

Bozio, Antoine; Laroque, Guy; O'Dea, Cormac

**Working Paper**

## Heterogeneity in time preference in older households

IFS Working Papers, No. W13/02

**Provided in Cooperation with:**

The Institute for Fiscal Studies (IFS), London

*Suggested Citation:* Bozio, Antoine; Laroque, Guy; O'Dea, Cormac (2013) : Heterogeneity in time preference in older households, IFS Working Papers, No. W13/02, Institute for Fiscal Studies (IFS), London,  
<https://doi.org/10.1920/wp.ifs.2013.1302>

This Version is available at:

<https://hdl.handle.net/10419/91504>

**Standard-Nutzungsbedingungen:**

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

**Terms of use:**

*Documents in EconStor may be saved and copied for your personal and scholarly purposes.*

*You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.*

*If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.*

# Heterogeneity in time preference in older households

IFS Working Paper W13/02

Antoine Bozio  
Guy Laroque  
Cormac O'Dea

# Heterogeneity in time preference among older households\*

Antoine Bozio<sup>†</sup>

Guy Laroque<sup>‡</sup>

Cormac O’Dea<sup>§</sup>

December 5, 2012

## Abstract

This paper suggests a method for estimating the distribution of discount rates using panel data on income and wealth. Using the English Longitudinal Survey of Ageing (ELSA), a representative sample of the English population over age 50, we generate panel data on total consumption from the intertemporal budget constraint. The distribution of consumption levels is shown to closely match that estimated using the UK’s household budget survey. Consumption transitions over time are then used to estimate the discount rates of households. We show that there is substantial heterogeneity in discounting behaviour and find that, among this older population, households with less education or numerical ability exhibit greater patience than those with higher education or numerical ability. The direction of this association is the opposite to that which has been found in experimental investigations of time preference.

*Keywords:* Time preference; Discount rate; Heterogeneity; Consumption

*JEL Classification Numbers:* D12; D31; D91; E21

## 1 Introduction

In a number of economic situations individuals make decisions that involve a comparison of present and future circumstances. They must decide how much to invest in education, how much to save for retirement, how much to invest in health, etc. In each case these decisions are based on some assessment of the potential welfare at different periods under different scenarios. Following Samuelson (1937), economists have largely adopted a discounted-utility model which assumes that preferences over time can be condensed into one major parameter, the geometric discount rate (see Frederick et al. (2002) for a critical review and Hall (2010) for a review of recent research developed using this approach).

The empirical estimate of the discount rate – and the extent of its heterogeneity – is a crucial element for many research topics in economics and can have important policy implications. For instance, heterogeneity in discount rates has been discussed in the context of capital income taxation (Saez, 2002; Golosov et al., 2011; Diamond and Spinnewijn, 2011), in the context of pension reforms (Samwick, 1998;

---

\*Co-funding is acknowledged from the European Research Council (reference ERC-2010-AdG-269440 - WSCWTBDS) and the Economic and Social Research Council (Centre for Microeconomic Analysis of Public Policy reference RES-544-28-5001). Thanks to James Banks, Richard Blundell, James Browne and Rowena Crawford and to seminar participants at the Institute for Fiscal Studies for valuable comments. Any errors are our own.

<sup>†</sup>Paris School of Economics and Institute for Fiscal Studies

<sup>‡</sup>University College London and Institute for Fiscal Studies

<sup>§</sup>Institute for Fiscal Studies and University College London

Gustman and Steinmeier, 2005), in the context of poverty dynamics in developing countries (Tanaka et al., 2010) and in the context of health care spending (Hall and Jones, 2007).

In the literature that seeks to estimate discount rates, experimental studies are by far the most numerous. Amongst the 42 studies surveyed by Frederick et al. (2002), 34 use experimental methods. A typical approach is for individuals to be offered a menu of (real or hypothetical) choices between obtaining a quantity of money now or a different quantity of money at some point in the future. Their choices can be used to estimate a discount rate. Our paper fits into a much smaller literature that estimates discount rates using field data on aspects of behaviour – consumption trajectories in our case. We investigate the association between time preference and both the level of education and numerical ability. It is hard to extract robust empirical findings from this literature. Lawrance (1991), using panel data on food consumption, finds that those with more education are more patient though Carroll (1997) critiques her approach and Gourinchas and Parker (2002) find the opposite.

We estimate discount rates using a lifecycle model of consumption and saving. A typical approach in estimating preference parameters in such models, though not the one that we will take, has been to solve numerically the intertemporal optimisation problem that the agents in a particular population are assumed to face. Estimates of parameters such as the discount rate are chosen such that the model’s predictions are close, in some metric and according to some data, to those seen in reality. Such studies vary in the extent to which heterogeneity in the discount factor is admitted into the model. At one extreme there are those papers that assume homogenous discounting behaviour. Examples are French (2005) where discounting is exponential and Laibson et al. (2007) where discounting is quasi-hyperbolic. More flexibility was allowed by Attanasio et al. (1999) who estimate a version of the lifecycle model where the discount rate varies stochastically with the composition of the household. A more flexible approach again (Samwick, 1998; Gustman and Steinmeier, 2005) allows each household to have its own (time-invariant) discount factor. Both these papers solve the intertemporal optimisation problem faced by each household and estimate the discount rate as that rate which sets equal assets predicted in the model and assets observed in the data.

Our approach is to estimate features of the distribution of the discount rate using the Euler equation (the first order condition of the life-cycle model of consumption and saving) and panel data on household consumption.<sup>1</sup> Such an approach has typically been precluded before by the absence of good quality panel data on consumption – a problem discussed in detail by Browning et al. (2003). Our data allows us to calculate longitudinal observations on consumption using comprehensive and high quality data on assets and income and the intertemporal budget constraint. The resulting distribution of consumption is shown to be remarkably similar to that derived from the UK’s household budget survey. Browning and Leth-Petersen (2003) carry out a similar exercise using Danish data on assets and income and also conclude that the practice of generating consumption in this way has promise.

We find that households with lower levels of education and lower levels of numerical ability tend to

---

<sup>1</sup>This empirical strategy has some parallels with a literature that estimates the coefficient of relative risk aversion) using the Euler equation (e.g. Campbell and Mankiw, 1989; Attanasio and Weber, 1995). A lively debate has taken place on whether it is possible to obtain consistent estimates of preference parameters in this manner. Carroll (2001) has argued that it is not possible due to approximation errors induced by the typical approach – which involves working with a linear approximation to the Euler equation. A number of other authors (e.g. Attanasio and Low, 2004; Alan et al., 2006) argue the contrary. As will be outlined in the next section, we do not take a linear approximation to the Euler equation, but work with its exact, non-linear form.

have lower discount rates (i.e. they tend to be more patient) than those who are more educated and have greater levels of numerical ability, respectively. This result is contrary to that found in the literature using laboratory experiments and field studies that measure the extent to which individuals discount future income streams (see for instance Warner and Pleeter (2001), Harrison et al. (2002) and Dohmen et al. (2010)). These papers differ in their empirical approach – the first uses data on the choices of departing military personnel over whether they will take their severance payment in a lump-sum or in the form of an annuity payment, while the second and third papers use laboratory experiments. However, they have an important feature in common – the discount rate is estimated using choices between obtaining a sum of money now or a different sum of money in the future.

This paper continues by outlining our empirical approach in Section 2. In Section 3 we describe the data and explain how we calculate consumption from assets and income. That Section also illustrates the correspondence between the distribution of consumption so derived and the distribution of consumption observed in the UK’s Household Budget Survey. Results are presented in Section 4 and Section 5 concludes.

## 2 Theory and empirical approach

In our estimation of discount rates, we start from a fairly standard life-cycle model in which each household (as a collective unit) maximises expected discounted utility by choosing their consumption and their holdings of each of  $J$  different asset or debt instruments each period. In period  $t$ , household  $i$  faces the following optimisation problem:

$$\begin{aligned} \max_{\{X_{is}^j, c_{is}\}_{s=t}^T} \quad & u_t(c_{it}) + \sum_{\tau=t+1}^T \left( \prod_{s=1}^{\tau-t} \frac{1}{1 + \rho_{it+s}} \right) E[u_\tau(c_{i\tau})] \\ \text{s.t.} \quad & (i) \quad p_\tau c_{i\tau} + \sum_j p_{\tau+1}^j X_{i\tau+1}^j = e_{i\tau} + t_{i\tau} + \sum_j r_{i\tau+1}^j X_{i\tau}^j + \sum_j p_{\tau+1}^j X_{i\tau}^j \quad \forall \tau \end{aligned} \quad (1)$$

$$(ii) \quad X_{i\tau+1}^j \geq b_{i\tau+1}^j \quad \forall \tau, j \quad (2)$$

where  $\rho_{it}$  is the discount rate for household  $i$  between period  $t$  and  $t+1$ . Equation (1) is the intertemporal budget constraint and equation (2) represents a borrowing constraint for asset  $j$  –  $b^j$  is the minimum level of that asset that must be held. This will be negative for debt instruments that households have access to and zero for non-debt instruments.<sup>2</sup> The other quantities in the model are consumption ( $c_t$ ), the instantaneous utility function in period  $t$  ( $u_t(\cdot)$ ), holdings of each of  $J$  assets ( $X_t^j$ ) which will take negative values in the case of debt instruments, the nominal income yield of asset  $j$  ( $r_t^j$ ), the price of asset  $j$  ( $p_t^j$ ), labour income ( $e_t$ ), income from transfers ( $t_t$ ) and the price of consumption ( $p_t$ ).

A standard result in the asset-pricing literature is that optimal consumption will satisfy an Euler equation (first order condition) for every asset that households can potentially hold (see for example Campbell (2000)). That is, for each asset  $j$ , and for each pair of consecutive periods  $t$  and  $t+1$ , the following inequality holds:

---

<sup>2</sup>Alternatively, we could specify a liquidity constraint that ensures that total debts are no greater than a certain quantity.

$$\frac{du_t(c_{it})}{dc_{it}} \geq \frac{1}{(1 + \rho_{it+1})} E \left[ \frac{(p_{t+1}^j + r_{t+1}^j)}{p_t^j} \frac{p_t}{p_{t+1}} \frac{du_{t+1}(c_{it+1})}{dc_{it+1}} \right] \quad (3)$$

The Euler equation will bind with equality for household  $i$  as long as they are not liquidity-constrained with respect to asset  $j$  (i.e. as long as  $X_{it+1}^j > b_{it+1}^j$ ). This equation will hold, in particular, for the safe asset (say  $j = 0$ ), for which  $p_t^0 = 1$  for all  $t$ . Optimal consumption for households who hold positive cash balances will satisfy:

$$\frac{du_t(c_{it})}{dc_{it}} = \frac{1}{(1 + \rho_{it+1})} E \left[ (1 + r_{t+1}^0) \frac{p_t}{p_{t+1}} \frac{du_{t+1}(c_{it+1})}{dc_{it+1}} \right] \quad (4)$$

To identify the discount rate between periods  $t$  and  $t + 1$ , we need to impose the restriction that the utility function doesn't change between the two periods – that is that  $u_t(\cdot) = u_{t+1}(\cdot)$ . The length of a period in our application is two years – the length of time between survey waves of our (panel) data. All features of the model other than the instantaneous utility function are allowed to vary freely over time – in particular we do not assume that the discount rate is time-invariant. Therefore, we do not rule out discounting behaviour that departs from the standard model which involves a discount rate between consecutive periods that doesn't vary with time. For example, quasi-hyperbolic discounting, proposed by Laibson (1997) can be admitted as a special case of the model we present here.

We specify the utility function as taking the familiar isoelastic form:

$$u(c) = \frac{c^{1-\gamma}}{1-\gamma} \quad (5)$$

where  $\gamma$  is the coefficient of relative risk aversion and continue to assume that liquidity constraints do not bind (suppressing  $i$  subscripts):

$$\rho_{t+1} = E \left[ (1 + r_{t+1}^0) \frac{p_t}{p_{t+1}} \left( \frac{c_t}{c_{t+1}} \right)^\gamma \right] - 1. \quad (6)$$

This equation forms the basis for our empirical approach. We group households according to, separately, their level of education and numerical ability and estimate the expectation in the equation above by using the sample average of the quantity in square brackets among households of a particular group.<sup>3</sup> Before taking equation (6) to the data, we need to specify an interest rate on cash and a coefficient of relative risk aversion. We now discuss each of these in turn.

Figure 2 shows the nominal pre-tax rate of return on two types of cash assets in the UK between 2002 and 2009 – instant access savings and time deposits. Until the large fall in in the last quarter of 2008, interest rates were relatively stable over this period, moving within a range of approximately a percentage point. In our estimation of the discount factor using equation (6) we use a nominal pre-tax interest rate of 3% – approximately the average rate of return on time deposits over the period. Our headline results refer to the consumption changes over the period 2004 to 2006 and will therefore be unaffected by the large fall in interest rates in 2008. With some exceptions, interest is taxable in the UK – and for each household we calculate their marginal tax rate. Therefore, we convert this pre-tax interest rate to a post-tax interest

---

<sup>3</sup>We take the education level and the level of numerical ability of a household to be the maximum level held by either member of the couple at the head of that household. We postpone further details on the categorisation of individuals into education and numeracy groups to the next section.

rate using the marginal rate of tax faced by the household in a particular year. For couples who face different marginal tax rates, we use the lower of the two rates, on the basis that efficient tax-planning in most cases will allow the couple pay that lower rate of tax on their asset income.

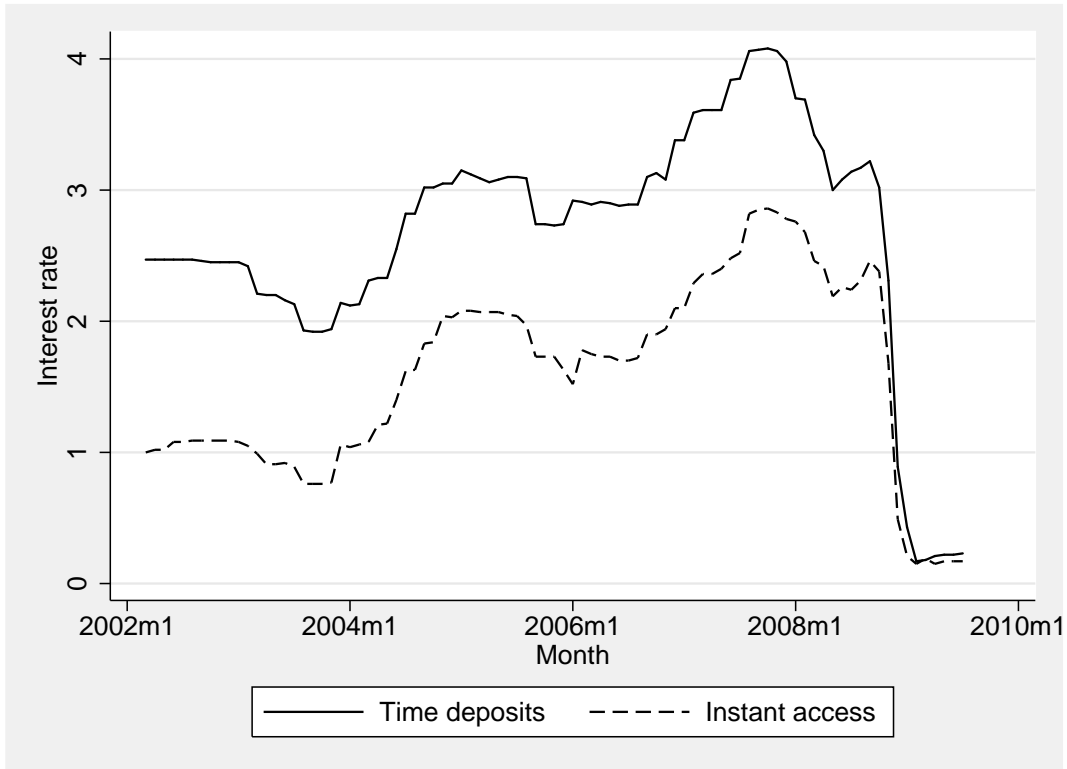


Figure 1: Nominal pre-tax rate of return on cash - 2002 to 2009

Source: Bank of England

In our empirical application we assume a homogenous coefficient of relative risk aversion ( $\gamma$  in equation (6)). The value we use in our estimation is 1.25 (consistent with the elasticity of intertemporal substitution as estimated by Attanasio and Weber (1993) using UK data and very close to that estimated by Gustman and Steinmeier (2005) in their estimation of the distribution of discount rates. We have generated results assuming alternative values of  $\gamma$ . While the mean discount rates that we estimate are sensitive to this choice, the ranking of households' discount rates (and therefore the direction of the association between average discount rates and education and numerical ability) is the same for all positive choices of  $\gamma$ .

Before proceeding, it is worth making explicit three restrictions implicit in the use of equation (6). These relate to liquidity constraints, household composition, and the potential non-separabilities between consumption and leisure in the utility function.

First, recall that the Euler equation (6) will only hold with equality when individuals aren't liquidity constrained. If we use equation (6) to estimate the discount rate for a group of individuals containing those who are liquidity constrained, then the estimate will be biased downwards. Concerns about the presence of liquidity constraints in our case are mitigated by the fact that the population covered by our data is those over the age of 50 and are at a point in the lifecycle where most have accumulated some liquid wealth – over 95% of our sample have positive gross liquid asset holdings. While the fact that a household has positive liquid wealth cannot be taken to conclusively show that they are not liquidity

constrained, it is certainly indicative of that fact. To assess whether our results are being driven by liquidity constraints we have confirmed that there are no substantial changes in focussing only on those with liquid assets above a certain minimum level.

Second, where individuals leave or join a household between periods the assumption of a constant instantaneous utility function at dates  $t$  and  $t + 1$  is no longer acceptable. As a result we do not include households whose composition changes in our estimating sample.

Finally, the utility function (equation (5)) that we have specified is quite restrictive. It does not depend on leisure. We could, without altering the Euler equation or affecting our empirical strategy, augment the utility function with an *additive* term representing the utility of leisure. However, a utility function in which consumption and leisure enter in a nonseparable manner would require changes to the approach. For example, if the utility function was:

$$U(c, l) = \frac{c^{1-\gamma}}{1-\gamma} g(l) \quad (7)$$

where  $g(l)$  is some function of leisure, then the Euler equation would be:

$$\rho_t = E \left[ (1 + r_{t+1}^0) \frac{p_t}{p_{t+1}} \left( \frac{c_t}{c_{t+1}} \right)^\gamma \left( \frac{g(l_{t+1})}{g(l_t)} \right) \right] - 1. \quad (8)$$

Given that separability of consumption and leisure is a strong, though not uncommon, assumption, we exclude from our sample any households where some member changed their labour supply (and therefore their time spent at leisure) in the period between waves of our data.

The previous two paragraphs have outlined two exclusions from our estimating sample. Some further exclusions are necessary due to the fact that we are not able to calculate consumption satisfactorily for every household in our sample. The extent of these exclusions is outlined in the Appendix. To deal with the fact that those omitted are unlikely to be a random sub-sample of our overall sample, we generate weights representing the probability of each household being included in our sample and, in our results we attach a weight to each household of the inverse of this probability. These probabilities are estimated as functions of marital status, education, age, income quintile and wealth quintile. Our results will, therefore, be representative of the entire population aged 50 and over if the selection into our sample can be adequately modelled as a function of these characteristics.

### 3 Data

Our data comes from the English Longitudinal Study of Ageing (ELSA). ELSA is a panel survey that is representative of the English population aged 50 and over. It started in 2002, and individuals have been re-interviewed every two years since then. The purpose and form of the survey is similar to the Health and Retirement Study (HRS) in the US and the Survey of Health, Ageing and Retirement in Europe (SHARE) in 20 European countries. In this paper we use the first four waves of ELSA, the data for which was collected in 2002/03, 2004/05, 2006/07 and 2008/9.

While ELSA contains questions on some components of expenditure (food, domestic fuel and clothing) it does not, unfortunately, contain data on total consumption, which is needed to estimate the discount rate. In fact there is no nationally-representative longitudinal survey that collects total consumption in the



UK and such data is rare internationally.<sup>4</sup> This lack of comprehensive longitudinal data on consumption has proved something of an obstacle to the practice of bringing Euler equations to data. The literature that has done this has either relied on aggregate data, or, following Browning et al. (1985), has used repeated cross-sections of consumption data to form a quasi-panel of birth cohort-level average consumption.

We proceed in a different manner; we use the extremely rich data on income and assets that is contained in ELSA to back consumption out from the intertemporal budget constraint. The rest of this section summarises this procedure – further details are given in the Appendix – and shows a close correspondence between features of the resulting distribution of consumption with those that are obtained using the UK’s Household Budget Survey.

An alternative approach to the creation of panel data on expenditure was suggested by Skinner (1987) and refined by Blundell et al. (2008). It involves estimating the relationship between food (and possibly other items of) expenditure and total expenditure using a household budget survey. As general-purpose panel surveys often contain data on food expenditure this estimated relationship can be used to impute total expenditure.

### 3.1 Calculating consumption using longitudinal data on assets and income

We use longitudinal data on assets and income along with the budget constraint to calculate consumption between two waves. Equation (1) can be re-arranged to get the value of consumption in period  $t$  as follows:

$$p_t c_t = e_t + t_t + \sum_j r_{t+1}^j X_t^j + \sum_j p_{t+1}^j (X_t^j - X_{t+1}^j) \quad (9)$$

The timing convention and how it relates to the data deserves some discussion. In ELSA interviews take place approximately every two years. So to be precise:

- Flow variables  $c_t$ ,  $e_t$ ,  $t_t$ ,  $r_t^j$  represent consumption, non-capital income, transfers and the asset yield over the entire two year period
- Stock variables  $X_t^j$  represent holdings of assets at the beginning of the period (i.e. at the time of the first interview);  $X_{t+1}^j$  represents asset holdings at the beginning of period  $t + 1$  or equivalently at the end of period  $t$  (i.e. at the time of the next interview)
- The asset prices  $p_t^j$  and  $p_{t+1}^j$  represent asset prices at the time of the first and second interview
- The overall price level  $p_t$  represents the average price level in the period between the two interviews.

To make precise the data requirements let us re-write equation (9) as:

$$p_t c_t = e_t + t_t + \sum_j r_t^j X_t^j + \sum_j \left( \frac{p_{t+1}^j}{p_t^j} p_t^j X_t^j - p_{t+1}^j X_{t+1}^j \right) \quad (10)$$

---

<sup>4</sup>A notable exception to this is the Spanish ‘Encuesta continua de presupuesto familiares’, a diary-based longitudinal survey of consumption. The Panel Study of Income and Dynamics in the US, since 1999, has collected consumption data that covers approximately 70% of average total consumption.

This is the equation that we will use to calculate consumption between the two waves. Having calculated consumption in this manner we will make one further adjustment and subtract mortgage repayments (both capital components and interest) from the resulting quantity. While these represent cash expenditure on housing, they are not necessarily indicative of consumption of the flow of housing services.

Some, but not all of the quantities on the right-hand-side of equation (10) can be directly read from the ELSA data. In each wave the value of each asset held ( $p^j X^j$ ) is recorded, as are non-capital income ( $e$ ), capital income ( $r^j X^j$ ) and lump-sum transfers ( $t$ ) in the period prior to the interview (where ‘period’ in the case of most forms of income and transfers represents 12 months). The gaps that we face are data on capital gains on assets held between the two waves ( $\frac{p_{t+1}^j}{p_t^j}$ ) and data on income and transfers for a period of approximately one year (recall that ELSA sample members are surveyed approximately every two years). The majority of assets in this population in the UK are held in safe forms – so there is no capital gain to be considered for these assets. For equity holdings we assume a capital gain (or loss) in line with the change in the FTSE index between the two interview dates. Estimating income in the missing year is facilitated by exploiting the longitudinal aspect of the survey data – we interpolate linearly between income in year  $y$  and income in year  $y + 2$  to obtain income in year  $y + 1$ . Finally, we assume that there are no lump-sum transfers in the missing year and we exclude from our sample those households where it is likely that some member received a lump-sum transfer (due to retirement or redundancy or the death of a spouse or parent). We give further details about all of these assumptions and their implications in the Appendix.

### 3.2 Comparing consumption in ELSA and in the EFS

In this section, we compare the distribution of consumption estimated in the manner described above with that estimated using the Expenditure and Food Survey (EFS).<sup>5</sup> The EFS is the UK’s household budget survey, and is used to calculate the commodity-weights that form the basis of the calculation of the UK’s inflation indices (the Consumer Prices Index and the Retail Prices Index). The data is collected annually, throughout the year and the survey is designed to be nationally-representative. Respondents are asked to record all purchases over a two week period in a diary and also to complete a questionnaire that seeks information on infrequently-purchased items. The combination of the diary and the questionnaire allows a comprehensive measure of consumption to be calculated.


Figure 2 shows the cumulative distribution function and the probability density function of total consumption in both surveys. The data shown is for calendar year 2003 for the EFS and for (annualised) calculated consumption between the surveys in 2002/03 and 2004/05 for ELSA. The EFS functions are estimated using only those households where the head is aged over 50 so that both samples are drawn from populations with the same age profile. Both distributions are shown net of mortgage repayments. As discussed in the Appendix, for a proportion of the ELSA sample, no consumption calculation is possible. These largely represent households for whom we have incomplete data on assets and income. Therefore, in our estimation of the distribution of consumption in ELSA we weight each observation by the inverse of the probability of being able to calculate consumption, where this probability is estimated using a probit on dummies representing age group, marital status, education, income quintile and wealth quintile. In

---

<sup>5</sup>This survey has, since 2008, been known as the Living Costs and Food Survey. However, the data that we show is from years prior to this, so we make use of the older name

both surveys we trim the most extreme values – showing the middle 80% of the distribution.

Figure 2: CDF and PDF of consumption in EFS and ELSA – 2003



The figure shows that there is a close correspondence between the distributions in both shape and location. The correspondence is closest at the bottom of the distribution (i.e. up to annual consumption of £10,000). At this point, the distributions diverge somewhat – with the distribution of consumption in ELSA lying to the right of that in the EFS. This divergence, which represents a tendency for consumption to be greater in ELSA than the EFS in the upper half of the distribution is consistent with the fact that consumption in the EFS (grossed up to national levels) is known to under-record the level of consumption calculated as part of the National Accounts with the degree of under-reporting thought to be greater for those who have higher levels of consumption (see Brewer and O’Dea (2012)). In light of this fact, it is somewhat reassuring that our estimated cumulative distribution function of consumption in ELSA lies to the right of that in the EFS.

Figures 3, 4 and 5 probe the comparison between the distribution of consumption in each survey a little more deeply. The figures each take a particular household characteristic (age, education and marital status respectively) and compare the conditional distributions of consumption in both surveys. We show the 25th, 50th and 75th percentiles as well as the mean. Of interest are both the difference between quantiles for a household of a particular type (for example, whether the shape and location of the distributions match for young people) and the differences for a particular quantile across household types (for example, whether the relationship between median expenditure and age is similar in both surveys).

These figures show that many of the points we made above in our comparison of the unconditional distributions of consumption between surveys are true for the distributions conditional on (at least these) household characteristics. For all age groups, for those with middle and higher levels of education and

for all marital statuses, the 75th percentile of consumption is higher in ELSA than in the EFS, and in most cases the median and mean are higher too, with smaller differences to be seen between the 25th percentiles in each survey. The socio-economic gradients observed in the EFS are closely replicated in the ELSA consumption data – that is consumption decreases with age, rises with education and married households consume more than single households (unsurprisingly as our measure of consumption here is not adjusted using an equivalence scale).

Figure 3: Comparing consumption in EFS and ELSA – by age – 2003

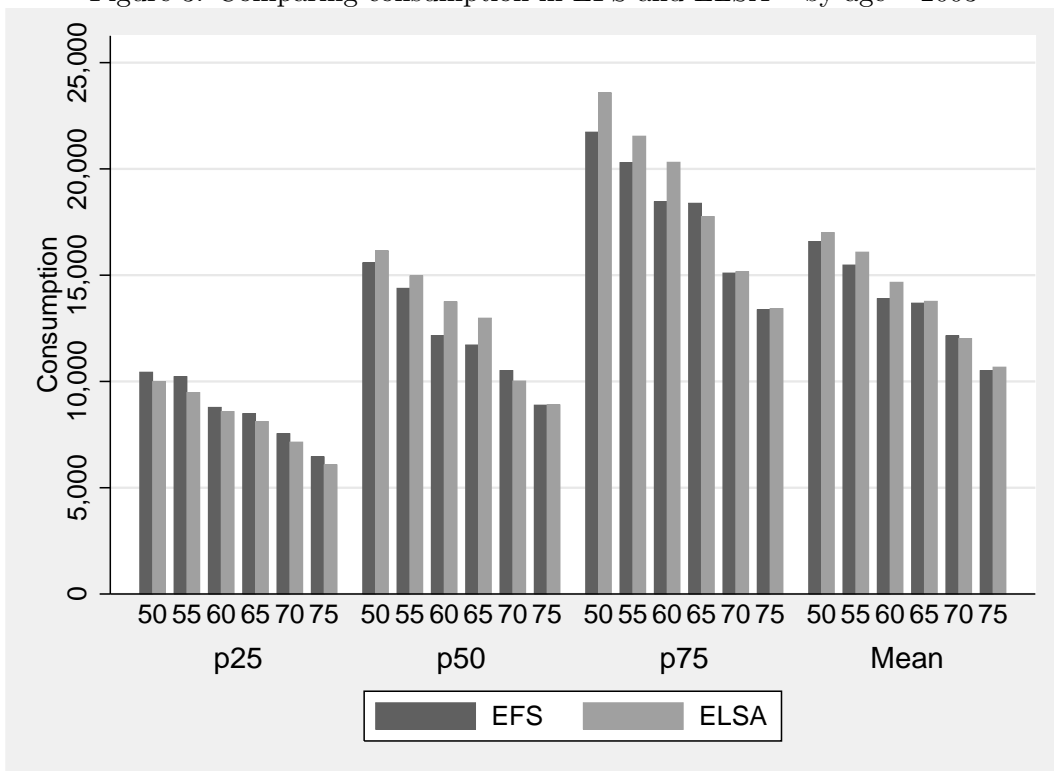


Figure 4: Comparing consumption in EFS and ELSA – by education – 2003

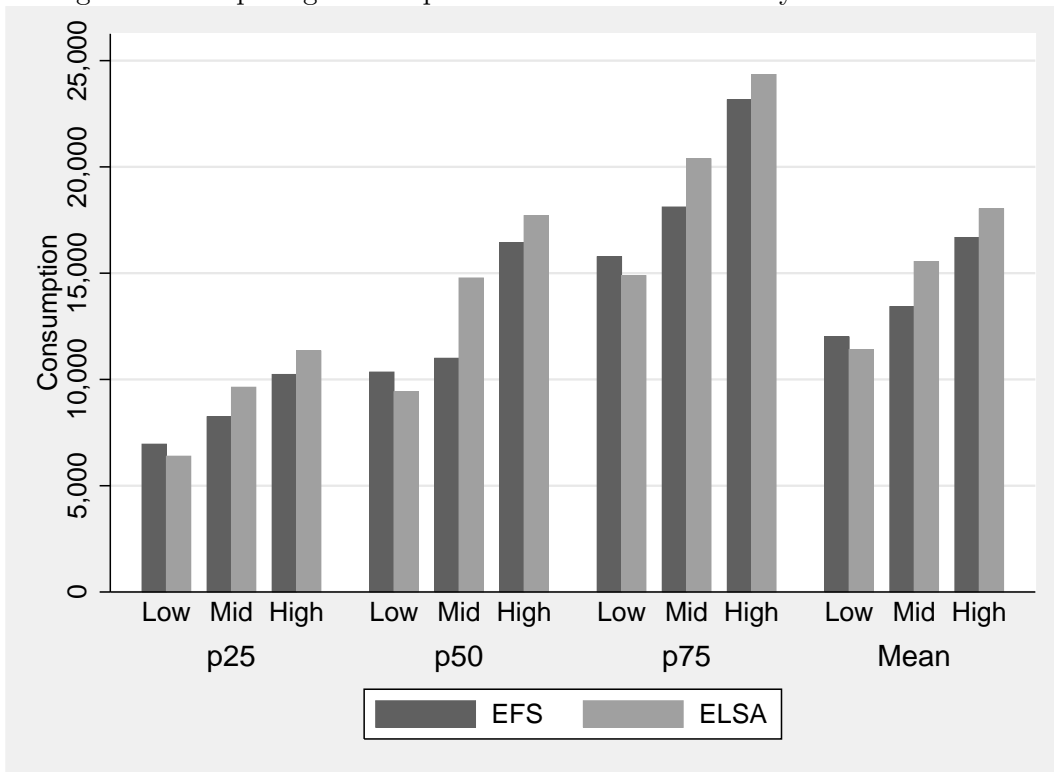
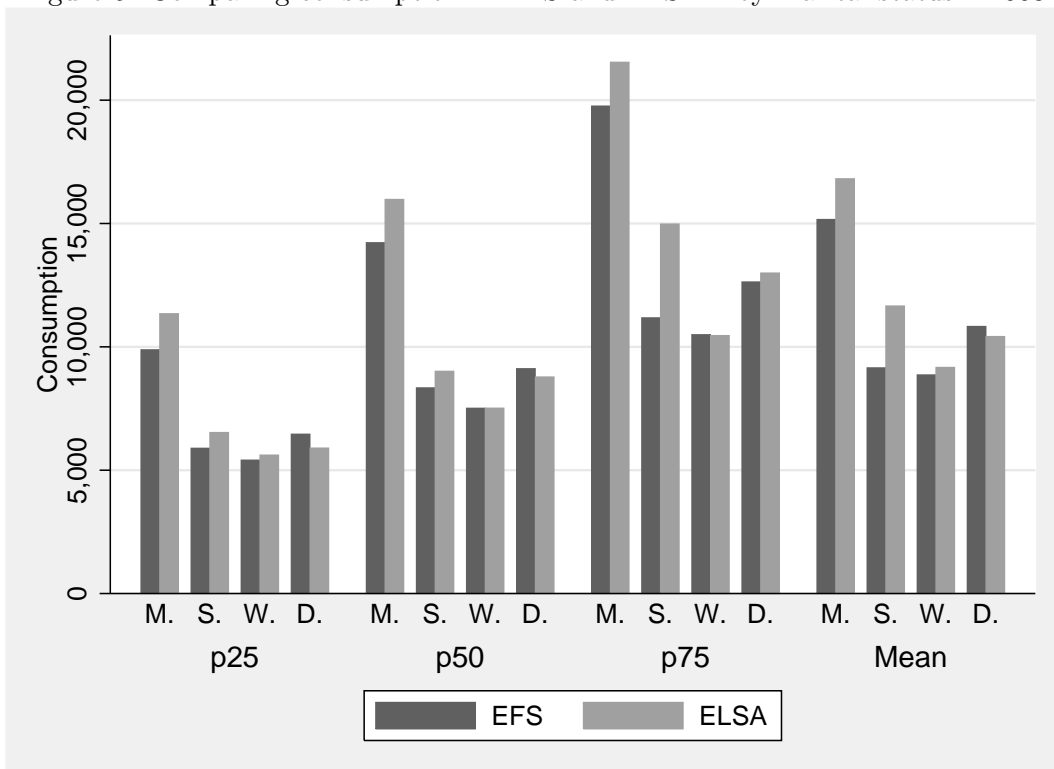



Figure 5: Comparing consumption in EFS and ELSA – by marital status – 2003



Notes: The labels on the horizontal axis mean as follows: M. – Married (including cohabitating); S. – Single, never married; W. – Widowed; D. – Divorced/Separated.

As a final comparison of our generated data on consumption with that in the EFS we make use of the fact that data on food spending is recorded in ELSA. In Figure 6 we plot the relationship between food spending and total spending in both surveys (again trimming the bottom and top 10% of consumption). The relationships shown are estimated using locally-weighted regressions. Food spending is estimated as greater in ELSA than in the EFS<sup>6</sup>; the slope of the relationship between it and total spending is, however, similar in both surveys from the point at which total consumption is equal to approximately £10,000. Below this level, food spending in ELSA doesn't vary much with our calculated measure of total consumption – an indication, perhaps, that some of the households in the left-tail of the distribution of consumption are there due to measurement error.

Figure 6: Comparison of relationship between food spending and consumption in EFS and ELSA - 2003



### 3.3 Data on education and numerical ability

In the next section, we will show how our estimated discount rates vary with education and numerical ability. We here outline how our measures of these variables are defined.

We categorise individuals into one of three education groups on the basis of the highest qualification that they have. We consider those who have a third-level degree or higher to have a ‘high’ level of education. Those who have A-levels (British school-leaving exams, taken at age 18) or equivalent but no university degree are in the ‘mid’ education group. All others are in the ‘low’ group. Our measure of consumption, and therefore our estimates of discount rates, are at the household (formally benefit unit<sup>7</sup>)

<sup>6</sup>One potential reason for this is the different methods by which the data on food spending is gathered. The ELSA data is taken from respondents’ responses to a question that asks how much they spend on food in a typical month. The EFS data is taken from respondents’ spending diary entries.

<sup>7</sup>A benefit unit is a single adult or couple along with any dependent children that they have. Relatively few members of

level. We need an education measure, therefore, at the household level. We take the education of a household containing a couple to be the greater of the two levels of education held by the adults in that couple.

Numerical ability is measured in ELSA using a series of six questions, which are reproduced in Box 1. The simplest of these questions requires the respondent to solve a very simple exercise in subtraction while the most difficult requires the respondent to solve a problem involving compound interest. We divide all individuals into one of four groups following the categorisation in Banks et al. (2010). Mirroring our approach with respect to education we take the numerical ability of a household containing a couple to be the greater of the two levels of numerical ability held by the two members of that couple.

**Box 1 Numerical Ability in ELSA**

1. If you buy a drink for 85 pence and pay with a one pound coin, how much change should you get?
2. In a sale, a shop is selling all items at half price. Before the sale a sofa costs £300. How much will it cost in the sale?
3. If the chance of getting a disease is 10 per cent, how many people out of £1,000 would be expect to get the disease?
4. A second hand car dealer is selling a car for £6,000. This is two-thirds of what it cost new. How much did the car cost new?
5. If 5 people all have the winning numbers in the lottery and the prize is £2 million, how much will each of them get?
6. Let's say you have £200 in a savings account. The account earns ten per cent interest per year. How much will you have in the account at the end of two years?

## 4 Results

We first summarise the distribution of the quantities represented by the right-hand side of equation (6):

$$\left[ (1 + r_{t+1}^0) \frac{p_t}{p_{t+1}} \left( \frac{c_t}{c_{t+1}} \right)^\gamma \right] - 1 \quad (11)$$

This quantity would be equal to the discount rate if  $c_{t+1}$  was perfectly forecasted by households at date  $t$ . Our presentation of the distribution of this quantity which we refer to below as the ‘ex-post’ discount rate is a useful preliminary step.

Our use of four waves of ELSA data gives us up to three observations on consumption for each household and therefore up to two observations on the ex-post discount rate. Figure 7 shows two distributions of the ex-post discount rates (trimming the bottom 10% and top 10% of the sample). The median discount rate is approximately -3% in the earlier period and is 0% in the later period. These median ex-post discount rates are low relative to estimates of the discount rate found in the literature that estimate such rates using field data and very low relative to those found in the experimental literature. The figure, though, shows substantial heterogeneity around these medians. This observed heterogeneity is likely our sample have dependent children, so our results can be thought of as representing single adults and couples.

driven by a number of phenomenon.

First, there is likely to be differences in the degree of patience both between individuals and over time for the same individual (recall that our model and estimation procedure allows for a discount rate that changes over time). Secondly, some households may be making expectational errors in which case the distribution of ex-post discount rates, shown below, will differ from the distribution of true (or ‘ex-ante’) discount rates. Thirdly, our consumption data is likely to include some measurement error, which will cause a divergence between the true distribution of the ex-post discount rates and that presented in Figure 7.

Figure 7: Distribution of ex-post discount rates – 2004/06 and 2006/08





Figure 8 adds to the figure of the distributions of the ex-post discount rate in two cross-sections the distribution of their geometric mean.<sup>8</sup> The variance of this distribution is substantially smaller than the variance of either cross-sectional distribution. This could be due to the averaging over time of the, perhaps changing, discount rate for each family and also to a diminished effect of measurement error once we take a time average.

<sup>8</sup>That is we show the distribution of  $\sqrt{(1 + \rho_t)(1 + \rho_{t+1})}$ .



Figure 8: Distribution of ex-post discount rates – 2004/06, 2006/08 and average



To the best of our knowledge, there have been only two papers that have, using field data, attempted to estimate the entire distribution of discount rates. These are Samwick (1998) and Gustman and Steinmeier (2005) (hereafter GS). The distributions shown above can be interpreted as estimates of the entire distribution of discount rates for our population under the assumption of perfect foresight. That assumption is reasonably close to that implicit in the models of both Samwick and GS. The former has uncertainty only over survival while the latter contains only income uncertainty. While income uncertainty will cause a divergence between expected consumption and realised consumption, and therefore between the ex-post discount rate and the true discount rate, its effect will be muted in the population that we study, the majority of whom are retired.

We therefore feel it is worthwhile to compare the results presented so far to those presented in these two papers. It is important to note, though, that we would not necessarily expect a close correspondence between our results and theirs as the populations on which the estimates are based are very different. Our results are for the English population aged over 50, while the results of both Samwick and GS are for the entire US adult population.

In Table 1 we compare our three distributions of the ex-post discount rate (the two cross-sectional distributions, and the distribution of their geometric average) to those found in Samwick and GS. The division of the sample into four groups follows the categorisation found in those papers. The most striking difference between our results and those of Samwick and GS is the substantial number of households that we find with ex-post discount rates of less than 5%. We find approximately 60% here in this portion of the distribution compared to approximately 40% in the distribution of discount rates in the two papers based on the US all-age population.

In our two cross-sectional distributions, we find a similar share of households in the right tail of the distribution (those with a discount rate greater than 15%) as do Samwick and GS, and find less mass in the region of 5% to 15%. On our average measure, the mass in the left tail of our distribution increases, largely at the expense, relative to either cross-sectional distribution, of that in the right tail.

The models of both Samwick and GS assume a discount rate for a particular household that doesn't change over time. If, as is likely, discount rates do vary over time, their estimates will represent some average of the lifetime sequence of discount rates. Therefore, the large mass that we find in the left tail of the distribution could be reconciled with the estimates of Samwick and GS if households have higher discount rates at younger ages than they do at older ages (at which point they enter our population at interest).

Table 1: Comparison of our results with those of Samwick (1998) and Gustmann & Steinmeier (2005)

<b>Discount rate</b>	<b>Samwick</b>	<b>GS</b>	<b>Ours 04-06</b>	<b>Ours 06-08</b>	<b>Ours Ave</b>
<5%	38%	40%	60%	56%	67%
5%-10%	25%	21%	5%	6%	8%
10%-15%	10%	6%	4%	6%	7%
>15%	25%	33%	30%	32%	18%

Table 2 gives estimates of the median ex-post discount rate ( $\hat{\rho}$ ) and associated standard errors of the medians ( $\sigma$ ), for groups defined according to education and numerical ability. We find that less educated families and families with lower levels of numerical ability tend to have lower median rates than those with more education and greater levels of numerical ability respectively.

Table 2: Median ex-post discount rate by education and numerical ability

<b>Education</b>	$\hat{\rho}$	$\sigma$	<b>Numerical ability</b>		
				$\hat{\rho}$	$\sigma$
Low	-3.4	1.0	1 (Lowest)	-2.9	2.0
Mid.	-1.8	2.3	2	-3.4	1.1
High	5.7	5.7	3	-0.8	2.7
			4 (Highest)	-1.3	4.0
All	-2.5	1.0	All	-2.3	1.0

The number of households in each of the education groups are: 749, 319, 133. The number of households in each of the four numeracy groups are: 150, 559, 346, 139. The median in the 'All' row differs slightly between columns as the number of households in each differs. In calculating the 'All' group median we exclude those with missing values of the covariate in question.

We now use a grouping estimator to estimate the average true (or ex-ante) discount rate for groups defined by levels of education and numerical ability. We have also estimated the average rate according to the age of the household and their marital status and find no robust differences between groups in either of those two cases. The estimator is based on equation (6). It estimates the expectation in that equation by its sample analogue for a particular group and weights the results to account for possibly non-random selection into our sample.

Table 3 summarises these results. For each group we trim the sample, removing those in the first and tenth decile of consumption growth. The relationship between education and numerical ability and the ex-ante discount rate is similar to that found with the ex-post rate – those with the less education are more patient than those with more education, and patience decreases with numerical ability.

Table 3: Mean ex-ante discount rate for groups defined according to education and numerical ability

<b>Education</b>	$\bar{\rho}$	$\sigma$	<b>Numerical ability</b>		
				$\bar{\rho}$	$\sigma$
Low	-2.5	0.7	1 (Lowest)	-3.5	1.4
Mid.	0.9	1.4	2	-2.3	0.8
High	6.7	3.0	3	1.7	1.5
			4 (Highest)	2.0	2.4
All	-1.0	0.6	All	-1.0	0.6

The sample sizes in each category are smaller here than in Tables 2 as, in each category we trim the top and bottom 10% of values. The number of households in each of the education groups are: 814, 361, 166. The number of households in each of the four numeracy groups are: 153, 607, 398, 177.

The magnitude of the differences between groups is large. For example, the difference between the point estimates of the mean discount rate for those with only compulsory schooling and those in the ‘mid’ education group is 6 percentage points. This means that, over this period and in this population, if those with more than compulsory schooling are to exhibit the same saving behaviour at the margin as those with only compulsory schooling, the former group will require a safe return of 6 percentage points greater than the latter group.

Table 4 investigates the joint association between average discount rates, education and numerical ability. Here those in numeracy groups 1 and 2 are categorised as having ‘low’ numerical ability and those in numeracy groups 3 and 4 are categorised as having ‘high’ numerical ability. The ‘low’ education group is defined as before, while the ‘med./high’ education group contains the upper two categories. The gradient of the association between education and average discount rate, conditional on level of numerical ability is particularly large. Among those with low levels of numerical ability, the average discount rates among the low education group is estimated at -3%, compared to 2.7% for those in the mid./high education group. For those with more numerical ability the differences according to education are starker – with average discount rates of -2.2% and 5% for those with less and more education respectively. There seem therefore to be separate negative associations between each of education and numerical ability and the discount rate.

Table 4: Mean ex-ante discount rate for groups defined by pairwise combinations of both education and numerical ability

	<b>Low education</b>	<b>Med./High education</b>
<b>Low numeracy</b>	-3.1 (0.8)	0.6 (1.8)
<b>High numeracy</b>	-1.7 (1.5)	4.7 (1.9)

Standard errors are in parentheses.

Tables 5 and 6 explore the robustness of our estimated mean discount rates for groups defined by

education and numerical ability respectively. Column 1 in each table is associated with less trimming than in our headline results. We trim those in the bottom and top 5% of the distribution of consumption growth instead of those in the bottom and top 10%. Column 2 shows the mean discount rate for the second of our two periods of data (2006 - 2008). Columns 3 and 4 apply successively stricter sample selection rules than are applied in our baseline sample. These are the ‘middle’ and ‘strict’ sample selection rules outlined in the Appendix. In each case the results that we previously emphasised still hold – the estimated mean discount rates are higher for those with more education and numerical ability than those with less. However, some of the results are suggestive of a non-monotonic association between education and the discount rate. Taking the two sets of results that apply a stricter sample selection, the lowest estimated discount rates are for those with the least education, the highest are for those in the middle group, with those with the most education having an intermediate value.

Table 5: Sensitivity of mean ex-ante discount rate for groups defined according to education

	(1)		(2)		(3)		(4)	
<b>Education</b>	$\bar{\rho}$	$\sigma$	$\bar{\rho}$	$\sigma$	$\bar{\rho}$	$\sigma$	$\bar{\rho}$	$\sigma$
Low	-1.1	0.9	-0.5	0.8	-1.9	0.8	-2.6	1.2
Mid.	3.8	1.8	1.0	1.5	1.9	1.7	3.2	2.5
High	10.5	3.6	12.2	3.4	2.2	4.1	-1.5	6.9
All	2.3	0.8	1.6	0.7	-0.9	0.8	-1.3	1.2

Table 6: Sensitivity of mean ex-ante discount rate for groups defined according to numerical ability

	(1)		(2)		(3)		(4)	
<b>Numerical ability</b>	$\bar{\rho}$	$\sigma$	$\bar{\rho}$	$\sigma$	$\bar{\rho}$	$\sigma$	$\bar{\rho}$	$\sigma$
1 (Lowest)	-1.5	1.7	0.3	1.7	-2.7	1.9	-5.5	2.3
2	-0.2	1.1	1.2	0.9	-1.0	1.0	0.1	1.5
3	7.1	2.0	2.1	1.2	-3.3	1.7	-0.4	2.8
4 (Highest)	4.1	2.9	6.9	2.2	4.1	3.4	6.9	5.6
All	2.3	0.8	2.0	0.6	-0.5	0.8	-0.3	1.2

## 5 Conclusion

This paper uses rich panel data on income and wealth to calculate consumption data for households with individuals over the age of 50. The distribution of consumption that we estimate is shown to closely correspond to that estimated using the UK’s household budget survey. Using this data and a life-cycle model of consumption and saving, we estimate features of the distribution of discount rates.

Relative to other papers that have attempted to estimate discount rates using field data and a lifecycle model our empirical strategy requires fewer assumptions and relies more on the data. The approach that we take allows the discount rate to change over time for a particular household and the population on which we focus in those over the age of 50 in England.

We find substantial heterogeneity in discounting behaviour. Our main result is that households with less education and lower numerical ability exhibit greater patience than, respectively, those with more

education and greater numerical ability. The direction of this association is the opposite to that found in a small number of recent experimental studies. Our results, like those in the experimental studies, are local ones in that they refer to the behaviour of individuals in one country over a short period of time. We further focus on households over the age of 50. Further study would be welcome to investigate the extent and manner of discounting over the lifecycle and in other populations.

## References

- Alan, Sule, Orazio Attanasio, and Martin Browning**, “Estimating Euler Equations with Noisy Data: Two Exact GMM Estimators,” Economics Series Working Papers 283, University of Oxford, Department of Economics 2006.
- Attanasio, Orazio, James Banks, Costas Meghir, and Guglielmo Weber**, “Humps and Bumps in Lifetime Consumption,” *Journal of Business & Economic Statistics*, January 1999, 17 (1), 22–35.
- Attanasio, Orazio P. and Guglielmo Weber**, “Consumption Growth, the Interest Rate and Aggregation,” *The Review of Economic Studies*, 1993, 60 (3), pp. 631–649.
- and —, “Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey,” *Journal of Political Economy*, 1995, 103 (6), pp. 1121–1157.
- and **Hamish Low**, “Estimating Euler Equations,” *Review of Economic Dynamics*, April 2004, 7 (2), 405–435.
- Banks, James, Cormac O’Dea, and Zoë Oldfield**, “Cognitive function, numeracy and retirement saving trajectories,” *The Economic Journal*, Nov 2010, 120 (548), F381–F410.
- Blundell, Richard, Luigi Pistaferri, and Ian Preston**, “Consumption Inequality and Partial Insurance,” *American Economic Review*, December 2008, 98 (5), 1887–1921.
- Brewer, Mike and Cormac O’Dea**, “Measuring living standards with income and consumption: evidence from the UK,” IFS Working Paper 12/12 Jul 2012.
- Browning, Martin and Soren Leth-Petersen**, “Imputing consumption from income and wealth information,” *Economic Journal*, 06 2003, 113 (488), F282–F301.
- , **Angus Deaton, and Margaret Irish**, “A Profitable Approach to Labor Supply and Commodity Demands over the Life-Cycle,” *Econometrica*, 1985, 53 (3), pp. 503–544.
- , **Thomas F. Crossley, and Guglielmo Weber**, “Asking consumption questions in general purpose surveys,” *Economic Journal*, November 2003, 113 (491), F540–F567.
- Campbell, John Y.**, “Asset Pricing at the Millennium,” *The Journal of Finance*, 2000, 55 (4), pp. 1515–1567.
- and **N. Gregory Mankiw**, “Consumption, Income and Interest Rates: Reinterpreting the Time Series Evidence,” in “NBER Macroeconomics Annual 1989, Volume 4” NBER Chapters, National Bureau of Economic Research, Inc, Julio Dic 1989, pp. 185–246.
- Carroll, Christopher D.**, “Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis,” *The Quarterly Journal of Economics*, 1997, 112 (1), 1–55.
- , “Death to the Log-Linearized Consumption Euler Equation! (And Very Poor Health to the Second-Order Approximation),” *The B.E. Journal of Macroeconomics*, 2001, 0 (1), 6.

- Diamond, Peter A. and Johannes Spinnewijn**, “Capital Income Taxes with Heterogeneous Discount Rates,” *American Economic Journal: Economic Policy*, 2011, 3 (4), 52–76.
- Dohmen, Thomas, Armin Falk, David Huffman, and Uwe Sunde**, “Are Risk Aversion and Impatience Related to Cognitive Ability?,” *The American Economic Review*, 2010, 100 (3), pp. 1238–1260.
- Frederick, Shane, George Loewenstein, and Ted O’Donoghue**, “Time Discounting and Time Preference: A Critical Review,” *Journal of Economic Literature*, 2002, 40 (2), 351–401.
- French, Eric**, “The Effects of Health, Wealth, and Wages on Labour Supply and Retirement Behaviour,” *Review of Economic Studies*, 04 2005, 72 (2), 395–427.
- Golosov, Mikhail, Maxim Troshkin, Aleh Tsyvinski, and Matthew Weinzierl**, “Preference Heterogeneity and Optimal Capital Income Taxation,” *mimeo*, 2011.
- Gourinchas, Pierre-Olivier and Jonathan A. Parker**, “Consumption Over the Life Cycle,” *Econometrica*, January 2002, 70 (1), 47–89.
- Gustman, Alan L. and Thomas L. Steinmeier**, “The social security early entitlement age in a structural model of retirement and wealth,” *Journal of Public Economics*, February 2005, 89 (2-3), 441–463.
- Hall, Robert E.**, *Foward-Looking Decision Making*, Princeton University Press, 2010.
- and **Charles I. Jones**, “The Value of Life and the Rise in Health Spending,” *The Quarterly Journal of Economics*, 02 2007, 122 (1), 39–72.
- Harrison, Glenn W., Morten I. Lau, and Melonie B. Williams**, “Estimating Individual Discount Rates in Denmark: A Field Experiment,” *The American Economic Review*, 2002, 92 (5), pp. 1606–1617.
- Laibson, David**, “Golden Eggs and Hyperbolic Discounting,” *The Quarterly Journal of Economics*, May 1997, 112 (2), 443–77.
- Laibson, David I., Andrea Repetto, and Jeremy Tobacman**, “Estimating Discount Functions with Consumption Choices over the Lifecycle,” NBER Working Papers 13314, National Bureau of Economic Research, Inc 2007.
- Lawrance, Emily C.**, “Poverty and the Rate of Time Preference: Evidence from Panel Data,” *Journal of Political Economy*, February 1991, 99 (1), 54–77.
- ONS**, “Wealth in Great Britain- Main Results from the Wealth and Assets Survey 2006/08,” Technical Report, Office for National Statistics 2009.
- Saez, Emmanuel**, “The Desirability of Commodity Taxation under Non-Linear Income Taxation and Heterogeneous Tastes,” *Journal of Public Economics*, 2002, 83 (2), 217–230.
- Samuelson, Paul**, “A Note on Measurement of Utility,” *Review of Economic Studies*, 1937, 4, 155–161.

**Samwick, Andrew A.**, “Discount rate heterogeneity and social security reform,” *Journal of Development Economics*, 1998, *57* (1), 117–146.

**Skinner, Jonathan**, “A superior measure of consumption from the panel study of income dynamics,” *Economics Letters*, 1987, *23* (2), 213–216.

**Tanaka, Tomomi, Colin F. Camerer, and Quang Nguyen**, “Risk and Time Preferences: Linking Experimental and Household Survey Data from Vietnam,” *American Economic Review*, March 2010, *100* (1), 557–71.

**Warner, John T. and Saul Pleeter**, “The Personal Discount Rate: Evidence from Military Downsizing Programs,” *American Economic Review*, March 2001, *91* (1), 33–53.



## Appendix

This appendix gives additional detail on the data that we use, our method for calculating consumption and the derivation of the sampling weights that we use. Section A presents statistics on the measures of wealth and debt in ELSA. Section B details the procedure by which we computed consumption from our panel data on wealth and income. Section C presents the characteristics of the final sample and, addressing concerns about our use of a non-random sub-sample of the population, discusses the construction of survey weights.

### A. Data

The data we utilise in this paper is the English Longitudinal Study of Ageing (ELSA). ELSA is a biennial longitudinal survey of a representative sample of the English household population aged 50 and over (plus their partners). The first wave was conducted between April 2002 and March 2003 and sampled 12,099 individuals (of whom 11,391 were core sample members; the remainder were individuals aged under 50 who were the partners of core sample members). There are 7,894 benefit units (i.e. a single person or couple along with any dependent children) where each member of the couple is a sample member. Our sample is drawn from these benefit units.

ELSA collects a wide range of information on individuals' circumstances. This includes detailed measures of their financial situation: income from all sources (including earnings, self-employment income, benefits and pensions), non-pension wealth (including the type and amount of financial assets, property, business assets and antiques) and private pension wealth (including information on past contributions and details of current scheme rules). ELSA also collects information on individuals' physical and mental health, cognitive ability, social participation and expectations of future events (such as surviving to some older age or receiving an inheritance).

Given the importance of the wealth measure for our estimation, it is worth detailing how it is measured in the survey. ELSA respondents are asked for details on 16 different financial asset and (non-mortgage) debt instruments.<sup>9</sup> For each asset (X), the 'main respondent' in each benefit unit is asked: "How much do you/you and your husband/wife/partner currently have in X?". If the respondent doesn't know or refuses to say, a series of questions is asked that attempts to put lower and upper bounds on these assets. An imputation procedure is then carried out that gives a point estimate for the asset level for these individuals<sup>10</sup>.

Table 7 summaries the holdings in each of these assets. Some are self-explanatory (e.g. cash savings), others are specific to the UK and deserve some comment. Cash ISAs (Individual Savings Account) and TESSAs (Tax-Exempt Special Savings Account) are tax efficient cash savings which are subject to annual limits on what can be paid in. Stocks and shares ISAs and Personal Equity Plans (PEPs) are stocks and shares held in a tax-efficient ISA (Individual Savings Account) and are also subject to annual limits on what can be paid in. Life insurance savings ISA are life insurance savings held in a tax-efficient vehicle.

---

<sup>9</sup>There are also questions on housing wealth, physical wealth and pension wealth.

<sup>10</sup>This imputation procedure, carried out by the ELSA team is called a conditional hot-deck. Given that we know the, say, cash holdings of a particular household are between £a and £b – that individual is assigned a random draw from the empirical asset distribution of those who report their assets exactly as between £a and £b and have the same characteristics along some dimensions: here it is age and household composition.

Bonds could be either savings bonds (with retail banks) or government/corporate bonds. According to the ONS (2009) (Table 4.1) fewer than 2% of households directly hold government or corporate bonds while over 8% hold fixed-term bonds with financial institutions. Therefore households in ELSA who report holdings of ‘bonds’ are more likely to be holding savings bonds (which are effectively fixed-term risk free savings accounts) than gilts or corporate bonds. National savings are cash savings held in the a government-owned agency (‘National Savings and Investment’). Finally, premium bonds are also issued by the government. Instead of yielding interest, the holders of these bonds are included in a monthly draw for large cash prizes.

Table 7: Holdings of different financial assets

	(1)	(2)	(3)	(4)	(5)	(6)
<b>Asset</b>	Mean (Uncond.)	Mean (Cond.)	Med (Cond.)	Proportion with asset	Mean port. share	Prop. unimp.
Cash Savings	12,111	13,474	4,000	90.1%	54.9%	80.4%
Cash ISAs	2,436	7,452	6,000	34.1%	10.0%	91.6%
TESSAs	1,457	9,796	9,000	16.7%	3.5%	95.3%
National savings	832	11,547	3,000	9.2%	1.8%	96.6%
Premium bonds	763	2,373	100	33.6%	2.6%	95.7%
Bonds	2,837	29,425	16,000	11.6%	3.6%	95.2%
Shares	6,650	22,087	3,500	31.6%	7.5%	88.8%
S&S ISAs	1,551	11,982	7,000	14.8%	2.9%	92.3%
PEPs	2,792	18,158	9,000	17.2%	3.7%	93.1%
Invest. trusts	2,379	26,483	12,000	10.9%	2.5%	94.8%
Life ins. savings	2,267	22,470	10,000	12.0%	4.5%	93.2%
Life ins. ISAs	91	9,974	2,000	3.0%	0.2%	95.4%
Other savings	2,179	40,458	15,000	7.4%	2.4%	96.9%
<b>Total</b>	<b>38,346</b>	<b>41,258</b>	<b>12,152</b>	<b>93.1%</b>	<b>100.0%</b>	<b>64.7%</b>

NOTES: Column (1): mean holdings in the asset, *unconditional* on having a positive holding (in GBP); column (2): mean holdings in the asset, *conditional* on having a positive holding; column (3): median holdings in the asset, *conditional* on having a positive holding; column (4): the proportion benefit units that holds this asset; column (5): the mean portfolio share (among those with positive total gross assets); column (6): the proportion of benefit units who report their exact holdings and so for whom no imputation is necessary.

SOURCES: ELSA, wave 1 (2002/03).

The ELSA survey asks also for details on three different types of non-mortgage debt. These are credit card debt, private debt (i.e. debts to friends and family) and ‘other debt’ (primarily overdrafts and personal loans). Table 8 summarises the holdings of different debt instruments and has a form similar to Table 7.

## B. Computing consumption

In this subsection we present some additional details on the consumption calculation procedure. We reproduce equation (10) here, which makes clear the data requirements:

Table 8: Balances of different debts

	(1)	(2)	(3)	(4)	(5)	(6)
<b>Debt type</b>	Mean (Uncond.)	Mean (Cond.)	Med (Cond.)	Proportion with ass.	Mean port. share	Prop. uinimp.
Credit card debt	369	1,989	800	20.3%	41.1%	98.8%
Private debt	68	5,612	1,000	3.3%	2.8%	99.3%
Other debt	846	3,904	1,500	23.3%	56.1%	98.9%
<b>Total</b>	<b>1,293</b>	<b>4,067</b>	<b>1,400</b>	<b>31.8%</b>	<b>100.0%</b>	<b>97.1%</b>

NOTES: See notes to Table 7.

SOURCES: ELSA, wave 1 (2002/03).

$$p_t c_t = e_t + t_t + \sum_j r_t^j X_t^j + \sum_j \left( \frac{p_{t+1}^j}{p_t^j} p_t^j X_t^j - p_{t+1}^j X_{t+1}^j \right) \quad (12)$$

Three issues that we now discuss in turn are:

1. How to estimate capital gains on assets held between the two waves ( $\frac{p_{t+1}^j}{p_t^j}$ )
2. How to estimate income and transfers for the first half of the period between waves. This is necessary as ELSA sample members are surveyed approximately every two years with the questionnaire seeking, in most cases, data on receipts in the past year, rather than on over the entire period between waves
3. Whether to make adjustments to the procedure when asset stocks in either wave are imputed

### B.1 Estimating capital gains

We need to estimate capital gains on assets held between two waves ( $\frac{p_{t+1}^j}{p_t^j}$ ). Depending on the type of asset we make different assumptions.

- i. For cash or most cash-like assets (savings, TESSAs, National savings, Prize bonds) and ‘other savings’, we assume no capital gain ( $\frac{p_{t+1}^j}{p_t^j} = 1$ ).
- ii. For Cash ISAs and Bonds<sup>11</sup>, we have a concern that some respondents, for whom their interest income is simply being rolled up in their account and not withdrawn, will report that their ‘income’ is zero, when in fact it is positive but simply saved. If income from the asset is reported as positive, we assume that there is no change in the value of the asset ( $\frac{p_{t+1}^j}{p_t^j} = 1$ ). If, however, individuals report zero income from their holdings (and approximately 39% of bondholders do in wave 1), we assume that they *are*, in fact, receiving interest and that this interest is simply accumulating in their account and they don’t as a result consider it as ‘income’. We assume that they earn a rate of return equal to the median rate of return for those holding similar assets who *do* report income and reflect this

<sup>11</sup>These are largely bank savings bonds (i.e. effectively fixed term savings accounts) rather than government or corporate bonds.

interest in an increase in the ‘value’ of the asset. That is, for individuals who hold the particular asset but who report no income, we assume that, if the median return on that asset in that wave was 2% then  $\frac{p_{t+1}^j}{p_t^j} = (1.02)^\tau$ , where  $\tau \approx 2$  is the length of time in years between the two interviews.

- iii. For equities and equity-like assets, we make a distinction whether the income from the asset is reported as positive or zero. If income is reported as positive, we assume that the value of the asset increased in line with the FTSE 100 price index (which excludes dividend payments) between the dates of the two interviews. If income is reported as zero, we assume that the value of the asset increased in line with the FTSE 100 *total return* index (i.e including dividend payments) between the dates of the two interviews.
- iv. For debt, we assume that the interest rate on credit card debt is 15%, the interest rate on ‘other debt’ (mostly overdrafts) is 8% and that there is a 0% interest rate on ‘private debt’.

## B.2 Imputing missing income

For each month between the waves where we don’t have income data, we interpolate linearly between the two income observations that we have. We carry out this procedure separately for each category of income (employment income, self-employment income, private pension income, state pension income, benefit income, asset income and other income).

We vary the procedure in two cases. The first of these is when respondents do not report some category of their income exactly in one (but not both) of the waves (i.e. they perhaps only give bounds). In these cases we assume that income has been equal over the period to the value in the year for which we have full information. The second case where we vary the procedure is when it comes to state pension income. Here we use the data on the age of the respondents and the state pension age to establish exactly when their state pension payments started.

## B.3 Whether to make adjustments to the procedure when asset stocks in either wave are imputed

In a number of cases survey respondents don’t know the exact amount of some asset holding or some income component. In that case ELSA questionnaire attempts to obtain bounds on the unknown amounts and exact amounts have to be imputed to carry out our estimation.

When individuals report that they do not know exactly their holdings in any particular wave  $t + 1$  (i.e. they don’t know  $X_{t+1}^j$  for some  $j$ ), we make the assumption that they made no payments into or withdrawals from their assets – i.e.  $X_t^j - X_{t+1}^j = 0$ . When individuals do not know some component of their assets in wave  $t$  (the prior of the waves), but do know it wave  $t + 1$ , we make the mirror assumption, i.e. assume that there were no flows into or out of that asset in the period between the two waves.

## C. Selection and composition of sub-sample

To balance considerations of maintaining as large a sample as is possible and of using as little inaccurate data as is possible, we exclude observations where we feel the data did not allow us to estimate consump-

tion. We provide here details on how our sample was selected (C.1) and then we describe the selected sample and the construction of weights used to correct for our use of a non-representative sample in (C.2).

### C.1 Selection of the sample

We exclude benefit units from the our estimating sub-sample if any of the following conditions hold.

1. If at least one component of income is not known up to a closed interval in *both* waves  $t$  and  $t + 1$ .
2. If a ‘large’ change in physical wealth was observed between the two waves. Changes in the financial asset values that we observe could then be the result of transferring asset holdings between types of assets rather than being indicative of consumption. We define a large change in ‘physical assets’ as occurring when there has *both* been a change in the value of holdings of at least £2,000 *and* a proportionate change in the value of the portfolio of at least 30%
3. If a ‘large’ change in uncategorised wealth was observed between the two waves. A small number of couples keep their finances largely separate. When these individuals are asked for any asset holdings that they hold jointly, only the total level of jointly-held asset is sought - not the level of holdings of particular assets. Estimating the capital gain between waves is thus not possible. A large change in the value of joint holdings is considered to have occurred when there has *both* been a change in the value of holdings of at least £2,000 *and* a proportionate change in the portfolio of at least 30%
4. If an individual bought or sold a house between the waves
5. If mortgage payments are missing
6. If some critical piece of data is missing (usually education or age), the imputation procedure used by the ELSA team to fill in missing income and asset holdings is not possible
7. If the benefit unit composition changed between the two waves
8. If the either partner in a couple does not respond to the survey
9. If a lump-sum payment has been received in the past year and we do not observe the exact amount
10. If either member of a benefit unit suffered the bereavement of their last remaining living parent between waves (and thus may have received an inheritance that we might not observe if was received over a year before the ELSA interview)

Table 9 summarises the selection operated by our procedure from the initial sample. We need two successive observations on assets to compute consumption - therefore the left panel shows the the proportions for whom we have a successful consumption calculation as a proportion of the 6,022 benefit units who make up the balanced panel for waves 1 and 2 (the panel on the right of the table shows these proportions out of all those observed in wave 1. The rules we have established above lead us to exclude 38% of the balanced panel. For an additional 2.3% of the sample, we estimate negative consumption level, indicative of a failure to impute consumption.

Table 9: Success rate of consumption calculation

Computation status	Proportions of balanced panel		Proportions of wave 1 sample	
	Obs.	Percentage	Obs.	Percentage
Have consumption	3,541	58.8%	3,541	44.8%
Calculation failed	2,298	38.2%	2,298	29.1%
Negative consumption	183	3.0%	183	2.3%
Attrited	-	-	1,872	23.7%
Total	6,022	100.0	7,894	100.0%

Table 10 summarises the proportions of benefit units in the balanced panel where a particular reason for the consumption calculation being unsuccessful is relevant.

Table 10: Summary of reasons for consumption calculation failing

Reasons	Percentage	
Missing income component	760	12.6
Large change in physical wealth	721	12.0
Large change in uncategorised wealth	85	1.4
Bought or sold property	321	5.3
Mortgage payments missing	48	0.8
Critical information missing	181	3.0
Benefit Unit composition changed	145	2.4
Non-responding partner	416	6.9
Missing data on lump-sum receipt	141	2.3
Last surviving parent died	352	5.9

We have described above how we deal with cases where some component of assets or income is not known exactly and have had to be imputed. Where assets aren't known in either wave we assume that there has been no flow in or out between waves. We thus do not use the imputed data on assets. With do, on the other hand, use the imputed data on income, but exclude benefit units where at least one component of income is not known up to a closed interval in *both* waves  $t$  and  $t + 1$ . We have assessed if (and confirmed that) our results hold if we apply stricter sample selection rules with respect to the use of imputed data. The third and fourth panels in Tables 5 and 6 for example, show the estimated discount rates with successively stricter sample selection rules than applied in the baseline (in the text of the paper we have referred to these as our 'middle' and 'strict' sample selection rules). These rules are as follows:

- **Middle:** Do not include individuals when the value of some asset holding in wave  $t$  is not known at least up to a closed interval. Do not include individuals where at least one component of income is not reported exactly in *both* waves  $t$  and  $t + 1$ .
- **Strict:** Do not include individuals when the value of some asset holding in wave  $t$  is not reported exactly. Do not include individuals where at least one component of income in wave  $t$  is not reported exactly in *either* wave  $t$  or  $t + 1$ .

The foregoing tables and discussion detail the extent to which we have been able to calculate consumption. To estimate a discount rate we need, of course, at least two successive observations on consumption (and therefore three successive observations in the data). There are 4,886 benefit units in the balanced panel of the first three waves of data. Of these we can include 1,578 (32%) of these in our estimating sub-sample. We exclude benefit units where the consumption calculation failed between either pair of waves (1 and 2 or 2 and 3), where calculated consumption is less than £3,000 in either case or where the benefit unit composition or labour supply activity of either member changed. Table 11 shows the number (and proportion) of cases in which these conditions are relevant.

Table 11: Summary of reasons for consumption calculation failing

<b>Reasons</b>	<b>Percentage</b>	
Consumption calculation failed between either pair of waves	2,532	51.8
Consumption less than £3000 between either pair of waves	780	16.0
Benefit unit composition changed	1,103	22.6
Labour supply changed	690	14.1

## C.2 Characteristics of the selected sample

With a successful calculation of consumption for approximately 60% of the balanced (wave 1 and 2) panel, and the inclusion in the estimating sub-sample of approximately 32% of the balanced (waves 1 to 3) panel, there might be some concern about the representativeness of the sample. In this section we describe the dimensions along which these samples are non-random.

Table 12 contains a multivariate analysis of the probability of our having successfully calculated consumption (the left hand panel) and the probability of being in our discount rate estimating sub-sample (the right hand panel). Each panel contains the results from the two logistic regressions – the first where the sample is all those benefit units observed in wave 1, the second where the sample is the relevant balanced panel.

The dependent variables include dummies for age, marital status, education, quintiles of equivalised income and quintiles of wealth (odd ratios are presented). For couple, age is taken as the age of the male. The only other variable included in the analysis which requires some explanation is education. We categorise individuals as having one of three levels of education: low, middle and high. Those in the ‘low’ group either have no formal academic or vocational qualifications or have a Certificate of Secondary Education, those in the ‘middle’ group have A-levels or O-Levels and those in the ‘high’ group have a higher-level degree. Some individuals have qualifications which don’t fit neatly into one of these categories (because for example, they have foreign qualifications or have some higher level education below degree). We place these individuals in the low, middle and high education group respectively if they left full-time education at or before the age of 16, at 17 or 18, or at or later than 19.

The sub-samples that we use are indeed non-random samples from the entire ELSA sample – with an under-representation of younger benefit units, married benefit units and benefit units with more education. To account for this, we weight the observations in each empirical exercise. The weights are equal the product of the ELSA survey weights and the inverse of the probability of being included in the

estimating subsample. This probability is estimated using the logistic regressions on the entire wave 1 subsample.

Table 12: Logit regression of successfully calculating consumption on sample on characteristics

	Probability of being in consumption sample		Probability of being in discount rate sample	
	(1)	(2)	(3)	(4)
Age 50 to 54 (omm.)				
Age 55 to 59	0.98	1.12	0.94	0.99
Age 60 to 64	1.00	1.11	0.99	1.07
Age 65 to 69	1.15	1.31***	1.77***	2.07***
Age 70 to 74	1.30***	1.69***	1.95***	2.56***
Age 75 to 79	1.21*	1.88***	1.66***	2.56***
Age 80 to 84	0.95	1.70***	1.35**	2.48***
Age 85 or over	0.59***	1.82***	0.76	2.44***
Single (omm.)				
Married	0.81**	0.77**	0.77**	0.83
Widowed	1.30***	1.12	1.10	0.88
Separated	1.29**	1.24	1.33**	1.29
Low Educ. (omm.)				
Mid. Educ.	1.15**	0.98	1.14	0.94
High. Educ.	1.03	0.89	0.96	0.80*
Inc. quint. 1 (omm.)				
Inc. quint. 2	1.01	0.94	0.94	1.00
Inc. quint. 3	1.09	0.97	0.90	0.83*
Inc. quint. 4	1.05	1.00	0.87	0.86
Inc. quint. 5	1.01	0.92	0.66***	0.61***
Wealth quint. 1 (omm.)				
Wealth quint. 2	0.83**	0.45***	0.71***	0.49***
Wealth quint. 3	0.72***	0.33***	0.57***	0.33***
Wealth quint. 4	0.56***	0.22***	0.41***	0.22***
Wealth quint. 5	0.35***	0.13***	0.21***	0.11***
Observations	7,253	5,476	7,253	4,439

\*\*\* p<0.01, \*\* p<0.05, \* p<0.1

NOTES: In specifications (1) and (2) the sample is all those sampled in wave 1, regardless of whether they attrited by wave 2 or not. In specification (2) the sample is the balanced (wave 1 and 2) panel. In specification (4) the sample is the balanced (waves 1, 2 and 3) panel