues for the use of a particular replacement rate benchmark. The most commonly-used benchmark in the UK literature over the past decade is that put forward by the Pensions Commission in 2004. This benchmark consists of a set of replacement rate thresholds for individual gross earnings that depend on an individual's pre-retirement earnings, as illustrated in Table 2.1.

TABLE 2.1

Pensions Commission (2004) benchmark of adequacy

Pre-retirement gross earnings	Replacement rate threshold
Less than £9,500	80%
£9,500 to £17,499	70%
£17,500 to £24,999	67%
£25,000 to £39,999	60%
£40,000 or more	50%

Source: Pensions Commission, 2004, appendix G (The Pensions Commission's 'Group Modelling'), table G.1.

In this paper, we investigate the proportion of households likely to be at risk of inadequate resources when adequacy is defined along the lines of the Pensions Commission benchmark. However, instead of allocating

[6] Aegon (2013) found that 16% of UK respondents expected to need a gross income in retirement of less than 40% of current earnings, 30% expected to need 40–59%, 33% expected to need 60–79% and 15% expected to need 80–100% of current earnings.

individuals a target replacement rate on the basis of their own pre-retirement gross earnings (as the Pensions Commission analysis did), we allocate individuals a target replacement rate on the basis of the average equivalised gross lifetime earnings of their household (defined in Section 2.1), though still using the same earnings levels and thresholds set out in Table 2.1. This benchmark is not without its own drawbacks, however (as discussed in Crawford and O'Dea (2012)), and therefore we also investigate the proportion of households likely to be at risk of inadequate resources when adequacy is defined according to two simple illustrative benchmarks: 67% replacement of gross lifetime household earnings and 80% replacement of gross lifetime household earnings.

3. Method for assessing optimality

We assess the optimality of household wealth using an economic model of consumption and saving. This section provides a non-technical summary of the model and a discussion of how the approach differs from the simple replacement rate benchmarking method for assessing adequacy described in Section 2. A comprehensive and more technical description of the model can be found in Crawford and O'Dea (2014).

3.1. Overview of our model

Our model considers household behaviour between age 20 and death. The basic structure of the life-cycle model is summarised in Figure 3.1. Each year, households in the model choose how much of their income in that year to spend and how much to save either in a safe asset or in a pension fund. In making this choice, households aim to choose their consumption and saving allocations so that they can keep their living standards as constant as possible across their life cycle. The 'world' that is modelled is risky – household members do not know when they will die, whether they will be in employment at each age in the future, what their earnings will be if they are in employment and what the return will be on their pension fund investments. Their saving choices will need to ensure that they have enough wealth both to buffer themselves against bad luck and to fund their retirement without an undesirable fall in consumption.

The level of wealth that is optimal is that which best (according to a standard criterion[7]) allows living standards to be kept constant across the lifecycle. Too much wealth is suboptimal as consumption opportunities during working life are being forgone; too little wealth is suboptimal as living standards will fall on retirement.

Households in the model can place their saving in one of two assets – a safe asset and a defined contribution (DC) pension. The safe asset offers a

[7] Formally, the criterion is the equality of expected discounted marginal utility over the life cycle. The interested reader is referred to Crawford and O'Dea (2014). The uninterested reader can safely ignore this footnote.

certain but relatively low rate of return (2.2%[8]), while the DC pension offers a higher expected return (4%) but that return is uncertain, and the DC pension is less flexible in that wealth cannot be drawn until age 65 and 75% must be used to purchase an annuity.[9] There are no defined benefit (DB) pensions in our model, despite these being an important feature of the UK pensions landscape. This affects the interpretation of the results of our modelling. The model essentially asks the question 'If individuals had access to a safe asset and a DC pension, how much private wealth would they want to accumulate in these assets?'. There is also no explicit modelling of housing choices in our baseline model – housing investments should be considered as part of the safe asset investments.[10] We devote some time to discussing the role that housing may play in our results in Section 7.

FIGURE 3.1

Summary of the life-cycle model

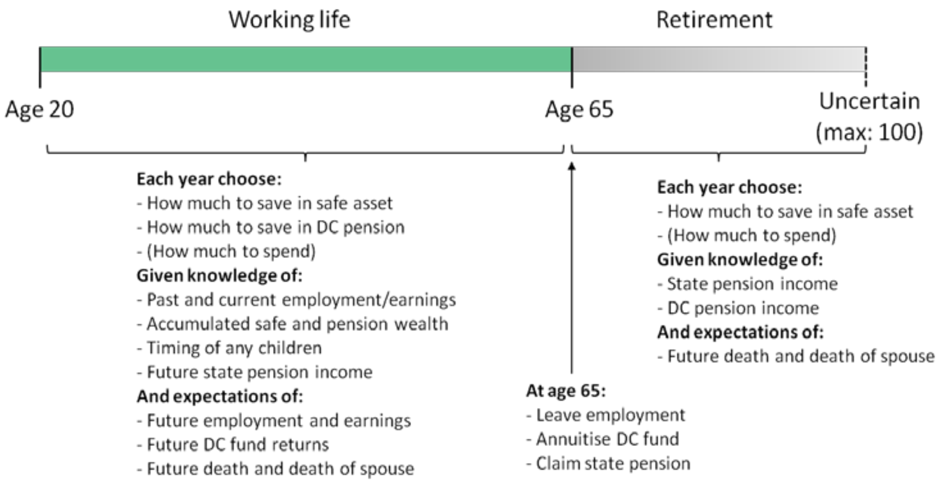

[8] This is the midpoint between the average real return on cash over the 50 years to 2012 (1.6% – see Barclays Capital (2012)) and the average real increase in house prices since 1974 (2.8%, calculated using the Nationwide price index).

[9] The compulsory annuitisation of at least 75% of DC pension funds was chosen to reflect existing UK legislation that imposes punitive tax rates on those who wish to withdraw more than 25% of their fund. However, in the March 2014 Budget, the government announced that from April 2015 it would change the tax rules so that individuals could access their DC pensions as they wish after a certain age (currently 55).

[10] The fact that we do not model housing wealth explicitly is an unfortunate but necessary modelling simplification. Our inclusion of a second asset (the pension fund) marks a substantial innovation relative to the literature that has examined this particular issue previously (primarily, Scholz et al. (2006)). The results reported in Crawford and O'Dea (2014) and referenced here take approximately three days to be generated, even when the program that generates them is written in an extremely efficient programming language and using a state-of-the-art computer. Adding a third asset would imply that the result would take a number of months to produce.

The policy environment in which households make these decisions is based on the UK tax and benefit system. Households must pay income tax on their income, and NICs on their earnings, in each year according to the parameters of the UK tax system in 2002–03.[11] When households annuitise their DC pension, they can take 25% of the fund value as a tax-free lump sum. Our model includes a flat-rate unemployment benefit, and a child benefit that is paid at a higher rate for the first child than for subsequent children. Households in our model are also entitled to a state pension from the age of 65. Households are assumed to know the level of this future state pension income from the age of 20. In other words, we are using our structural model to answer the question 'Conditional on their future state pension income, what is the optimal amount of private saving (in a DC pension and a safe asset)?'.

Structural models of this type are computationally intensive to solve for households' optimal behaviour[12] and therefore a number of simplifications have to be made. The main simplifications are described below.

Household composition

The model only considers couples (though in principle it would be possible to solve a similar model for single households). The only changes in household composition we allow for in the model are the arrival of children (the timing of which is assumed to be known by the household from the start of working life) and the death of each adult in the household (which we assume is uncertain and related only to age and sex). We assume that all couples have been partnered since the age of 20 and will stay together until death (i.e. we do not model household formation or separation).

Labour market activity

We model employment and earnings at the household level. Employment therefore represents whether either individual in the couple is in work, and changes in employment of a second earner in the couple would appear as changes in the level of household earnings. We assume all households enter the labour market at age 20 and leave at age 65. Conditional on being in the labour market, household unemployment happens (i.e. neither spouse receives any earnings) with a known probability (6.2%[13]).

[11] We use this tax system as it is the one that households faced in the year that we observe their wealth levels in our data. The model does not consider a time-varying tax system – households are assumed to have faced this tax system for their entire working life.

[12] Indeed, the true solution to the 'optimal' level of wealth given by this model is impossible to determine exactly without access to an infinite quantity of computing power – and such computing power does not, of course, exist. Therefore, the results in Crawford and O'Dea (2014), like those of all similar papers, give approximate solutions to the model that the paper presents.

[13] This is the incidence in our data of a household aged between 25 and 50 having total earnings of less than £4,402 – the level of unemployment benefit payable to an unemployed couple in 2002–03.

Employer contributions

Pension wealth observed in the data includes funds accumulated as a result of contributions made on an individual's behalf by their employer. The measure of past earnings that we include in our model includes an estimate of the value of employer pension contributions made in each year – that is, our measure of earnings captures not only current pay but also deferred remuneration received in the form of future pension promises. This allows us to use the model to ask 'If a household had free choice over what to do with its total remuneration (i.e. it received additional earnings instead of pension contributions), how much would be saved optimally?'. If we find that individuals have 'oversaved' for retirement, one explanation could be that they did not in fact have a free choice about when they consumed their earnings but instead were forced to save too much in a private pension by their employer making larger contributions than they desired (and could not offset this by reducing their own pension saving or other wealth accumulation).

There is one feature of employer pension contributions that we do not capture in our model. In reality, employer pension contributions may affect not just a household's total level of remuneration, but also their incentives to save. For example, employers may provide incentives to households to make contributions to a DC pension, such as matching employee contributions with employer contributions. We do not attempt to include such features in our model, but to the extent that they increase the return to pension saving, our model is likely to understate optimal pension wealth accumulation (though this does not necessarily mean that the model would understate optimal *total* wealth accumulation, as additional pension saving could be at the expense of non-pension saving).

Annuitisation behaviour

We assume that all households annuitise their accumulated DC pension fund at age 65. Households are assumed to purchase an annuity whose purchasing power does not change over time (i.e. a pension income stream that increases in line with inflation). If both members of the couple are alive by the time the man reaches age 65, the household is assumed to purchase a joint life annuity (that is, one that continues to pay out 50% of the original annuity value once the first member of the couple dies, until the second member of the couple also dies). If only one member of the couple is still alive at the point that the annuity is purchased, it is assumed that a single life annuity is bought. Annuity rates are assumed to be actuarially fair (given the average life expectancy of the cohort we consider), after 10% has been deducted to meet the administrative costs (including profits) associated with the provision of annuities. We take this estimate from Murthi et al. (1999), who apply the methodology of Mitchell et al. (1999) to the UK.

Defined benefit pensions

There are no DB pensions in our model, despite these being an important feature of the UK pensions landscape. This affects the interpretation of the results of our modelling. The model essentially asks the question 'If individuals had access to a safe asset and a DC pension, how much private wealth would they want to accumulate in these assets?'. One reason why observed wealth may differ from what our model suggests is optimal could therefore be that individuals had access to an alternative savings vehicle (such as a DB pension) that affected how much they wanted to save for retirement.

3.2. Using the model to assess optimality

The model produces estimates for a measure of optimal wealth holdings each year between age 20 and death (comprised of optimal safe asset holdings and optimal DC pension wealth). To assess the optimality of households' actual wealth holdings, we can then compare the level of wealth they report in one of our data sets – the English Longitudinal Study of Ageing (ELSA) in 2002–03 – with the optimal level of wealth that our model simulates they should have at that age.

3.3. Differences from a simple replacement rate approach

Despite the necessary simplifications described above, the modelling approach of estimating optimal wealth has a number of advantages over the simple replacement rate approach of assessing adequacy that we described in Section 2.

First, the model explicitly takes account of households' needs for precautionary saving and saving for retirement during working life. Optimal wealth is simulated taking into account that income saved during working life for these purposes does not need to be replaced in retirement.

Second, the model takes account of when earnings occur in the life cycle. Two households with the same average lifetime earnings can have very different earnings profiles – for example, one household could have received all their earnings in the first half of working life, while the second could have received all their earnings in the second half of working life. The replacement rate approach would hold these two households against the same benchmark, while the model would likely simulate that optimal wealth would be higher for the household that received their earnings in the first half of working life since they have more time to accumulate a real return on any earnings they save.

Third, the model is better able to take account of changes in household composition over time by adjusting preferences over consumption each period to take into account family composition at that time. This is a significant improvement over the average equivalisation factor applied to pre-retirement earnings in the replacement rate approach.

3.4. Additional factors that may influence wealth accumulation

There are a number of factors that are not included in our structural model that might be important in determining how much wealth households optimally want to accumulate throughout their lives. A number of these, and the way in which they could influence our results, are discussed below.

Bequests

Our model does not include an explicit bequest motive. To the extent that households derive well-being from leaving wealth to their heirs when they die, our model will likely understate optimal wealth holdings. However, from the perspective of the design and evaluation of policy, we might be more concerned with the ability to achieve a desired standard of living than with the ability to satisfy bequest motives.

Housing

In our model, there are only two assets: a DC pension and a risk-free asset. In reality, however, housing wealth is a very important component of wealth portfolios in the UK. Housing wealth differs from other assets in a number of ways. First, it has a consumption value. In other words, we tend to think a household that owns and lives in a £200,000 house has greater well-being (all else equal) than a household that rents accommodation and has £200,000 in a savings account – in large part, because it does not incur any rental costs, but also because the household may derive well-being from owner-occupation in itself (for example, being able to decorate as desired). Second, house price growth has been substantial (in a perhaps unexpected manner) for most of those in our sample who purchased a property. Third, housing assets are purchased by households in a different manner from most other assets – with a mortgage. This has the effect of magnifying the gains (and losses) on the initial investment. By not taking into account these benefits associated with housing wealth, our model may not accurately determine optimal wealth accumulation. We test the sensitivity of our main findings to the role of housing wealth in Section 7.

Maintaining relative standards of living

Our model assumes that households' objective is to maintain their standards of living in retirement by smoothing risk-adjusted consumption throughout life. However, as discussed in Section 2, households may be concerned with maintaining their standards of living relative to other working-age individuals, rather than just maintaining their average consumption bundle (which, since average earnings typically grow faster than prices, implies a higher level of consumption). If this were the case, households would need to accumulate greater wealth in order to fund higher consumption in retirement, and our model would understate optimal wealth.

Changing price of consumption on retirement

Our model assumes that the price of consumption is constant throughout life. However, there is literature that suggests that in retirement a given level of consumption can be obtained at a lower price (see, for example, Aguiar and Hurst (2005)). Once households have stopped working, they have more time to shop around for better value and have the ability to substitute home production for the purchase of some goods and services. To the extent that this is the case, our model would overstate optimal wealth accumulation for retirement, although the effect is likely to be small.

Changing well-being from consumption on retirement

Our model assumes that the well-being generated from consumption is the same in working life as it is in retirement. However, retired individuals have considerably more leisure time than they do when they are working. If leisure and consumption are complements (so that an increase in leisure would increase the consumption needed to obtain a given level of well-being), then our model would understate optimal wealth accumulation. On the other hand, if leisure and consumption are substitutes (so that an increase in leisure would reduce the consumption needed to obtain a given level of well-being), then our model would overstate optimal wealth accumulation. Plausible arguments can be given for each case: an individual with cycling as a hobby may prefer to buy a much more expensive bike when he is retired and has more time to cycle – an example of leisure and consumption being complements; an individual who enjoys hiking may prefer to give up more expensive leisure pursuits when he has more time to dedicate to walking – an example of leisure and consumption being substitutes. However, the academic economics literature to date suggests that, overall, consumption and leisure are substitutes (see, for example, Browning and Meghir (1991)).

Precautionary saving for long-term care needs

In the US, one factor that is argued to explain much of household wealth accumulation is precautionary saving for medical costs at older ages (see, for example, DeNardi et al. (2010)). In the UK, this is much less likely to be an important driver in households' wealth accumulation decisions since the National Health Service (NHS) provides healthcare free at the point of consumption. However, long-term social care needs are not fully met by the state and the risk of these costs could affect household wealth accumulation. It is not immediately clear what the impact of incorporating such risks in the model would be though. In England, state support for social care costs is subject to an asset-based means test. Therefore, while the risk of needing social care might induce some households to increase wealth accumulation for precautionary reasons (optimal wealth would be higher), other households might actually have a reduced incentive to accumulate wealth so as to avoid the means test (optimal wealth would be lower).

4. Data

The considerable data requirements in simulating households' optimal wealth accumulation are met using a powerful and under-exploited new data set – survey data from the English Longitudinal Study of Ageing linked to administrative data from National Insurance (NI) records.

While the ELSA data have been used to consider the adequacy of retirement resources in previous work, the availability of the linked NI data enables us to make two significant methodological improvements over our own previous analysis. First, instead of estimating future state pension entitlement on the basis of reported work history and numerous assumptions about past and future behaviour, we can use the NI data to calculate each individual's actual state pension entitlement given their contribution history to date (these are, after all, the very same data that the government uses to calculate entitlements). Second, the NI data also enable us to consider the adequacy question with respect to replacement of lifetime earnings rather than simply earnings in one particular year. This has the advantage of being less volatile, and a much better indicator of the average standards of living that a household is likely to want to maintain in retirement.

The particular data requirements for the two strands of analysis are described in Table 4.1, along with the data source we use. In the following subsections, we describe the ELSA data and NI data in more detail.

TABLE 4.1

Data requirements and sources

Analysis	Data required	Source
Adequacy	Average lifetime earnings	Estimated from NI data
	Number of children	Reported in ELSA
	Future state pension income	Calculated from NI data
	Future private pension income	Estimated from ELSA
	Future income from annuitised wealth	Estimated from ELSA
Optimality	Employment and earnings each year	Estimated from NI data
	Employer pension contributions each year	Estimated from ELSA & NI data
	Number and timing of children	Reported in ELSA
	Average DC fund return	Estimated from FTSE DC-isions and FTSE all-share indices
	Non-pension wealth	Reported in ELSA
	Pension wealth	Estimated from ELSA

4.1. The English Longitudinal Study of Ageing

The English Longitudinal Study of Ageing is a biennial longitudinal data set, broadly representative of the household population of England aged 50 and over. It began in 2002–03 with a sample of around 12,000 individuals, and to date there are five subsequent 'waves' of data available (collected in 2004–05, 2006–07, 2008–09, 2010–11 and 2012–13).

Respondents to the ELSA survey are also asked for their permission to access their NI records. For those individuals who gave permission in wave 1 (2002–03), their linked NI records up to 2003–04 are available. The analysis in this paper is therefore based on the data from the first wave of ELSA (2002–03).[14]

[14] While we do have additional information from subsequent waves of ELSA for these same respondents, we do not use these data because attrition from the survey reduces the sample size over time, and if attrition were non-random it could bias our results.

The ELSA survey collects a large amount of data on demographics, labour market circumstances, financial circumstances, subjective and objective measures of health, and individuals' expectations about various future events. Crucially for our purposes, the survey asks respondents about their income and wealth holdings, and there is an extensive set of questions asked about any private pension arrangements that respondents might have, which enables us to estimate future private pension income in retirement.[15]

The definitions of pension income and wealth that are used in the analysis of adequacy are described in Table 4.2. All financial variables are described at the household level. The year-of-retirement income that we are interested in for our definition of the replacement rate lies in the future for all of the households in our sample – we therefore need to estimate future wealth and future pension income on the basis of what is reported in ELSA in 2002–03. We estimate pension income given each individual's reported scheme rules and contribution/entitlement history. For non-pension wealth, we simply assume that all wealth grows by 2.2% per year in real terms between 2002–03 and the year of interest.[16] For state pension income, we calculate individuals' entitlement given their NI contribution history up to 2002–03, and then estimate how that entitlement would increase assuming those in work in 2002–03 continue in employment until the SPA.

TABLE 4.2

Definitions of income and wealth

Measure	Comprised of:
Pension income	State pension income (basic state pension plus state earnings-related pension / state second pension)
	Private pension income (employer pensions plus personal pensions)
Total wealth	Non-housing wealth
	Net housing wealth
Non-housing wealth	Net financial wealth (interest-bearing accounts at banks and building societies (including ISAs and TESSAs); National Savings accounts and personal equity plans; stocks and shares; government, corporate and local authority bonds; investment trusts and unit trusts; less outstanding loans and non-mortgage debts)
	Net physical wealth (net non-owner-occupied housing wealth; property and land; antiques and collectables; covenants and trusts; net business wealth)
Net housing wealth	House value (principal residence)
	Less outstanding mortgage debt

In our analysis of optimality, we compare the simulation from our model with observed wealth holdings in 2002–03. That observed wealth comprises both non-pension wealth and pension wealth. Non-pension wealth is as

[15] See Crawford (2012) for a full description of the estimation methodology.

[16] This is the same as the interest rate assumed for the risk-free (non-pension) asset in the model for assessing optimality. It is the midpoint between the average real return on cash over the 50 years to 2012 and the average real increase in house prices since 1974.

described in Table 4.2 (the sum of non-housing wealth and net housing wealth). Pension wealth is taken to be the accumulated fund value in the case of DC pensions, and the sum of the discounted stream of future pension income in the case of DB pensions.

4.2. National Insurance data

The NI data are the administrative record of individuals' NICs and form the data set that is used by the Department for Work and Pensions to establish individuals' rights to claim contributory benefits such as the state pension.

For each year since 1975, the NI data record the level of earnings of employed individuals (though for the period before 1997 the data are top-coded at the upper earnings limit (UEL)). For the period between 1948 and 1975, the NI data record the number of weeks that an individual earned above the lower earnings limit (LEL).[17] For self-employed individuals, the NI data only ever contain information on the number of weeks of self-employment that an individual has done in a given year and no measure of self-employment income.[18]

While the NI data provide extremely detailed and accurate information on the earnings histories of most ELSA respondents, there is still some estimation that needs to be done to construct full lifetime earnings histories.

There are two limitations that need to be overcome in order to construct full lifetime earnings histories from the NI data: the top-coding of earnings at the UEL between 1975 and 1996, and the lack of data on the level of earnings prior to 1975. To estimate the level of earnings for those who earn more than the UEL, we use a regression technique (a fixed effects tobit). The interested reader is referred to Crawford and O'Dea (2014) for further details. To simulate earnings before 1975, we follow broadly the methodology of Bozio et al. (2011): we calculate mean earnings for each individual over the years 1975 to 2004 in which they were observed working, and then estimate potential previous years' earnings by adjusting for average economy-wide earnings growth and individual-level earnings growth given their age, sex and education level. The NI data record how many weeks the individual made NI contributions between 1948 and 1975. For men we assume they worked those weeks immediately prior to 1975 (therefore any periods not working were at the start of working life), while for women we assume that they worked those weeks from the point of leaving full-time education (therefore any periods not working were immediately prior to 1975). The combination of the estimates of potential

[17] Prior to the introduction of the State Earnings-Related Pension Scheme in 1978, entitlement to the state pension depended only on the number of weeks' worth of flat-rate NICs that had been made and therefore that is all that is recorded in the NI data for that period.

[18] This is because the self-employed accrue entitlement to benefits through their payment of flat-rate (class 2) NICs.

earnings in a particular year for each individual and the years in which they were working yields our earnings estimates for years prior to 1975.

The data tell us which years household members were in self-employment but do not allow us to estimate their self-employment income. Our working-life earnings measure does not, therefore, include self-employment income, though our main results are robust to excluding households where either spouse is recorded as being in self-employment in more than five years.[19]

5. Sample

In this paper, we investigate the wealth holdings of couples where the man was born in the 1940s. These households are 'near retirement' when observed in the 2002–03 ELSA data, being aged between 52 and 63 (the state pension age for this cohort is 60 for women and 65 for men), and therefore we only need to make limited assumptions about future work and pension saving behaviour.

There are 1,615 couples in ELSA where the man was born in the 1940s.[20] We focus only on couples for two reasons: first, for technical reasons, our structural model is currently designed to model only the behaviour of couples; second, the majority (76%) of individuals in this cohort observed to be single in 2002–03 were either widowed, divorced or separated. Since they previously had a partner, this means that their own lifetime earnings in isolation will not be representative of their actual access to resources, and therefore comparing their current wealth levels and future pension rights with their own past earnings may give a misleading impression of their ability to replace their previous standards of living in retirement.

There are three other restrictions we make to our sample. First, we drop households who are missing any 'core' ELSA data (age, education or wealth). Second, we drop households where not all members granted access to their NI records or where not all individuals could be linked to their NI records (normally when the individual was unable to provide their NI number). Finally, we focus our analysis on households for which we observe at least five years of earnings. The reasons for the first two restrictions are fairly obvious. The last restriction is imposed because we need to input into our structural model an estimated earnings process for each household, and to estimate such a process we need repeated observations of earnings over time for each household.

The impact of these restrictions on our sample size is illustrated in Table A.1 in the appendix. Of potentially greater concern is the impact of these restrictions on the representativeness of our sample. However, analysis of the characteristics associated with being in the final sample suggests that our sample of couples is broadly representative on observables, with the

[19] Results are available from the authors on request.

[20] There are also 322 single men and 526 single women in ELSA who were born in the 1940s.

particular exception that it is under-representative of the self-employed.[21] This is consistent with the findings of Bozio et al. (2010), and is illustrated – along with some summary statistics on the characteristics of our sample – in Table A.2 in the appendix.

6. Results: the adequacy of retirement resources

In order to relate our analysis on the optimality of household saving to the existing literature on adequacy, we first present analysis of the adequacy of retirement resources, using the replacement rate approach described in Section 2. It should be kept in mind throughout that this analysis is conducted for our particular sample of households – couple households in which the man was born in the 1940s, where we have NI data for both individuals, and where the household has at least five years of positive earnings in the NI data. The results are therefore not representative of the whole 1940s cohort, and are not directly comparable to the results published in previous work such as Crawford and O'Dea (2012).

6.1. Income in retirement

Figure 6.1 starts by illustrating the distribution of estimated state and total pension income at age 65. Income from the state pension is relatively tightly distributed, since the rules of the UK state pension system mean that entitlement is only weakly related to contributions made (or earnings received). Two-thirds of households have gross state pension income of between £6,000 and £10,000 per year (in 2002−03 prices). Median state pension income is around £8,700 per year. There is much more variation, however, in how much households have saved in private pensions. Looking at total pension income, the distribution is much wider and median total pension wealth is also obviously higher, at around £15,600 per year.

Many households also have other forms of wealth that could be used to fund consumption in retirement. Figure 6.1 also illustrates how the distribution of retirement incomes would look were we to include as income the annuitised value of non-housing wealth or the annuitised value of total wealth. In each of these cases, household annual gross income at age 65 would clearly be much greater than pension income – median income would be around £18,600 were the annuitised value of non-housing wealth included and around £26,000 were the annuitised value of total wealth included.

Table 6.1 summarises the proportion of households that would have gross income at age 65 below various different levels.[22] Virtually all households have no more than £15,000 per year of state pension income, but the majority of households (54%) have more than that amount from state and

[21] As discussed above, the self-employed are somewhat problematic in our analysis in any case since the NI records do not contain data on their earnings during periods of self-employment.

[22] The cumulative distribution of gross income at age 65 under the different income definitions is illustrated in Figure A.1 in the appendix.

private pensions combined. Including the annuitised value of non-housing wealth, nearly half of households have more than £20,000 per year, while if we also include the annuitised value of housing wealth 41% of households would have more than £30,000 per year.

FIGURE 6.1

Distribution of estimated gross income in retirement

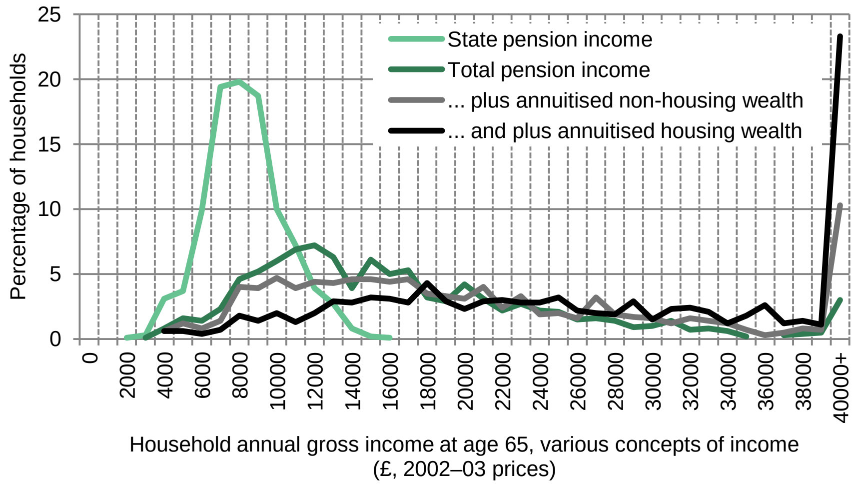

TABLE 6.1

Distribution of estimated gross household income in retirement

Percentage of households with annual gross income at age 65:	State pension income	Total pension income	... plus annuitised non-housing wealth	... and plus annuitised housing wealth
≤£5,000	3.5%	0.9%	0.6%	0.6%
≤£7,500	25.4%	4.8%	3.2%	1.8%
≤£10,000	75.0%	16.1%	11.9%	5.5%
≤£15,000	99.7%	46.5%	33.9%	16.6%
≤£20,000	100.0%	69.1%	54.4%	32.9%
≤£30,000	100.0%	91.1%	79.6%	59.0%
≤£40,000	100.0%	97.0%	89.7%	76.7%
>£40,000	0.0%	3.0%	10.3%	23.3%

If we were concerned about the adequacy of retirement incomes in terms of households having an income above a poverty line, we could investigate this simply by comparing the estimate of pension income with such a poverty line. For example, in 2002−03, the minimum income guarantee for pensioners (the state means-tested income floor for those aged over 65) was £149.80 per week for a couple, or £7,790 per year. Comparing our estimated incomes at age 65 with this threshold suggests that 5.6% of our households would have annual pension income of less than that amount and so would be reliant on means-tested benefits to avoid poverty.[23]

6.2. Replacement rates

We turn now to the replacement rates that are calculated according to equation 2 in Section 2. Figure 6.2 illustrates the distribution of replacement rates obtained by dividing retirement income by average lifetime earnings for each household. Table 6.2 summarises the proportion of households that would have replacement rates below certain levels.[24] For most households, state pension income at age 65 replaces between 20% and 80% of average lifetime earnings (only 16% of households have state pension income in excess of 80% of average lifetime earnings). Replacement by total pension income is higher, and for over 40% of households it is above 100%. Only 20% of households in our sample are estimated to have total pension income of less than 67% of average lifetime earnings and 35% to have total pension income of less than 80% of average lifetime earnings. The proportion of households falling below these common replacement rate benchmark thresholds would fall were the annuitised value of wealth to be included as income – we estimate that if all wealth were annuitised, only 2% of households would replace less than 67% of average lifetime earnings and 5% would replace less than 80% of average lifetime earnings.

TABLE 6.2

Distribution of estimated gross household income replacement in retirement

Percentage of households with:	State pension income	Total pension income	... plus annuitised non-housing wealth	... and plus annuitised housing wealth
<=40% replacement	28.5%	2.0%	0.9%	0.4%
<=50% replacement	46.5%	4.9%	2.5%	0.5%
<=67% replacement	72.7%	19.6%	10.0%	2.3%
<=80% replacement	84.3%	35.0%	19.9%	5.3%
<=100% replacement	93.5%	58.6%	41.0%	16.0%
>100% replacement	6.5%	41.4%	59.0%	84.0%
<= PC replacement	77.1%	17.9%	10.1%	2.7%

Note: Replacement rate is defined as the ratio between household income in the year in which the household reaches age 65 and average lifetime earnings of the household, adjusted for household size (as defined by equation 2 in Section 2).

Thinking about the Pensions Commission threshold for adequacy (where the benchmark replacement rate varies between 50% and 80% depending on

[23] The proportion of households at risk of poverty would be lower if income were to include the annuitised value of wealth, but since the means test for income support only includes non-housing wealth above a certain threshold, households are unlikely to have a financial incentive to annuitise their wealth to increase their incomes as that would simply reduce their benefit entitlement rather than increasing their total income.

[24] The cumulative distribution of replacement rates under the different income definitions is illustrated in Figure A.2 in the appendix.

pre-retirement earnings), we find that 77% of households would have inadequate incomes from their state pension income alone, and 18% would find their total pension income to be inadequate. If we were also to include annuitised wealth in income, the proportion with inadequate resources would fall to 10% if non-housing wealth were included and 3% if housing wealth were also included.

FIGURE 6.2

Distribution of estimated household gross income replacement in retirement

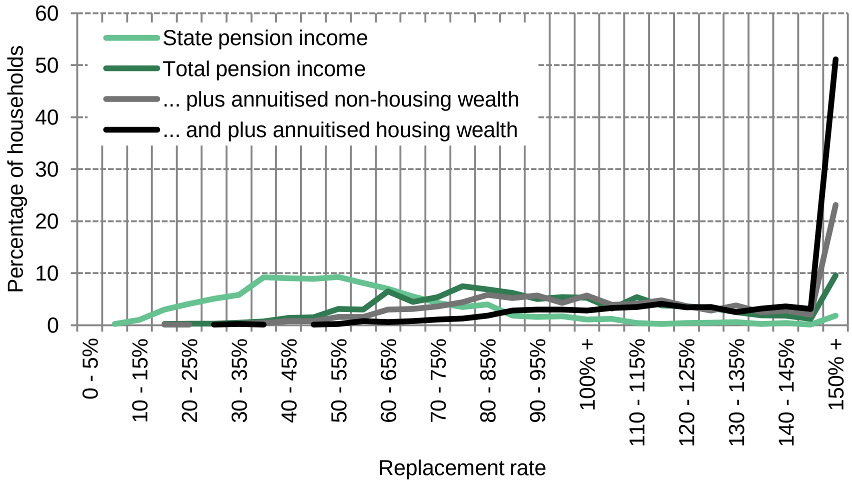

Note: Replacement rate is defined as the ratio between household income in the year in which the household reaches age 65 and average lifetime earnings of the household, adjusted for household size (as defined by equation 2 in Section 2).

6.3. Characteristics associated with adequacy

If some households are accumulating inadequate resources for retirement, it is important to understand why in order to make the best policy recommendations to address the undersaving problem. The first step in attempting to do that is to identify who is at risk of having inadequate resources.

Perhaps one of the most important characteristics is lifetime earnings. Households with higher lifetime earnings have a greater capacity to save than less well-off households, but they are also trying to replace a higher level of income in retirement and therefore need to save more in absolute terms. Table 6.3 shows how the proportion of households estimated to replace less than 67%, 80% and the Pensions Commission benchmark proportion of their average lifetime gross earnings varies by quartile of the average lifetime earnings distribution.[25] Focusing only on pension income in

[25] Quartiles of the lifetime earnings distribution are defined for the whole population of households in the cohort, including both single and couple households (earnings are equivalised). There are not equal numbers of couple households in each quartile because couples tend to be higher up the (equivalised) lifetime earnings distribution than single households. We exclude years of self-employment when

retirement, the proportion of households with inadequate resources is significantly higher among higher average lifetime earners than among lower average lifetime earners. For example, only 2.4% of our households who are in the lowest-earning quarter of the distribution are estimated to have pension income of less than 67% of their average lifetime earnings, compared with 29.8% of our households in the richest quarter of the distribution. That those with the highest incomes are more likely to have inadequate resources is even the case under the Pensions Commission benchmark for adequacy, where the replacement rate threshold is lower for those with higher earnings.

TABLE 6.3

Percentage of households with less than 67%/80%/Pensions Commission (PC) replacement of average lifetime earnings – by quartile of lifetime earnings

Percentage of households with:	Lowest av. lifetime earners	Quartile 2	Quartile 3	Highest av. lifetime earners	All
Pension income					
<67% replacement	2.4	12.0	23.7	29.8	19.6
<80% replacement	4.9	26.1	40.9	50.2	35.0
<PC replacement	4.9	14.8	23.7	20.4	17.9
Pension income and non-housing wealth					
<67% replacement	2.4	6.4	13.4	13.4	10.0
<80% replacement	4.1	17.3	24.7	24.1	19.9
<PC replacement	4.1	9.2	14.4	9.4	10.1
Total income					
<67% replacement	1.6	1.4	3.8	2.0	2.3
<80% replacement	3.3	4.2	7.2	5.4	5.3
<PC replacement	3.3	1.8	4.5	1.7	2.7
N	123	283	291	299	996

Note: Quartiles of the average lifetime earnings distribution are defined for the whole population of households in the cohort, including both single and couple households (earnings are equivalised). There are not equal numbers of couple households in each quartile because couples tend to be higher up the (equivalised) lifetime earnings distribution than single households. We exclude years of self-employment when calculating average lifetime earnings, and therefore the ranking of individuals according to their average lifetime earnings may differ slightly from their ranking according to their total lifetime earnings.

In large part, this pattern is driven by the structure of the UK state pension system: benefits are only somewhat related to earnings/contributions, and therefore state pension income generally represents much higher income replacement for those on lower incomes than for those on higher incomes. However, households with higher average lifetime earnings tend to accumulate greater stocks of wealth, and so once we include the annuitised value of non-housing wealth in retirement income, and the annuitised value of all wealth in retirement income (‘total income’), the relationship between average lifetime earnings and the proportion of

calculating average lifetime earnings, and therefore the ranking of individuals according to their average lifetime earnings may differ slightly from their ranking according to their total lifetime earnings.

households falling below the 67%/80%/Pensions Commission replacement rate thresholds is less stark.

Table 6.4 presents the results of multivariate regression and shows that the negative relationship between average lifetime earnings and the replacement rate remains even after controlling for other characteristics such as education, numeracy, number of children and current region of residence. The only characteristic (of those we consider) other than lifetime earnings that is correlated with the replacement of average lifetime earnings by pension income in retirement is the number of children. However, this result should be interpreted with caution since it will be sensitive to the method used to equivalise earnings to take account of the changing composition of the household between working life and retirement.

TABLE 6.4

Characteristics associated with replacement rate (median regression)

	Median replacement from: Pension income	Median replacement from: Total income
Low education	–	–
Mid education	–0.005 (0.033)	0.103 (0.076)
High education	–0.022 (0.032)	0.293*** (0.075)
Lowest numeracy	–	–
Second-lowest numeracy	–0.004 (0.060)	0.050 (0.137)
Second-highest numeracy	0.055 (0.060)	0.228 (0.139)
Highest numeracy	0.077 (0.064)	0.316* (0.146)
No children	–	–
1 or 2 children	0.155*** (0.047)	0.160 (0.107)
3 or more children	0.203*** (0.050)	0.223 (0.115)
Lowest average lifetime earnings	–	–
Quartile 2	–0.299*** (0.042)	–0.370*** (0.098)
Quartile 3	–0.403*** (0.044)	–0.517*** (0.102)
Highest average lifetime earnings	–0.450*** (0.047)	–0.647*** (0.107)
North East	–	–
North West	–0.022 (0.057)	0.057 (0.130)
Yorkshire and The Humber	–0.056 (0.058)	–0.022 (0.135)
East Midlands	–0.052 (0.059)	0.115 (0.136)
West Midlands	–0.070 (0.062)	0.041 (0.142)
East of England	–0.066 (0.058)	0.254 (0.133)
London	–0.039 (0.064)	0.612*** (0.146)
South East	–0.039 (0.055)	0.377*** (0.128)
South West	–0.003 (0.060)	0.339*** (0.138)

Note: Education and numeracy are for the man in the couple. Low education is defined as less than O level or equivalent, mid education is defined as O level or equivalent, and high education is defined as A level or equivalent and above. Numeracy groups are defined following the methodology of Banks et al. (2010) using ELSA questions designed to gauge the numerical ability of respondents. Figures are marginal effects, with standard errors in parentheses. ***, ** and * indicate statistical significance at the 1%, 5% and 10% level respectively.

If we consider the replacement of average lifetime earnings by total income in retirement (in other words, including the annuitised value of total wealth), we also find that higher education is associated with a higher replacement rate, as is living in certain parts of the country (notably, London, the South East or South West). This last factor is likely to be due to high house price growth in these areas in the past having increased households' wealth holdings, and therefore increased their potential retirement resources relative to their lifetime earnings.

The association between household characteristics and the probability of the replacement rate being less than the 67%, 80% or Pensions Commission thresholds is described in Tables A.3A and A.3B in the appendix. As would be expected given the results described above, after controlling for other household characteristics, having higher average lifetime earnings is associated with a greater chance of having 'inadequate' resources when retirement income is narrowly defined to include income only from pensions but not when the annuitised value of wealth is also included in retirement income.

6.4. Sensitivity

The results presented above suggest that, on the basis of their pension income alone, only around one-fifth of households are at risk of having less than what is commonly thought to be an adequate replacement rate in retirement. If we take into account other household wealth that could be used to finance retirement, virtually no households have inadequate resources.

These findings, however, are very sensitive to how we make income received in different periods comparable. Throughout the results presented above, we have assumed that households are concerned with replacing their standards of living in terms of their purchasing power, and therefore we have uprated earnings in line with inflation. If instead we were to assume that households were concerned with replacing their standards of living in terms of their income relative to others, and we uprate earnings in line with average earnings growth, the estimated distribution of income replacement rates is as shown in Figure 6.3.[26] Rather than only 20% of households having less than 67% replacement in retirement from their pension income, we would find that 70% of households are in that position, and rather than 35% of households being unable to achieve 80% replacement, we would find that 85% of households are in that position.

This indexation assumption is so important because we are considering the replacement of lifetime earnings and earnings growth has been substantially greater than growth in prices over the period we consider.[27] The different indexation assumption will have the most impact on earnings received in years furthest from 2002–03, which are equally important when

[26] The full counterpart to Table 6.2 is provided in Table A.4 in the appendix.

[27] For example, while prices rose by a factor of 12 between 1965 and 2002, earnings rose by a factor of 24.

we are considering replacement of average lifetime earnings but might not be important at all if we were, for example, considering replacement of average earnings between the ages of 40 and 50. While this sensitivity of our results is, perhaps, undesirable, the advantage of estimating the replacement of average lifetime earnings is that we can avoid the problems of volatility that comparisons with any particular year of earnings involve, and we do not have to make an arbitrary assumption about which period of life a household would like to replicate the standards of living of in retirement.

Whether households are actually concerned with replacing their purchasing power or their relative income in retirement is an open question.

FIGURE 6.3

Distribution of estimated gross income replacement in retirement – by uprating assumption

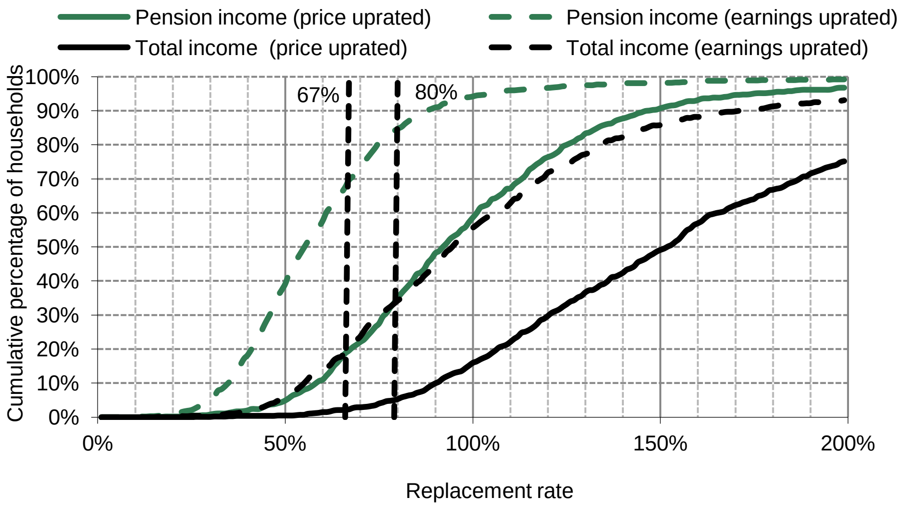

7. Results: the optimality of saving

To assess the optimality of household saving, we use our model to simulate how much wealth households should have accumulated by the age they are in 2002–03 given their lifetime circumstances, and compare that with how much wealth households report having accumulated in ELSA in 2002–03. The results from this analysis are presented in this section. We start in Section 7.1 by describing the wealth levels that are simulated to be optimal by our model, before turning to a comparison with actual wealth holdings in Section 7.2. We discuss some sensitivity analyses in Section 7.3.

7.1. Simulated optimal wealth

The cumulative distribution of optimal private wealth holdings simulated by our model is illustrated in Figure 7.1. All households are simulated to have optimal total wealth holdings that are positive, although for some households optimal wealth holdings are very small (12% of households have optimal wealth of less than £1,000). The level of total private wealth that is considered optimal has a wide distribution across the different households in our sample: median optimal wealth is £76,990, but 25% of households are simulated to have optimal wealth of less than £21,000, while 25% of households are simulated to have optimal wealth of more than £201,000. Recall that this is not optimal wealth at retirement, but optimal wealth when observed in the data, at which point households are aged between 52 and 63. (We describe the distribution of optimal wealth at retirement in Section 7.4.)

Our model allows households to choose between two types of asset: a safe asset and a DC pension. Figure 7.1 also illustrates the distribution of optimal saving in each asset. Our model suggests that 42% of households would optimally hold no private pension wealth by the time they are observed (though recall that households are aged between 52 and 63 at this point, so have up to 14 years left before they reach their SPA). The median optimal private pension wealth holding would be just £21,339, but 25% of households would optimally hold more than £172,000 in this form. In contrast, optimal non-pension wealth is positive but low for all households. Median optimal holding of the safe asset is simulated to be £27,708, and only 10% of households are simulated to have optimal non-pension wealth in excess of £80,000.

FIGURE 7.1

Simulated optimal wealth

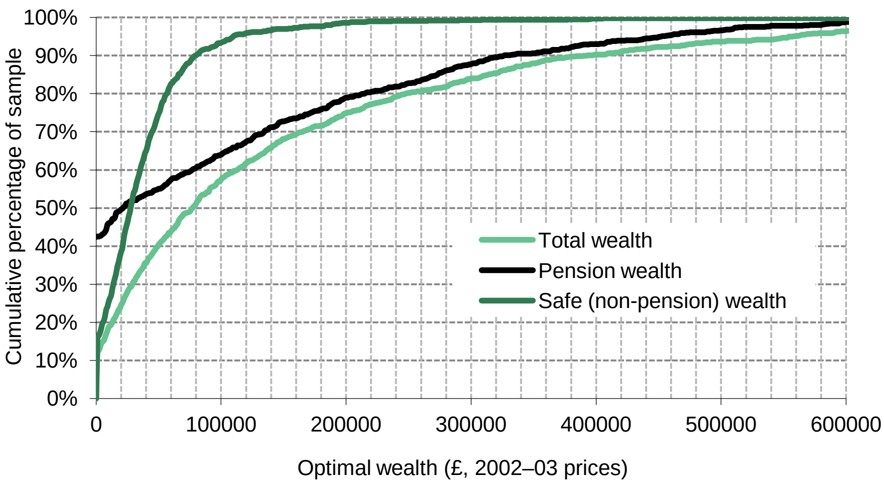

Our model suggests that the majority of households should optimally hold considerably more wealth in a DC pension than in the safe asset for three reasons. First, the pension is expected to offer a higher return than the safe asset. Second, the pension is tax advantaged since contributions are paid gross of income tax, and one-quarter of the pension fund can be taken tax free at the age of 65. (While the DC pension is more illiquid than the safe asset – in that it cannot be accessed until age 65 and 75% of the fund must be used to purchase an annuity – this does not particularly deter households in our model from saving in this vehicle because saving behaviour is primarily motivated by retirement saving rather than precautionary saving.) The third reason is that saving in a pension gives households access to the annuity market. Annuities are a very effective way of insuring against longevity (i.e. ensuring that accumulated assets do not run out before death).[28]

One of the main reasons that there is such a wide distribution for simulated optimal wealth is that different households have very different lifetime earnings histories. Table 7.1 summarises how median simulated optimal wealth differs for households in different quartiles of the lifetime earnings distribution. Optimal private wealth holdings are simulated to be increasing, on average, with average lifetime earnings: among the quarter of households with the lowest average lifetime earnings median optimal wealth is simulated to be £3,617, while among the quarter of households with the highest average lifetime earnings median optimal wealth is simulated to be £305,066. This is intuitive, since households with higher average earnings would need to accumulate greater wealth (in absolute terms) in order to maintain their consumption in retirement. What is perhaps surprising is the relatively low levels of wealth that our model suggests are optimal for some households; for example, among the quarter of households with the lowest average lifetime earnings, median optimal private pension wealth is simulated to be zero (as is also the case for the second-poorest quarter of households). However, it is important to remember that our model is simulating optimal private wealth holdings, conditional on a household's state pension entitlement. As indicated in the penultimate column of Table 7.1, median state pension wealth, at over £100,000, is relatively high among low-average-lifetime-earning households, and therefore relatively little private saving may be required to provide these households with an optimal retirement income.

TABLE 7.1

Median optimal wealth – by average lifetime earnings

	Total (private) wealth	Pension wealth	Safe (non-pension) wealth	State pension wealth	Median average lifetime earnings
All	£76,990	£21,339	£27,708	£120,184	£20,992
Lowest av. lifetime earners	£3,617	£0	£947	£103,978	£10,140
Quartile 2	£32,669	£0	£23,132	£119,913	£16,866
Quartile 3	£90,446	£45,316	£31,125	£124,115	£21,695
Highest av. lifetime earners	£305,066	£265,825	£41,861	£125,534	£31,866

Note: Figures are in 2002–03 prices. State pension wealth is observed pension wealth; optimal private (pension and safe) wealth is simulated conditional on that state pension wealth.

[28] In spite of the theoretical attractiveness of annuities, economists have documented an 'annuity puzzle' in that in countries other than the UK (where the purchase of annuities was, until very recently, compulsory for those with DC pensions), very few annuities are purchased. See Brown (2007) for a discussion of this.

7.2. Have households saved optimally?

The comparison between observed levels of wealth and what our model simulates would be optimal is shown for all the households in our sample in Figure 7.2 (each dot represents a household). The dashed 45-degree line illustrates where households would lie if their observed wealth were equal to their optimal wealth and the continuous line represents a smoothed line of best fit. The majority of households (92%) lie above the 45-degree line – in other words, they have observed wealth that is greater than what our model suggests is optimal. In fact, they hold considerably more wealth than is optimal: conditional on being an ‘oversaver’, the median surplus is £226,491. As a proportion of optimal wealth, the median surplus is 316% – in other words, more than half of oversaving households hold more than four times as much wealth as our model suggests is optimal. Those who are ‘undersaving’, on the other hand, are doing so to a much lesser extent.

FIGURE 7.2

Comparing observed and optimal wealth holdings

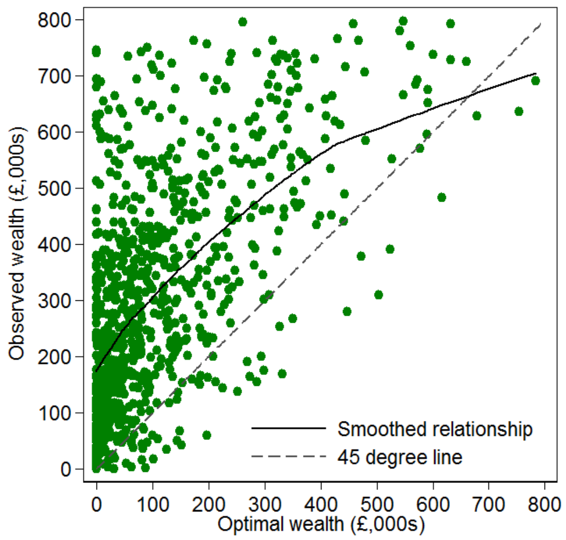

Source: Crawford and O’Dea, 2014.

The comparisons between observed and optimal wealth for the two types of asset in our model are shown in Figure 7.3 and Table 7.2. The left-hand panel of the figure shows how households’ observed pension wealth holdings (note that this includes all pension wealth, not just that held in DC pensions) compare with what our model suggests they should hold in a DC pension. Households are more evenly distributed around the 45-degree line than is the case for total wealth, with 75% of households holding more in pension wealth than our model suggests is optimal. The right-hand panel draws a very different picture for non-pension wealth (recall that in our model there is only one non-pension asset – the safe asset – whereas observed non-pension wealth will include financial wealth, housing wealth, business wealth and physical wealth). Virtually all households (91%) hold greater-than-optimal non-pension wealth. Our model suggests that the median optimal amount of safe wealth was £27,708, while the median observed amount of non-pension wealth is £164,797. The median surplus among households who have oversaved in non-pension wealth is £146,080; on average, households have saved nearly six times as much non-pension wealth as our model suggests is optimal.

FIGURE 7.3

Comparing observed and optimal wealth holdings – by asset type

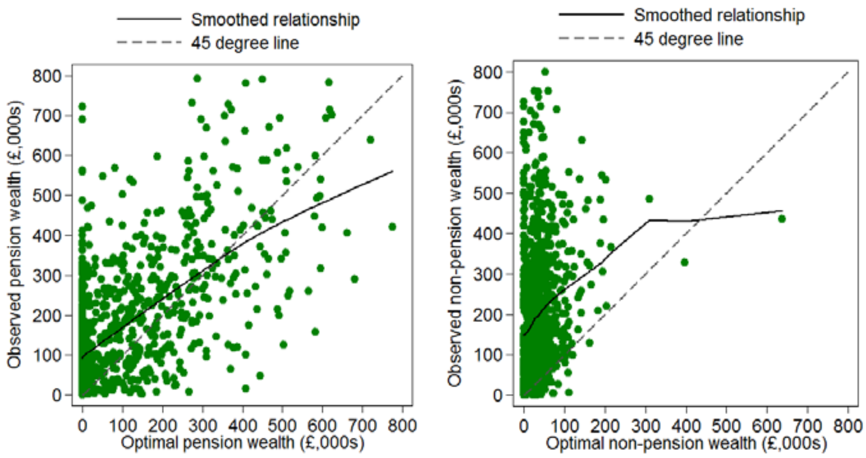

Note: Observed pension wealth includes all pension wealth (DC and DB), while optimal pension wealth is in DC pensions only. Observed non-pension wealth includes financial wealth, housing wealth, business wealth and physical wealth, while optimal non-pension wealth is in the safe asset only.
Source: Crawford and O'Dea, 2014.

TABLE 7.2

Summary statistics for comparison of observed and optimal wealth

	Total wealth	Pension wealth	Safe (non-pension) wealth
Median optimal wealth	£76,990	£21,339	£27,708
Median observed wealth	£324,135	£123,358	£164,797
Percentage oversaving	92%	75%	91%
Median £ surplus (cond. on oversaving)	£226,491	£86,335	£146,080
Median % surplus (cond. on oversaving)	316%	94%	598%
Median £ deficit (cond. on undersaving)	£38,747	£71,987	£16,333
Median % deficit (cond. on undersaving)	40%	45%	100%

Note: 5% of households have optimal and observed pension wealth of zero and therefore are neither undersaving nor oversaving.

Who has more than optimal wealth?

The apparent oversaving of households is a feature that appears across the distribution of households' lifetime earnings. Figure A.3 in the appendix reproduces Figure 7.2, but with households split into four groups according to their average lifetime earnings (shown in separate panels for clarity). Households in the lowest quarter of the average lifetime earnings distribution (illustrated in panel 1) are simulated by the model to have very low levels of optimal wealth, but are still observed with relatively large positive wealth holdings. Many households in the highest quarter of the average lifetime earnings distribution (panel 4) are simulated by the model to have much higher optimal levels of wealth, but are still typically observed to hold even greater wealth levels than this.

TABLE 7.3

Characteristics associated with the probability of having more than optimal wealth, and the median surplus conditional on being an oversaver

	Probit model (probability of oversaving)	Median regression (conditional on £ surplus)
Low education	–	–
Mid education	0.022 (0.022)	16.666 (21.073)
High education	0.002 (0.023)	49.780** (20.918)
Lowest numeracy	–	–
Second-lowest numeracy	0.078 (0.055)	18.533 (40.227)
Second-highest numeracy	0.092 (0.056)	58.421 (40.550)
Highest numeracy	0.121* (0.057)	92.677* (42.545)
No children	–	–
1 or 2 children	−0.053 (0.044)	22.395 (29.296)
3 or more children	−0.128*** (0.044)	6.920 (31.801)
Lowest average lifetime earnings	–	–
Quartile 2	0.055 (0.029)	54.610 (28.112)
Quartile 3	0.036 (0.032)	88.393*** (29.275)
Highest average lifetime earnings	0.010 (0.036)	140.673*** (30.536)
North East	–	–
North West	−0.038 (0.037)	−1.817 (36.626)
Yorkshire and The Humber	0.001 (0.035)	−32.212 (37.553)
East Midlands	0.033 (0.032)	9.466 (37.709)
West Midlands	−0.048 (0.043)	−18.683 (40.216)
East of England	−0.053 (0.040)	21.201 (37.460)
London	−0.037 (0.043)	151.981*** (41.121)
South East	−0.010 (0.035)	33.461 (35.592)
South West	0.016 (0.035)	79.679 (38.241)

Note: Education and numeracy are for the man in the couple. Low education is defined as less than O level or equivalent, mid education is defined as O level or equivalent, and high education is defined as A level or equivalent and above. Numeracy groups are defined following the methodology of Banks et al. (2010) using ELSA questions designed to gauge the numerical ability of respondents. Figures are marginal effects, with standard errors in parentheses. ***, ** and * indicate statistical significance at the 1%, 5% and 10% level respectively.

In fact, there are very few household characteristics that appear to be associated with a household’s propensity to have more than optimal wealth.

Table 7.3 presents the results of a probit regression on the probability of having more than optimal wealth. Having 3 or more children (relative to having no children) is associated with a 13 percentage point lower probability of having greater than optimal wealth. But, perhaps surprisingly, education and numeracy do not appear to play a significant role, and nor does current region of residence. These characteristics are, however, associated with the size of the wealth surplus among oversavers – shown in the final column of Table 7.3. Interestingly, currently living in London is also associated with a significantly larger wealth surplus among those who have greater than optimal wealth. This could be explained by large, unanticipated house price growth resulting in households' wealth holdings increasing in a way they did not expect when they made their saving decisions.

One feature that does appear to be associated with whether a household is 'oversaving' is whether the household is an owner-occupier.[29] Only 10% of our sample are not owner-occupiers, but Figure A.4 in the appendix illustrates that these households cluster around the 45-degree line to a greater extent than owner-occupiers. However, it is difficult to interpret this, since owner-occupation is itself a household choice. On the one hand it could be high wealth levels that induce households to become owner-occupiers, while on the other it could be owner-occupation itself that subsequently influences households' profile of wealth accumulation. Here we simply note the association; in Section 7.3, we discuss the impact housing wealth may have on our main conclusions.

7.3. Why do households have greater than optimal wealth?

The important question then – particularly if we want to understand the policy implications or make inferences about how other cohorts might compare – is 'Why do households have far more wealth than our model suggests is optimal?'. In this section, we consider several possible answers to this question. First, we consider the possibility that our model understates optimal saving because of the parameters chosen in the model for key assumptions such as households' risk aversion and the rate at which households discount the future. Second, we consider timing effects. Third, we consider the importance of housing wealth. We then end by reiterating some of the factors discussed in Section 3 that our model does not account for but that might influence wealth accumulation.

Sensitivity to chosen parameters

In order to construct the life-cycle model used to simulate optimal wealth, we need to make a number of assumptions about household preferences and

[29] We do not include owner-occupation as an explanatory variable in the regressions presented in Table 7.3 because owner-occupation is a choice, and may be a symptom of oversaving as much as a cause of oversaving. Here we are interested in the association between underlying individual characteristics and oversaving, rather than whether or not that oversaving is exhibited through home-ownership.

expectations and a number of simplifying assumptions about the choices available to households (described in more detail in Section 3). The results obtained for optimal wealth (and therefore the comparison between optimal and observed wealth) will be sensitive to these assumptions.

In general, we have chosen assumptions for household preferences and expectations that are within the range suggested by previous studies, but, in some cases, biased towards our model simulating higher optimal wealth holdings (and therefore they bias the model away from finding the oversaving result that we find). It is despite this approach to the assumptions used that we find observed wealth holdings are in excess of what is simulated to be optimal. Were we to use what might be considered more 'central' assumptions, our model would simulate lower optimal wealth and the comparison would be even starker.

TABLE 7.4

Sensitivity of results to selected modelling assumptions

	Median optimal wealth	Percentage oversaving	Median £ surplus (cond.)	Median £ deficit (cond.)
Baseline				
Total wealth	£76,990	92%	£226,491	£38,747
Pension wealth	£21,339	75%	£86,335	£71,987
Safe (non-pension) wealth	£27,708	91%	£146,080	£16,333
More patient				
Total wealth	£300,946	57%	£137,657	£94,241
Pension wealth	£261,746	23%	£61,480	£136,783
Safe (non-pension) wealth	£34,203	87%	£138,049	£26,596
Retire at age 60				
Total wealth	£80,660	90%	£207,752	£58,385
Pension wealth	£0	78%	£98,790	£157,973
Safe (non-pension) wealth	£47,844	84%	£132,474	£30,922
Higher risk aversion				
Total wealth	£75,041	92%	£223,293	£33,725
Pension wealth	£0	80%	£98,634	£73,113
Safe (non-pension) wealth	£50,194	85%	£133,527	£24,129

Note: In the baseline model 5% of households have optimal and observed pension wealth of zero and therefore are neither undersaving nor oversaving. When we assume households are more patient, this proportion is 2%; when we assume households retire at age 60, it is 6%; and when we assume households are more risk averse, the proportion is 5%.

Table 7.4 illustrates the sensitivity of our results to the main assumptions in the model. Assuming households are more patient increases optimal wealth (intuitively, all else equal, a more patient household will want to save more for retirement than a less patient household) and therefore reduces the proportion found to be oversaving.[30] However, even if households did not discount the future *at all* (as shown in the second panel of Table 7.4), 57% of our sample would still be assessed to be oversaving (by an average £137,657). Assuming households retire at age 60 rather than 65 makes little difference to optimal total wealth. Assuming households are more risk averse than in our basic model also makes little difference to optimal total wealth, although the optimal composition is shifted somewhat in favour of the safe asset (which offers a lower but certain rate of return).

Overall, these sensitivity analyses demonstrate that our main finding – that most households hold considerably more than optimal wealth – is robust to the main assumptions made in the model.

Timing effects

Another reason why households are found to hold more wealth than our model suggests is optimal could be the timing of the comparison. Our model simulates not just how much wealth households should accumulate for their retirement, but also *when* they should accumulate that wealth. For many households, that optimal accumulation will happen towards the end of working life since children are more likely to have left home, and giving up consumption to fund saving is always more palatable in the future. Optimal wealth will be highest at age 64, just before modelled retirement.

TABLE 7.5

Summary statistics for comparison of observed and optimal wealth at age 64

	Total wealth	Pension wealth	Safe (non-pension) wealth
Median optimal wealth at age 64	£153,609	£88,696	£3,819
Median observed wealth	£324,135	£123,358	£164,797
Percentage oversaving	71%	52%	85%
Median £ surplus (cond. on oversaving)	£190,828	£90,045	£165,000
Median £ deficit (cond. on undersaving)	£105,016	£201,963	£30,824

Note: 3% of households have optimal pension wealth at age 64 and observed pension wealth of zero, and are therefore neither undersaving nor oversaving.

Our sample are aged between 52 and 63 when observed and so throughout we have compared observed wealth with what our model suggests would be optimal wealth at their age. However, if the members of our sample have done all their retirement saving by the time they are

[30] Our model assumes that households are impatient, but no more impatient than they are compensated for by the risk-free interest rate. Formally, the discount rate assumed is equal to the risk-free interest rate (2.2%). Available evidence suggests that households are, if anything, less patient than this (see a review of the literature by Frederick et al. (2002)). Incorporating this assumption into the model would *lower* optimal household wealth.

observed (i.e. saved 'early' relative to what our model suggests would be optimal), then this comparison may overstate the extent of oversaving, since optimal wealth will increase with age while actual wealth will not. This may be a concern, as modelling the precise timing of retirement saving is substantially more challenging than modelling its overall level. To investigate this concern, Table 7.5 presents the results of comparing observed wealth for each household with what our model suggests they should optimally hold at age 64. Median optimal wealth at age 64 is twice median optimal wealth at the age attained in 2002–03 (at £153,609 compared with £76,990). The proportion of households who are found to be oversaving is therefore lower, but 71% of our sample are still found to have wealth in excess of the optimal level, and the median conditional surplus is £190,828. In our model, households primarily use saving in the pension asset to fund retirement and saving in the safe asset for precautionary reasons (to fund consumption in periods of unemployment or low earnings), and therefore at age 64 the balance of optimal wealth is shifted more towards pension wealth rather than the safe asset. However, over half of the sample are still found to have pension wealth greater than what the model suggests is optimal and 85% of the sample are still found to have above optimal non-pension wealth. Timing effects therefore provide little explanation for the oversaving found among households.

Housing

An obvious limitation of our model is that we do not model housing as a separate asset. As described in Section 3, housing wealth is different from most other assets in that it provides a consumption value, and it has also benefited from very high returns in the UK due to both high house price growth and the leveraged nature of mortgages. These factors would all increase the incentive for households to hold housing wealth in a way that our model does not capture, and could therefore lead us to understate optimal wealth. This is likely to be a particular issue for our sample of households, 90% of whom are owner-occupiers and among whom median net housing wealth is £120,000 (as shown in Table A.2 in the appendix).

We cannot include housing wealth in our model as a separate asset choice due to computational constraints. However, we can consider the following thought experiment: suppose that housing wealth is obtained at no cost but cannot be used to fund consumption and that owning that house does not affect the well-being households derive from other consumption. Under these conditions, have households saved optimally for retirement? To answer this question, we can simply compare the optimal wealth generated by our model and households' observed *non-housing* wealth holdings. This comparison is illustrated in Figure 7.4 and Table 7.6.[31]

As expected, when housing wealth is excluded, the differences between households' optimal wealth and observed wealth holdings are much smaller.

[31] Figure A.5 in the appendix provides the comparison between observed and optimal pension and safe (non-housing) wealth separately (the counterpart to Figure 7.3).

However, the key thing to note is that 75% of our sample are still found to have greater than optimal wealth, and the median surplus among those oversavers is still in excess of £120,000. This suggests that, while the omission of housing wealth from the model may be an important reason why many households in our sample have saved more than our model suggests is optimal, it is not the only (or, perhaps, even the most important) reason.

FIGURE 7.4

Comparing observed and optimal wealth holdings – excluding housing from observed wealth

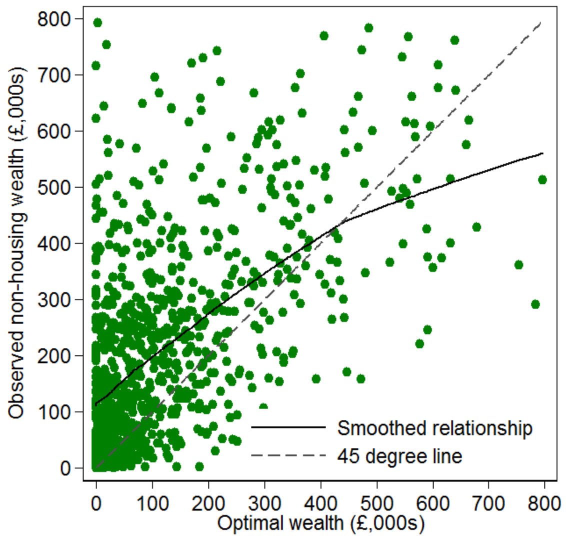

Source: Crawford and O'Dea, 2014.

TABLE 7.6

Summary statistics for comparison of observed and optimal wealth – excluding housing from observed wealth

	Total wealth	Pension wealth	Safe (non-pension) wealth
Median optimal wealth	£76,990	£21,339	£27,708
Median observed (non-housing) wealth	£188,930	£123,358	£32,140
Percentage oversaving	75%	75%	55%
Median £ surplus (cond. on oversaving)	£121,475	£86,335	£48,062
Median £ deficit (cond. on undersaving)	£43,701	£71,987	£25,782

Note: 5% of households have optimal pension wealth at age 64 and observed pension wealth of zero, and are therefore neither undersaving nor oversaving.

Other important factors

If the above explanations are not responsible for all of households' observed 'oversaving', then what else might cause our model to understate optimal wealth to the observed extent? As discussed in Section 3, there are a number of other features that are not incorporated in our model that might influence wealth accumulation. In particular, households may be concerned with maintaining their average position in the income distribution, they may be planning to leave bequests, they may be expecting to enjoy consumption more during retirement than during working life and/or they may have accumulated wealth against the risk of long-term care costs in later life. Investigating these potential explanations will prove interesting areas for future research.

8. How well does the replacement rate approach approximate the optimality approach?

The analysis presented in this paper is the first in the UK that has used a life-cycle model of consumption and saving to estimate optimal wealth and assess observed wealth holdings among households approaching retirement. Given that all analysis on the preparedness of households for retirement to date has used a replacement rate approach to assess the adequacy of wealth, and that this approach is likely to continue to be commonly used going forward due to its relative simplicity, it is useful to consider how the results of the two approaches would compare. That is therefore the focus of this section.

In Section 6, we considered a number of ways of measuring retirement income when calculating the replacement rate of average lifetime earnings. The definition that is most comparable to the results produced by our model of optimal retirement saving is the measure of replacement of average lifetime earnings by total income in retirement (where total income includes private and state pension income and the annuitised value of total wealth). On this measure of retirement income, we found that 97.7% of couple households had more than 67% replacement of average lifetime earnings, 94.7% had more than 80% replacement and 97.3% had greater replacement than the relevant Pensions Commission benchmark threshold. These are all somewhat higher than the 92.1% of the same households that our structural model suggested had accumulated more than optimal wealth.

Table 8.1 illustrates how the labelling of households as having 'inadequate' resources (less than a certain replacement rate) interacts with whether they are found to have more or less than optimal wealth. Given the high proportion of the sample found to have 'adequate' wealth and the high proportion of the sample found to have more than optimal wealth, there is obviously considerable overlap and so it is difficult to tell from this table how the two approaches compare.

TABLE 8.1

Overlap between 'adequate' and 'optimal' wealth

Results from benchmarking:	All	Results from modelling: Less than optimal wealth	Results from modelling: More than optimal wealth
Less than 67% replacement	2.3%	47.8%	52.2%
Greater than 67% replacement	97.7%	7.0%	93.0%
Less than 80% replacement	5.3%	41.5%	58.5%
Greater than 80% replacement	94.7%	6.0%	94.0%
Less than PC replacement	2.7%	40.7%	59.3%
Greater than PC replacement	97.3%	7.0%	93.0%
All		7.9%	92.1%

FIGURE 8.1

Implied optimal replacement rates from the modelling approach

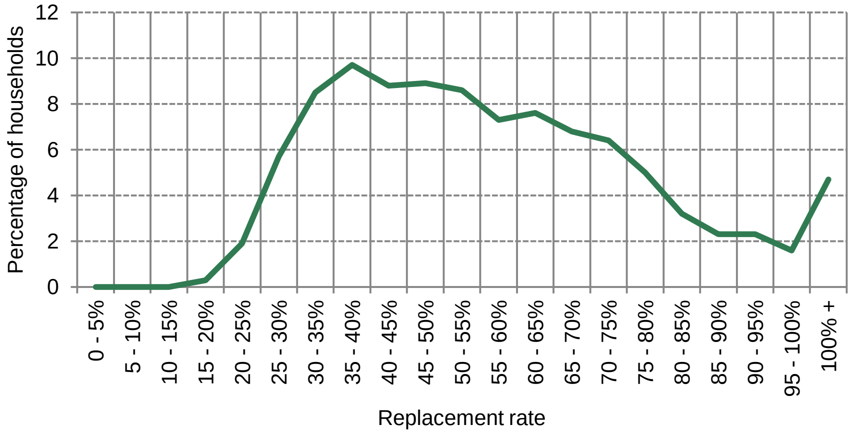

Greater insight can be gained from Figure 8.1, which illustrates the distribution of replacement rates (total income at age 65 divided by average lifetime earnings) implied by the profile of optimal wealth accumulation simulated for each household by our model.[32] The most important thing to note is the fairly wide distribution of optimal replacement rates: one-quarter of households have an implied optimal replacement rate of less than 40%, while one-quarter of households have an implied optimal replacement rate in

[32] These replacement rates are calculated excluding means-tested benefits from total income at age 65. They are therefore comparable to the actual replacement rates estimated in Section 6, which also do not include means-tested benefits. However, they are estimated conditional on means-tested pension credit being available to eligible households. In the absence of such benefits, the optimal replacement rate for each household may differ.

excess of 70%. The implied optimal level of earnings replacement is also relatively low for many individuals – in particular, over two-thirds of households have an implied optimal replacement rate of less than 67% (and over four-fifths of households less than 80%). This suggests that, for many individuals, commonly-used replacement rate benchmarks are a more demanding test of their preparedness for retirement than a comparison with the optimal wealth simulated by our structural model.

FIGURE 8.2

Implied optimal replacement rates from the modelling approach – by position in the average lifetime earnings distribution

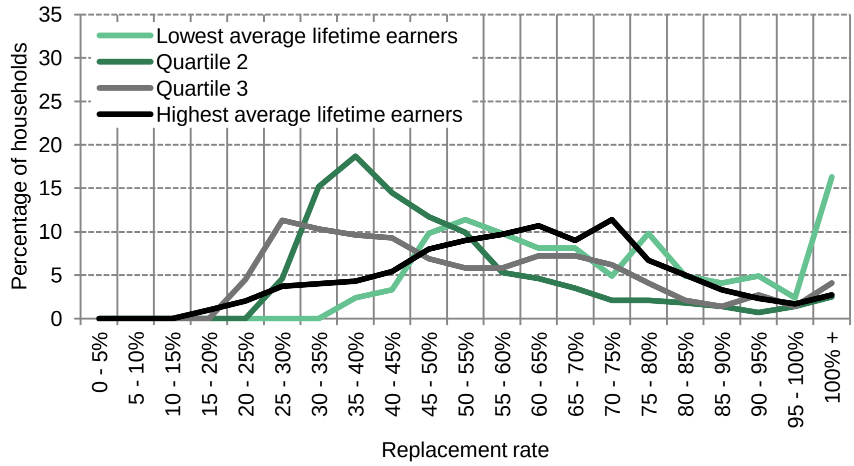

The wide distribution of implied optimal replacement rates suggests that using a single benchmark threshold to assess the adequacy of resources is unlikely to give an accurate picture of the proportion of individuals who have or have not saved sufficiently to maintain their standards of living in retirement. The Pensions Commission benchmark, with its target replacement rates that vary depending on earnings, has commonly been used in the UK literature to address this suspected (but not previously quantified) problem. However, Figure 8.2 illustrates that this is unlikely to provide much of a solution. There is a wide distribution of implied optimal replacement rates even within quartiles of the average lifetime earnings distribution, and it is not clear that the range of optimal values is distinct for households in different parts of the average lifetime earnings distribution. In addition, while the median implied optimal replacement rate is highest for those in the lowest quarter of the lifetime earnings distribution (at 68%), it is then increasing among the other wealth quartiles (44%, 49% and 61% for the second, third and richest quartiles respectively) in contrast to the pattern assumed for the Pensions Commission benchmark.[33] This strongly suggests

[33] This pattern between quartiles is driven largely by the differing extent (across quartiles) to which average tax rates change at the age of 65. In 2002–03 (the year from which the model's tax system is taken), the tax-free allowance rose significantly at the age of 65. Additionally, National Insurance

that it is not just the level of a household's average lifetime earnings that matters for how much they would optimally like to replace in retirement, but also the timing of that income and other household circumstances such as the number and timing of children.

9. Conclusions

Much of the UK policy debate over the past decade has been on the adequacy (or otherwise) of individuals' private retirement saving. At a time when defined benefit pensions are increasingly scarce, and the UK state pension system is moving away from providing any earnings-related benefits, this matter is becoming ever more important.

Existing analysis of this issue to date has attempted to assess the likely adequacy of household retirement resources by comparing projected income replacement rates in retirement to some illustrative benchmarks. However, these benchmarks are necessarily somewhat arbitrarily chosen, and are often justified based on the retirement experience of previous cohorts without any attempt to assess whether these cohorts themselves 'did the right thing'.

In this paper, we present the first assessment of the *optimality* of households' saving, focusing on couple households born in the 1940s. We use a life-cycle model of consumption and saving to determine what households ought to have saved, given their life circumstances, if their objective is to maintain their living standards over time, and compare that optimal wealth with their observed wealth holdings. We find that the vast majority of households (92%) hold more wealth than our model suggests is optimal, and that that would still be true for 75% of households were we to exclude housing wealth from observed wealth holdings. This perhaps suggests a role for some combination of bequest intentions, precautionary saving against long-term care costs, households expecting their well-being from consumption to change in retirement and households being concerned with maintaining their position in the income distribution in addition to maintaining their average consumption bundle. These are all features that are not currently captured by our model, and exploring these possibilities will be an interesting avenue for future research.

To place our results in the context of the existing adequacy literature, we also calculated the earnings replacement these same households are likely to enjoy in retirement. We find that, once non-pension wealth is included, only 2% of households would be unable to replace 67% of their average lifetime earnings at age 65, only 5% would be unable to replace 80% and only 3% would be unable to achieve the replacement rate assumed to be adequate by the Pensions Commission benchmark. However, for most households, such thresholds for adequacy are actually a more demanding test of the adequacy of their resources than a comparison with the optimal wealth suggested by

contributions are not payable on pension income but are on earnings. Those likely to see the biggest falls in their average tax rates (those in the second quartile, on average), have a lesser need to replace their gross income to maintain a given level of net income.

our model. In particular, over two-thirds of households have an implied optimal replacement rate of less than 67% and over four-fifths have an implied optimal replacement rate of less than 80%.

However, what is more important than the average level of earnings replacement implied by our model to be optimal is the range of the distribution. One-quarter of households are found to have an optimal replacement rate of less than 40%, while one-quarter have an optimal replacement rate of over 70%. Even once households are grouped according to their average lifetime earnings, there is a wide distribution of optimal replacement rates *within* these groups. It is not just the level of average earnings that matters, but also the timing of earnings and other aspects of the households' circumstances such as their number and timing of children.

This suggests that caution should be used when using simple replacement rate benchmarks to assess the adequacy of households' resources. Such analysis could give a misleading picture of the preparedness of households for retirement, since it cannot capture the vast heterogeneity in households' circumstances. While this does not appear to matter so much for the sample of households analysed in this paper (couple households born in the 1940s) – among whom there was considerable overlap in the households found to have adequate resources and the households found to have optimal resources – this is largely driven by this cohort holding considerably in excess of optimal wealth. For later cohorts, where there is anecdotally much greater concern about undersaving, the divergence between households assessed to have optimal wealth and households assessed to have adequate wealth using a simple replacement rate benchmark would be expected to be greater.

References

Aegon (2013), 'The changing face of retirement: the Aegon retirement readiness survey 2013'.

Aguiar, M. and Hurst, E. (2005), 'Consumption versus expenditure', *Journal of Political Economy*, 113, 919–48.

Banks, J., O'Dea, C. and Oldfield, Z. (2010), 'Cognitive function, numeracy and retirement savings trajectories', *Economic Journal*, 120, F381–410.

Banks, J., Emmerson, C., Oldfield, Z. and Tetlow, G. (2005), *Prepared for Retirement? The Adequacy and Distribution of Retirement Resources in England*, Report 67, London: Institute for Fiscal Studies.

Barclays Capital (2012), 'The Equity Gilt Study 2012', tech. rep., Barclays Capital, London.

Bardasi, E., Jenkins, S. and Rigg, J. (2002), 'Retirement and the income of older people: a British perspective', *Ageing and Society*, 22, 131–59.

Binswanger, J. and Schunk, D. (2012), 'What is an adequate standard of living during retirement?', *Journal of Pension Economics and Finance*, 11, 203–22.

Blundell, R. and Tanner, S. (1999), ‘Labour force participation and retirement in the UK’, paper prepared for National Academy of Science, Institute for Fiscal Studies, mimeo.

Bozio, A., Crawford, R., Emmerson, C. and Tetlow G. (2010), ‘Retirement outcomes and lifetime earnings: descriptive evidence from ELSA-NI data’, Department for Work and Pensions (DWP), Working Paper 81.

Bozio, A., Emmerson, C., and Tetlow, G. (2011), ‘How much do lifetime earnings explain retirement resources?’, Institute for Fiscal Studies, Working Paper 11/02.

Brown, J.R. (2007), ‘Rational and behavioral perspectives on the role of annuities in retirement planning’, National Bureau of Economic Research, Working Paper 13537.

Browning, M. and Meghir, C. (1991), ‘The effects of male and female labor supply on commodity demands’, *Econometrica*, 59, 925–51.

Crawford, R. (2012), ‘ELSA pension wealth derived variables (waves 2 to 4): methodology’, Institute for Fiscal Studies, mimeo.

Crawford, R. and O’Dea, C. (2012), *The Adequacy of Wealth among Those Approaching Retirement*, Report 72, London: Institute for Fiscal Studies.

Crawford, R. and O’Dea, C. (2014), ‘Cash and Pensions: Have the elderly in England saved optimally for retirement?’, Institute for Fiscal Studies, Working Paper WP14/22.

Crawford, R. and Tetlow, G. (2012), ‘The evolution of pension wealth and contribution dynamics’, in J. Banks, J. Nazroo and A. Steptoe (eds), *The Dynamics of Ageing: Evidence from the English Longitudinal Study of Ageing 2002–10*, London: Institute for Fiscal Studies.

De Nardi, M., French, E. and Jones, J.B. (2010), ‘Why do the elderly save? The role of medical expenses’, *Journal of Political Economy*, 118, 39–75.

Department for Work and Pensions (2012), ‘Estimates of the number of people with inadequate retirement incomes’.

Department for Work and Pensions (2013), ‘Framework for the analysis of future pension incomes’.

Frederick, S., Loewenstein, G. and O’Donoghue, T. (2002), ‘Time discounting and time preference: a critical review’, *Journal of Economic Literature*, 40, 351–401.

MacDonald, B-J. and Moore, K.D. (2011), ‘Moving beyond the limitations of traditional replacement rates’, Society of Actuaries.

Mayhew, V. (2002), *Pensions 2002: Public Attitudes to Pensions and Saving for Retirement*, DWP Research Report 193, Leeds.

Mitchell, O.S., Poterba, J.M. and Warshawsky, M.J. (1999), ‘New evidence on the money’s worth of individual annuities’, *American Economic Review*, 89, 1299–318.

Munnell, A.H., Webb, A. and Golub-Sass, F.N. (2007), ‘Is there really a retirement savings crisis? An NRRI analysis’, Brief IB7/11, Center for Retirement Research at Boston College.

Munnell, A.H., Webb, A. and Golub-Sass, F.N. (2012), 'The National Retirement Risk Index: an update', Brief IB12/20, Center for Retirement Research at Boston College.

Murthi, M., Orszag, J.M. and Orszag, P.R. (1999), 'The value for money of annuities in the UK: theory, experience and policy', *Journal of Pensions Management*.

Pensions Commission (2004), *Pensions: Challenges and Choices – The First Report of the Pensions Commission*, London: The Stationery Office.

Pensions Policy Institute (2012), *Retirement Income and Assets: The Implications for Retirement Income of Government Policies to Extend Working Lives*.

Pensions Policy Institute (2013), *Automatic Enrolment: What Level of Pension Contribution is Needed to Obtain an Adequate Retirement Income?*

Scholz, J.K., Seshadri, A. and Khitatrakun, S. (2006), 'Are Americans saving "optimally" for retirement?', *Journal of Political Economy*, 114, 607–43.

Appendix

TABLE A.1

Sample restrictions

	Reduction in sample size		Remaining percentage of original sample
	#	%	
Initial sample	1,615		
Missing 'core' ELSA data (age, education, wealth)	–89	–5.5%	94.5%
Not linked to NI data	–462	–28.6%	65.9%
Observe fewer than 5 years of positive earnings	–68	–4.2%	61.7%
Resulting sample	996		

TABLE A.2

Characteristics and representativeness of the sample

	Couples in our sample		Couples in ELSA[a]	
	Men	Women	Men	Women
Mean age	56.8	54.1	56.9	51.6
Low education	35.3%	43.6%	36.5%	41.9%
Mid education	27.3%	32.1%	26.6%	31.1%
High education	37.3%	24.3%	36.8%	22.4%
Employee	60.6%***	59.3%***	55.8%	52.8%
Self-employed	10.7%***	3.7%***	14.9%	4.7%
Retired	14.9%	11.0%	14.3%	10.8%
Other	13.8%*	25.9%***	15.0%	31.7%
Has children	92.0%		91.5%	
Owner-occupier	89.6%***		87.6%	
Median total income (£pw)	433.5		425.8	
Median earned income (£pw)	343.5		330.6	
Median asset income (£pw)	3.6		3.5	
Median total net wealth (£)	164,797		166,100	
Median net non-housing wealth (£)	32,140		30,625	
Median net housing wealth (£)	120,000		120,000	
Sample size	996		1,615	

[a] With man born in the 1940s.

Note: ***, ** and * indicate a statistically significant difference between those in our sample and those in the initial ELSA sample (at the 1%, 5% and 10% level respectively). Generated using sex-specific probit regressions on each characteristic.

TABLE A.3A

Characteristics associated with probability of having less than 67%/80%/Pensions Commission replacement – from pension income

	Probability of less than:		
	67% replacement	80% replacement	PC replacement
Low education	–	–	–
Mid education	−0.017 (0.033)	0.017 (0.039)	−0.007 (0.033)
High education	−0.013 (0.032)	−0.023 (0.038)	−0.049 (0.031)
Lowest numeracy	–	–	–
Second-lowest numeracy	−0.037 (0.074)	0.023 (0.079)	−0.010 (0.068)
Second-highest numeracy	−0.048 (0.075)	−0.023 (0.080)	−0.015 (0.069)
Highest numeracy	−0.092 (0.077)	−0.109 (0.082)	−0.110 (0.069)
No children	–	–	–
1 or 2 children	−0.187*** (0.038)	−0.185*** (0.051)	−0.151*** (0.038)
3 or more children	−0.175*** (0.043)	−0.205*** (0.056)	−0.150*** (0.043)
Lowest average lifetime earnings	–	–	–
Quartile 2	0.100*** (0.024)	0.209*** (0.033)	0.099*** (0.027)
Quartile 3	0.220*** (0.029)	0.359*** (0.036)	0.193*** (0.032)
Highest average lifetime earnings	0.270*** (0.033)	0.456*** (0.039)	0.174*** (0.034)
North East	–	–	–
North West	0.039 (0.056)	0.069 (0.066)	0.032 (0.051)
Yorkshire and The Humber	0.030 (0.058)	0.081 (0.068)	0.057 (0.054)
East Midlands	0.092 (0.061)	0.112 (0.069)	0.141** (0.059)
West Midlands	0.024 (0.060)	0.155* (0.073)	0.038 (0.056)
East of England	0.030 (0.056)	0.125 (0.067)	0.022 (0.052)
London	0.049 (0.063)	0.098 (0.074)	0.053 (0.059)
South East	0.025 (0.054)	0.075 (0.063)	0.039 (0.050)
South West	0.055 (0.059)	0.074 (0.069)	0.058 (0.055)

Note: Figures are marginal effects, with standard errors in parentheses. ***, ** and * indicate statistical significance at the 1%, 5% and 10% level respectively.

TABLE A.3B

Characteristics associated with probability of having less than 67%/80%/Pensions Commission replacement – from total income

	Probability of less than:		
	67% replacement	80% replacement	PC replacement
Low education	–	–	–
Mid education	−0.002 (0.013)	−0.020 (0.019)	0.005 (0.014)
High education	−0.008 (0.012)	−0.027 (0.019)	−0.014 (0.012)
Lowest numeracy	–	–	–
Second-lowest numeracy	−0.019 (0.034)	0.027 (0.032)	−0.006 (0.028)
Second-highest numeracy	−0.016 (0.034)	0.019 (0.032)	−0.005 (0.029)
Highest numeracy	−0.036 (0.034)	−0.017 (0.032)	−0.027 (0.029)
No children	–	–	–
1 or 2 children	−0.046*** (0.014)	−0.086*** (0.020)	−0.060*** (0.016)
3 or more children	−0.042*** (0.015)	−0.079*** (0.023)	−0.051*** (0.016)
Lowest average lifetime earnings	–	–	–
Quartile 2	0.006 (0.011)	0.017 (0.018)	−0.005 (0.015)
Quartile 3	0.029 (0.015)	0.048* (0.021)	0.025 (0.019)
Highest average lifetime earnings	0.008 (0.013)	0.029 (0.021)	−0.011 (0.016)
North East	–	–	–
North West	−0.001 (0.019)	0.004 (0.038)	−0.018 (0.022)
Yorkshire and The Humber	0.050 (0.028)	0.003 (0.039)	0.041 (0.030)
East Midlands	−0.006 (0.019)	−0.039 (0.036)	−0.019 (0.030)
West Midlands	−0.004 (0.020)	−0.019 (0.039)	−0.006 (0.026)
East of England	0.022 (0.024)	−0.003 (0.038)	0.019 (0.028)
London	0.005 (0.023)	−0.022 (0.039)	−0.010 (0.026)
South East	−0.011 (0.017)	−0.060 (0.032)	−0.024 (0.021)
South West	−0.003 (0.020)	−0.029 (0.037)	0.017 (0.023)

Note: Figures are marginal effects, with standard errors in parentheses. ***, ** and * indicate statistical significance at the 1%, 5% and 10% level respectively.

TABLE A.4

Distribution of estimated gross household income replacement in retirement – when income is revalued in line with average earnings rather than prices

Percentage of households with:	State pension income	Total pension income	... plus annuitised non-housing wealth	... and plus annuitised housing wealth
<=40% replacement	73.6%	18.3%	9.6%	2.2%
<=50% replacement	88.0%	39.2%	23.9%	5.7%
<=67% replacement	96.4%	69.7%	51.6%	20.7%
<=80% replacement	98.1%	84.7%	67.8%	34.2%
<=100% replacement	98.9%	94.2%	83.0%	55.6%
>100% replacement	1.1%	5.8%	17.0%	44.4%
<= PC replacement	98.1%	70.2%	50.9%	22.6%

Note: Replacement rate is defined as the ratio between household income in the year in which the household reaches age 65 and average lifetime earnings of the household, adjusted for household size and average earnings growth.

FIGURE A.1

Cumulative distribution of estimated gross income in retirement

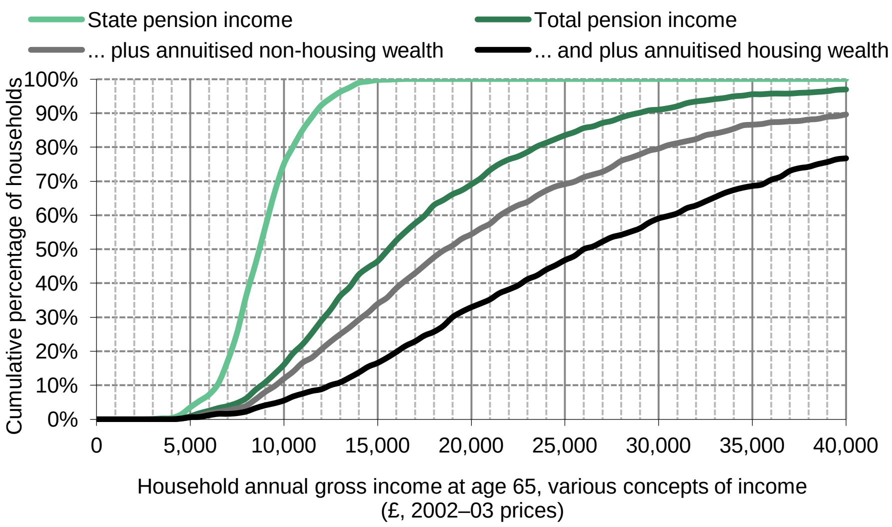

FIGURE A.2

Cumulative distribution of estimated gross income replacement in retirement

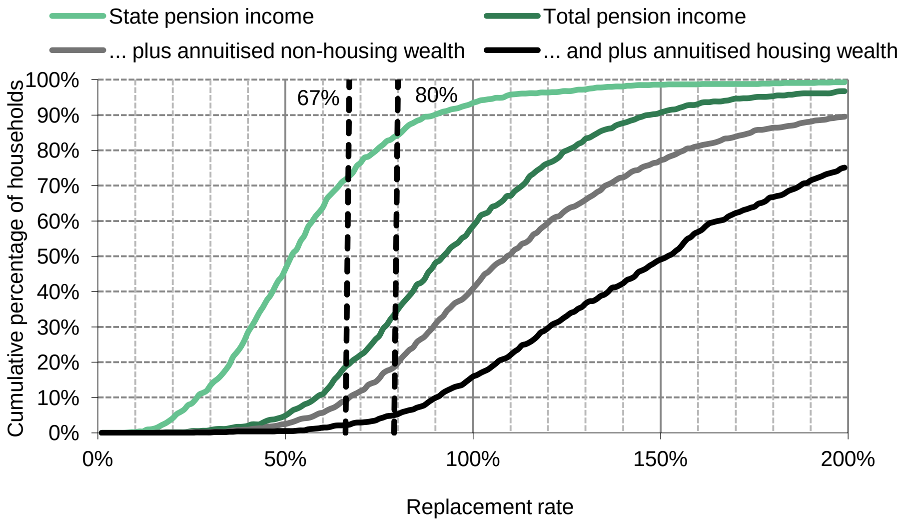

Note: Replacement rate is defined as the ratio between household income in the year in which the household reaches age 65 and average lifetime earnings of the household, adjusted for household size (as defined by equation 2 in Section 2).

FIGURE A.3

Comparing observed and optimal wealth holdings – by quartile of average lifetime earnings

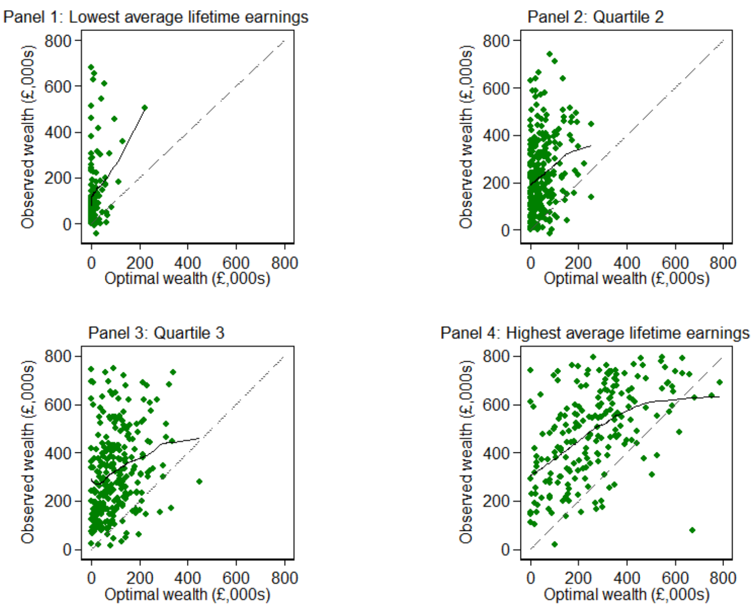

Source: Crawford and O'Dea, 2014.

FIGURE A.4

Comparing observed and optimal wealth holdings – by housing tenure

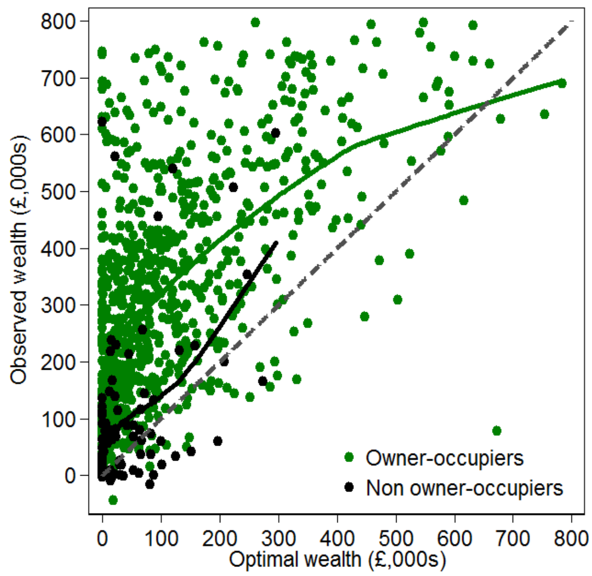

Note: Green circles and smoothed relationship are for owner-occupiers. Black circles and smoothed relationship are for non-owner-occupiers.

FIGURE A.5

Comparing observed and optimal wealth holdings – by asset type – excluding housing from observed wealth

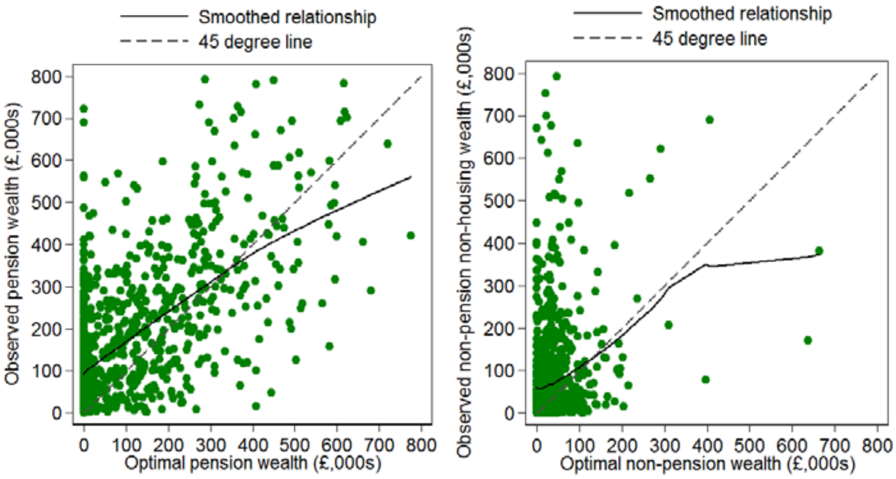

Note: Observed pension wealth includes all pension wealth (DB and DC), while optimal pension wealth is in DC pensions only. Observed non-pension wealth includes financial wealth, housing wealth, business wealth and physical wealth, while optimal non-pension wealth is in the safe asset only.