Savings Age Profiles in the UK

David Demery
Nigel W. Duck[1]

May 2001

Department of Economics
University of Bristol
8 Woodland Road
Bristol BS8 1TN
UK

Discussion Paper No. 01/518

Corresponding Author: David Demery
E-mail: David.Demery@bristol.ac.uk
Telephone: +44 (0)117 928 8437
Fax: +44 (0)117 928 8577

Abstract

In this paper we examine the relationship between age and savings in the UK using data from the Family Expenditure Survey over the period 1969-1998. We address two unresolved problems with cohort methods of estimating age-saving profiles: the potential sample selection problem encountered when analysing household-level savings; and the mismeasurement of pension income. The first is resolved by deriving and analysing *individual* savings rates; the second is resolved by calculating the 'interest component' of pension income. We derive a savings-age profile which is much more consistent with the life-cycle model: in particular, the saving rate rises to a peak in middle age and is generally negative in retirement.

Classification Code: E21
Key Words: Individual and household income, consumption and saving, Life-cycle savings, Family Expenditure Survey

[1] The research reported in this paper is part of the project, *The Macroeconomy and Demographic Change* undertaken as part of the ESRC programme *Understanding the Evolving Macroeconomy.* Financial support from the ESRC is gratefully acknowledged. Material from the Family Expenditure Surveys used in this paper is Crown Copyright; it has been made available by the Office for National Statistics through The Data Archive; and has been used by permission. Neither the ONS nor The Data Achive bear any responsibility for the analysis or interpretation of the data reported here. The authors are grateful for helpful comments from Cliff Attfield, David Winter, Edmund Cannon, Olwen Renowden, Elisabeth Dedman, Ian Tonks and Caroline Joll. We accept full responsibility for all remaining errors.

1 Introduction

In this paper we examine the relationship between age and savings using data from the UK Family Expenditure Survey (FES) over the period 1969-1998. Blundell *et al.* (1994) and Attanasio and Weber (1993) have shown that, after controlling for demographics in preferences and nonseparabilities with labour supply, it is possible to explain observed consumption profiles for *working-age* households. However, the savings behaviour of the elderly has been more problematic: in particular the savings rates of the retired are generally found to be positive whereas the life-cycle model suggests that the retired will dissave; and, at least in the case of the UK, the savings rates of households tend to rise with the age of the household head (Attanasio and Banks (1998)).

A number of explanations have been suggested in the recent literature for the observed age-savings profiles of the elderly. Börsh-Supan and Stahl (1991), for example, argue that the marginal utility of consumption falls significantly after the age of 70, so the elderly accumulate more during retirement than the simple life-cycle model implies.[1] Weil (1994) stresses *inter-vivos* transfers and the bequest motive as possible reasons for high savings amongst the elderly. Attanasio and Hoynes (1995) suggest that because longevity and wealth are positively correlated, older age groups are increasingly populated by the richer (and higher-saving) households. Deaton and Paxson (2000) suggest that those with especially high savings propensities are likely to survive longer as separate households and not be obliged to live within households headed by their children or move to non-household institutions; hence the surviving households with elderly heads will tend to be those with above-average propensities to save. Banks *et al* (1998) suggest that consumption falls around the age of retirement on the 'arrival of new and unfavourable information at retirement' (p.770). Miles (1999) attributes part of the measured high savings rates amongst the elderly to the mismeasurement of pension income, as do Bosworth *et al* (1991).

In our exploration of the relationship between savings and age we touch on some of these themes but primarily we address the two 'major unresolved problems with cohort methods of estimating age-saving profiles' (Deaton and Paxson (2000)). The first is that the life-cycle model is cast in terms of the individual whereas the data are generally available only at the household level and many studies, for the lack of anything better, treat the household as an individual whose age is that of the household head. As suggested above, this practice may pose special problems for the analysis of the behaviour of the elderly: those who, compared with others in their cohort, have a low propensity to save are more likely, when they become elderly, to be obliged to live with their children or in non-household institutions. In the

[1] They offer support for this view using West German data.

latter case they will not be sampled (at least, not in the UK FES); in the former case their behaviour will be attributed to the (younger) head of the household in which they live. As a result, households with elderly heads appearing in the survey are more likely to be those with a high savings propensity. To overcome this potentially serious sample-selection problem we use techniques designed to extract from *household* data information on the savings behaviour of *individuals*.

The second unresolved problem is the mismeasurement of pension income. Most measures misclassify *as income* that component of pensions that comes from running down the underlying asset, with the result that saving is overstated among the elderly. We derive an adjustment factor based on survivorship rates and assumptions about real interest rates and expected inflation that allows us to estimate the pure income component of total pension income.

Our main findings are the following. First, the use of household-level data does introduce some distortion into the UK age-savings profile: in particular it dampens life-cycle variations in the savings rate, exaggerating the rate in early adulthood and post-retirement ages whilst underestimating it in the 45 to 60 age group; however, in both household and individual approaches the savings rates of the elderly are both positive and rising with age. Secondly, once adjustments are made to the measurement of pension income, a savings-age profile emerges which is much more consistent with the life-cycle model: in particular, the saving rate rises to a peak in middle age and is generally negative in retirement. Thirdly, we find that the use of regression methods to extract estimates of individual income and consumption from household data is potentially misleading: in our own case the incomes of the very young and old are underestimated whilst the incomes of young to middle-aged adults are overestimated.

The paper is in four further sections. In the first we briefly present our theoretical framework; in the second, we present evidence on the relationship between the savings rate and age when the household is treated as an individual whose age is that of the household head; in the third we derive estimates of individual savings rates from household data and estimate the relationship between age and individual savings rates; and in the fourth we re-estimate the relationship between savings and age after making the adjustment to total pension income mentioned above. There is finally a brief set of conclusions.

2 The Life-Cycle Framework

We adopt a simple life-cycle framework in which consumption by individual i at age a, c_{ia}, is assumed to be proportional to his or her lifetime resources, W_i. The individual's labour income at age a can also be written as proportional to lifetime resources. The lifetime budget constraint - the present value of consumption must equal the present value of earnings and assets - and the process by which assets evolve allows y_{ia}, total income (the sum of labour and asset income), to be expressed as an age-determined proportion of lifetime resources. The ratio of an individual's savings to income is then simply a function of age: the terms in lifetime resources in the consumption and income relationships cancel out. Formally:

$$c_{ia} = f_i(a) W_i \tag{1}$$

$$y_{ia} = h_i(a) W_i \tag{2}$$

where $f_i(a)$ and $h_i(a)$ are the age-determined proportions of lifetime resources which individual i respectively consumes or receives as income when aged a.

Defining an individual's savings rate as the log of income minus the log of consumption gives:

$$\ln y_{ia} - \ln c_{ia} = \theta_i(a) \tag{3}$$

where $\theta_i(a) \equiv \ln h_i(a) - \ln f_i(a)$.

From equation (1) the *average* log of the consumption of all individuals aged a who were born in year b is:

$$\overline{\ln c_{ab}} = \overline{\ln f(a)} + \overline{\ln W_b} \tag{4}$$

where $\overline{\ln f(a)}$ is the *average* log of the proportion of their life-time resources consumed by individuals aged a; and $\overline{\ln W_b}$ is the *average* log of the lifetime resources of this particular cohort of individuals.

A similar expression for the *average* log of the income of all individuals aged a who were born in year b can be written:

$$\overline{\ln y_{ab}} = \overline{\ln h(a)} + \overline{\ln W_b} \tag{5}$$

From equations (4) and (5) we can derive the following expression for the *average* savings rate of individuals aged a born in year b:

$$\overline{\ln y_{ab}} - \overline{\ln c_{ab}} = \overline{\ln h(a)} - \overline{\ln f(a)} \equiv \theta(a) \tag{6}$$

Note that this savings profile is predicted to be independent of the 'cohort effect' - the log of the life-time resources of individuals born in year b.[2]

[2] In equation (5) we make the explicit assumption that cohort membership does not affect the income/age profile. However there is nothing in theory to prevent the patterns of labour income over the life cycle varying by cohort and this will mean that $\overline{\ln h(a)}$ and $\theta(a)$ will also vary by cohort.

3 Estimation using household data

The life-cycle model sketched above is cast in terms of an individual, yet most surveys provide consumption and income data only at the household level. Furthermore, since most household surveys themselves are not true panels - they sample *different* households over time - many studies construct 'pseudo-panels' by taking sample means of households over time. Consumption, income and savings profiles are then estimated from regressions which use as the dependent variable observations on the mean consumption, income and savings rates of all households with a head of a particular age and date of birth. These regressions typically also control for variables such as household size and demographic composition (such as the number of children) by including these as additional regressors.

In the case of the savings rate a typical regression (omitting any additional control variables) is:

$$\overline{\ln y} - \overline{\ln c} = D^a \left(\alpha_y - \alpha_c\right) + D^b \left(\gamma_y - \gamma_c\right) + \left(u_y - u_c\right) \tag{7}$$

where D^a and D^b are matrices of age and year of birth (cohort) dummies; α_y and γ_y are the age and cohort effects for income; α_c and γ_c are the age and cohort effects for consumption; u_y and u_c are the error terms in the income and consumption regressions respectively; $\overline{\ln y}$ and $\overline{\ln c}$ are stacked vectors of the *average* log of income and consumption respectively of each age and cohort group.

With the savings rate as the dependent variable, only the combined effects of α_y and γ_y or α_c and γ_c, are identified. Furthermore, if the assumptions of the simple life-cycle model were valid, then there should be no cohort effects since γ_y and γ_c should be the same. We present the savings-age profiles derived from estimates of equation (7) below but first we briefly describe our data.

3.1 Data

Our data are derived from the Family Expenditure Survey (FES), an annual cross-section survey of around 7,000 UK households (a sample of around one in 3,000 UK households) which has been extensively used in the analysis of consumption behaviour.[3] The FES is a voluntary survey with a response rate

[3] Notable recent examples are Banks et al. (1998), Paxson (1996) and Deaton and Paxson (1997). The FES has been extensively researched as it is considered sufficiently accurate for the analysis of consumption and savings. Atkinson and Micklewright (1983) suggest that there is little evidence of under-reporting in the income series, with the exception of investment income. Attanasio and Weber (1993) have more recently suggested that for consumption, 'under-reporting is noticeable only on alcohol, a relatively small item. Expenditure on other items is thought to be accurately recorded, thanks to the careful sampling design' (p. 633).

of around 70%, collecting detailed information on household and individual consumption and income, and on household and individual characteristics. Until 1993 the data were released on a calendar year basis; since then the data are organised by financial year. Our source allows us to derive 30 consecutive calendar years of household and individual data from 1969 to 1998.

The measure of household income we use in this study is defined in the FES as 'normal gross income, excluding tax and National Insurance contributions but including income in kind'. It therefore includes private pension contributions but it excludes 'contributions' to the state-funded pension scheme, which we treat in the same way as income tax.

Importantly for what follows, the FES provides details of the income of each *individual* in the household, which permits us to calculate each individual's income on the same basis as household income, i.e. normal gross income excluding tax and National Insurance contributions but including income in kind. Household-level income that is not directly attributable to individuals (for example, rental income and some income in kind)[4] is assigned in the FES to the household head. We attribute such 'common' income to all income-recipients within the household in proportion to their share in other income components, though adopting the FES approach does not materially affect the conclusions we reach in this paper. Income from non-state pensions and annuities are separately identified as an income source which allows us to make the adjustment to pension income which we explain in more detail below.

The definition of *household* consumption we use is total expenditure by household members. Calculating *individual* consumption is more problematic since much consumption is organised at the household level (e.g. housing, utilities, household durable goods, food for consumption within the home etc.) and such spending cannot easily be attributed to individuals, even where the individual who made the expenditure is identified. Because of this, we use regression methods to calculate, from the household data on consumption, the average consumption of individuals of each age and cohort. We explain this approach in more detail below.

Other variables retrieved from the FES are: the number of children (defined as those aged 17 years and under) in the household; the number of adults in the household; and the ages of the household head and other household members.

[4]Such income is a minor component of total household income.

3.2 Results

Figure 1 shows, for the age groups 25 to 80, the relationship between age and the savings rate predicted by two estimates of equation (7).[5] In both we allow for the demographic characteristics of the household by including the number of adults and the number of children in the household as additional regressors. In one we allow for cohort and year effects;[6] in the other we exclude the cohort effects. The profiles shown in Figure 1 are for a household consisting of two adults and one child and, where relevant, for a household head born in 1953.

Neither set of results suggests the relationship between age and the savings rate expected from the simple life-cycle framework described in section 2: in particular the savings rate is positive and increasing in old age; there is at best only mild evidence that the savings rate is especially high in middle age; and an F-test on the significance of cohort dummies strongly rejects the null hypothesis that they are jointly zero.[7] However, there are a number of problems with this approach to estimating the relationship between the savings rate and age. The concept of 'head of household' is a contrived, arbitrary and an increasingly anachronistic one (see Presser (1999)). There is also the possibility of selection bias mentioned earlier: there is evidence that elderly people prefer to live independently but those who have failed to save in their early years may be unable to do so and be obliged to live either in other households (for example with their children) or in non-household institutions; they will therefore either not appear in the survey at all (since individuals in non-household institutions are not sampled in the FES), or not appear as household heads.[8] The elderly who do appear in the sample as household heads will tend to be those with a higher propensity to save.

[5] In all our results we defined the final age group as those aged 80 and over since the number of household heads/individuals in each yearly age group above 80 becomes very small. Our choice of the lowest age group was similarly determined by our desire to have a minimum cell size of 60.

[6] Year effects were captured by including a set of $T-2$ year dummies defined from $t = 3, ..., T$ as $d_t^* = d_t - (t-1)d_2 - (t-2)d_1$where d_t is equal to one if the year is t and is zero otherwise. See Deaton (1997). Estimation of equation (7) using household-level data was by weighted least squares. The weights are inversely proportional to the standard deviation of the cohort means of each variable. Weighted and unweighted regressions produced very similar results.

[7] The F-value of the null hypothesis that the cohort effects were zero was 3.116 which is distributed as $F(79, 1540)$. We discuss possible reasons for the significance of cohort effects in section 4.

[8] In our data, the average age of the head of house of individuals aged 70 years and over is below the age of the individual, though the difference is markedly smaller than that reported by Deaton and Paxson (2000) for the cases of Taiwan and Thailand. This is to be expected given the importance of the extended family in those countries.

4 Estimating individual effects from household data

To overcome some of these problems we use an approach adopted in a similar context by Deaton and Paxson (2000)[9] which allows us to extract estimates of *individual* income (and consumption) from a time series of household surveys by regressing *household* income, y_{ht}, (or consumption c_{ht}) in year t, on n_{aht}, the number of people in the household in each possible age group in year t, where a indexes age groups 1 to A, the maximum age of household members. The coefficient estimated on the number in each age group is interpreted as the *average* contribution to income (or consumption) made by people of that age in year t. More formally, as an identity we can write the income of household h as:

$$y_{ht} \equiv \sum_{a=1}^{A} n_{aht}\overline{\beta}_{aht} \tag{8}$$

where A is the maximum age in the sample; and where $\overline{\beta}_{aht}$ (from equation (2)) equals $h_h(a)\,W_{ht-a}$ - the *average* contribution (for this particular household) to this household's income of people aged a born in year $t-a$.

Write $\overline{\beta}_{aht} = \overline{\beta}_{at} + \varepsilon_{aht}$ where $\overline{\beta}_{at}$ is the average contribution (across all households) to household income of individuals aged a in year t, and ε_{aht} is interpreted as a deviation from this mean peculiar to household h in period t. It follows that we can write:

$$y_{ht} = \sum_{a=1}^{A} n_{aht}\overline{\beta}_{at} + \eta_{ht} \tag{9}$$

where $\eta_{ht} \equiv \sum_{a=1}^{A} n_{aht}\varepsilon_{aht}$. The average contribution of those aged a in period t is identified as the coefficient estimated on n_{aht}. This estimate will be unbiased if there is no correlation between n_{aht} and the error term and this requires that $E\left(n_{jht}\varepsilon_{iht}\right) = 0$ for all j and i. We discuss the likely failure of this assumption below. The estimated value of each $\overline{\beta}_{at}$ is an estimate of $\overline{h(a)}\,\overline{W_{t-a}}$, i.e. the average contribution to income of someone aged a born in period $t-a$. It clearly combines the average age effect, $\overline{h(a)}$, and the average cohort effect $\overline{W_{t-a}}$. To disentangle the two one can take logs of the estimates of $\overline{\beta}_{at}$ and regress them on age and cohort dummies.[10] The same method can be used to derive an estimate of the average contribution each age group and cohort makes to household *consumption*.

[9] The approach we adopt has also been used recently by Chesher (1997, 1998) when analysing individual nutrient intake. For earlier applications to consumer behaviour see Mankiw and Weil (1989) and Weil (1994).

[10] Note that this procedure differs from that adopted in section 3 where we took means across households of log income and consumption. Here we take logs of mean income and consumption.

An attractive feature of this approach is that it should reduce the sample-selection bias mentioned above which might result from the low-saving elderly 'disappearing' from the sample as they cease to be household heads. To the extent that this sub-group have not been 'institutionalised' but live with their children or in other households their behaviour will still contribute to the estimate of the relevant $\overline{\beta}_{at}$ and hence this term will not be so dominated by those with an especially high propensity to save. The obvious weakness of this approach is its failure to model household formation: it treats households as arbitrary collections of individuals. Following Deaton and Paxson (2000), we adopt this approach as 'a first step away from the household version of the model, which we hope will serve as a base from which to make further exploration' (p. 219).

As mentioned above, it is possible to obtain from the FES a direct observation on the *income* of each member of each household and hence the average income of each age group in each year and so, at least in the case of income, the approach we have just described is unnecessary. However the observation of individual incomes does allow us to check the validity of the approach by comparing the regression coefficients $\overline{\beta}_{at}$ with the sample means of individual incomes. We estimate equation (9) across households in the FES for each year between 1969 and 1998 and compare the estimates with the known mean incomes of individuals by age. We present in Figure 2 the averages of both series over the whole period - the actual series is given by the bold continuous line. Three features stand out: first, using the regression method we obtain negative estimates for the mean incomes of individuals less than 15 years old, whereas in fact their incomes were zero (or close to zero).[11] Secondly, the regression method over-states the incomes of individuals in the child-bearing ages, 20-45 years. Thirdly the regression method underestimates the incomes of the elderly - by around 8-10% in the over-60 year age groups. The fact that the regression bias is strongly associated with age is of particular concern given our aim of deriving age profiles for income, consumption and savings.

One reason for this bias may lie in the fact that the numbers of individuals within the household by age is not independent of the residual in equation (9). For example, the presence of children in the household may prevent one parent from working full time and hence lower household income. There may, of course, be other mechanisms which work in the opposite direction: for example households that are richer than average may have more children, leading to a positive correlation with the error term. The downward bias in the estimates of the incomes of elderly individuals is more difficult to explain. If richer households are more likely to place their

[11] Deaton and Paxson, in their analysis of Taiwanese and Thai households, also obtain negative income (and consumption) estimates for children and this forced them to use procedures that avoided the need to take logarithms of negative numbers (see Deaton and Paxson (2000), p. 220).

elderly relatives in institutional care, we may observe more elderly members of poorer households and this may explain the observed bias. Again other mechanisms may also be at work which invalidate the assumption that the regressors in equation (9) are independent of the error.

In theory it would be possible to overcome the endogeneity of the n_{aht} terms by adopting an appropriate estimation technique such as instrumental variables, but the lack of suitable instruments prevented us from adopting this particular approach. The first two problems of the regression method are easily overcome by omitting the numbers of children from equation (9),[12] and this improves the accuracy of the regression approach significantly as is clear from the dashed line in Figure 2. However, the estimates of incomes of those over 60 years old are still seriously underestimated by the regression method.

Since we can directly observe *actual* individual income, the problems associated with the regression method are not directly relevant in the case of income, but they may also apply to our estimates of individual consumption where we are obliged to use the regression approach. We therefore estimate the consumption equivalent of equation (9) by including only those aged 16 and above, since we believe it highly likely that the residual in the consumption equation is also likely to be correlated with the number of children. Our approach implies that we treat consumption as something only adults do; they may choose to consume 'child-related' goods instead of 'adult' goods, but we assume that *total consumption* by individuals is unaffected by the presence of children. However, using the regression approach, we may still underestimate the consumption of the elderly - and therefore overestimate their savings - if the mechanisms that distort the estimated income profiles also apply to consumption.[13] As an initial attempt to address this potential bias we adjusted our estimates of the contribution to household consumption of each age group by the (across-years average) ratio of the known actual to predicted contribution of each age group to household income. In what follows, we report results based on 'corrected' and 'uncorrected' individual consumption coefficients.

4.1 Results

Corresponding to the corrected and uncorrected individual consumption measures we derive two measures of the savings rate: the first uses actual individual income averages for each age and cohort group and the associated *uncorrected* consumption coefficients; the second uses the same income

[12] That is restricting $\overline{\beta}_{at} = 0$ for $a = 1, 2, ..., 15$.

[13] For example, if the *income* estimates of the elderly are biased because poorer households are more likely to have elderly parent members, we would expect the estimated *consumption* of the elderly to be similarly biased. However there may be mechanisms at work in household formation that apply to income but not consumption.

figures but the associated *corrected* consumption coefficients[14]. The age profiles are obtained as the coefficients on age dummies in regressions that also include cohort and year dummies. Figure 3 shows from ages 20 to 80 the relationship between the savings rate and age for those born in 1953.[15]

Again, the relationships between the savings rate and age shown in Figure 3 are not entirely consistent with the life-cycle model. In particular both measures of individual savings show a sharp upward movement in the savings rate at around the age of retirement whereas the life-cycle model suggests a fall - indeed even a negative value for the savings rate in retirement. The estimates based on corrected consumption suggest that for this particular cohort an *individual's* saving rate will increase from around 3.5% when he or she is aged 65 years to around 30% when aged 80. However, the savings-age profile based on the corrected consumption series has, at least over the age range 20-60, the pronounced 'hump-shape' usually predicted by life-cycle models.

Figure 4 highlights the differences between the savings-age profile estimated using household data and that derived from individual observations. The main effects are: the household approach tends to dampen the profile, particularly over the age groups 25-60; the individual profile is closer to the standard life-cycle model, at least up to the age of 65 with the savings rate rising to a peak at around 50 and then declining; in both cases however, the savings rate of the elderly is consistently positive and generally rising with age, contrary to the life-cycle model predictions.

5 Pension income adjustment

The fact that a significant proportion of the working population save for their retirement through non-state pension funds leads to two measurement problems that may well seriously distort the true relationship between savings and age: specifically, savings during working life are *underestimated* whilst savings during retirement are *overestimated.* The savings of those who contribute to pension funds whilst working are likely to be underestimated for two reasons: first, in household surveys such as the FES, *employer* contribu-

[14]The consumption estimates we derive for each age and year (i.e. the coefficient estimates in equation (9), with consumption as the dependent variable) are estimates of the *arithmetic* means, whereas in equation (4) the dependent variable is strictly the average log of consumption. To ensure consistency, we measure savings below as the difference between the log of mean income and the log of mean consumption.

[15]We exclude the earlier age group because their age effects are erratic and would dominate the graph. The age profiles estimated from regressions without cohort and year dummies were similar to those in Figure 3: most notably the rise in the savings rate amongst the elderly was still pronounced. The year and cohort dummies were statistically significant: the F values are 2.114 and 4.136 for the cohort and year dummies respectively, distributed as $F(93, 1764)$ and $F(28, 1764)$. We therefore report results derived from regressions including all dummies.

tions to an employee's pension are not recorded and therefore not included in the calculation of an individual's income (or therefore saving); secondly, the interest on the current value of an individual's pension fund is not imputed as income to the individual and so, again, individual income and savings are underestimated for workers who are saving via pension funds. For both these reasons the data from household surveys will underestimate the savings rate of working individuals. Our data do not permit us to measure or estimate the employers' pension contributions nor impute to individuals the interest income from the current value of their pension fund, and so we can do little but recognise that our data are likely to underestimate the savings rates of the working population.

The second measurement problem specifically affects the elderly and suggests that their savings rate is likely to be overestimated. It arises because of the convention of treating the entire income received by a retired person from an annuity or pension fund as income, whereas that component of the payment to the individual which is attributable to the running down of the underlying asset should be viewed as dissaving.[16] We attempt to correct for this potentially serious overestimate of retirement income by making making an adjustment to reported non-state pension and annuity income. In effect we replace the reported figure with our estimate of the pure interest income from pensions and present a set of results based on that measure of pension income. To do this we first calculate the present value of an annuity stream to an individual j years into retirement with known probabilities of surviving each subsequent period. We then treat the rate of interest times this present value as the interest income from the annuity. More formally, assume an individual who buys an annuity which promises to pay the sum $P\left[1+\gamma\pi\right]^i$ at the end of period $t+i$ for $i \geq 0$, where P is some given nominal sum,[17] π is the (assumed constant) rate of inflation, and γ is the degree of inflation indexation. Let f_{t+j} represent the probability that this individual will be alive at the end of period $t+j$ given that he or she is alive at the beginning of that period. At the beginning of period $t+j$ the nominal present value of this income stream will be:

$$NPV_{t+j} = \frac{P\left[1+\gamma\pi\right]^{j+1} f_{t+j}}{1+r^n} + \frac{P\left[1+\gamma\pi\right]^{j+2} f_{t+j}.f_{t+j+1}}{\left[1+r^n\right]^2} +$$

[16] Bosworth *et al* (1991) attempt to correct for these distortions using US data. They adjust the incomes of those in employment by estimates of the employer pension contributions plus imputed interest income but they appear to exclude all pension payments during retirement (p.203). Applying these adjustments, the savings rates of those in the 45-64 age group in 1982-85 are estimated to be 15.2% (compared with 13% unadjusted) and the savngs rate of those over 64 is -3.9% (compared with 11.5% unadjusted).

[17] Some annuities are 'backloaded' so the nominal payments, though not strictly indexed, are tilted to give a greater share of payments in later years.

$$\frac{P\left[1+\gamma\pi\right]^{j+3} f_{t+j}.f_{t+j+1}.f_{t+j+2}}{\left[1+r^n\right]^3} + ... \tag{10}$$

where r^n is the nominal rate of interest. Our estimate of the interest income from the annuity will be $r^n NPV_{t+j}$. The actual income the individual will be reported as receiving in period $t+j$ will be $P\left[1+\gamma\pi\right]^{j+1}$. Hence the *fraction* of reported income that we will attribute to interest income in period $t+j$ for someone who retired in period t will be:

$$\begin{aligned} F_{t+j} &= \frac{r^n.NPV_{t+j}}{P\left[1+\gamma\pi\right]^{j+1}} = r^n \frac{f_{t+j}}{1+r^n} + \frac{\left[1+\gamma\pi\right] f_{t+j}.f_{t+j+1}}{\left[1+r^n\right]^2} \\ &\quad + \frac{\left[1+\gamma\pi\right]^2 f_{t+j}.f_{t+j+1}.f_{t+j+2}}{\left[1+r^n\right]^3} + ... \end{aligned} \tag{11}$$

In the results we report here we apply this adjustment factor to reported non-state pension and annuity income.[18] This requires assumptions about the nominal interest rate, the expected rate of inflation, the degree of indexation and survivorship rates. We assume that in each year the nominal interest rate is that year's average Treasury bill rate; that there is a constant expected real interest rate of 3%; that the inflation rate expected in year x is the difference between the nominal interest rate and the assumed real interest rate; and that the inflation rate expected in year x is also expected to continue thereafter.[19] Survivorship rates were derived as follows: for the years 1980-1999 they were provided by the Government Actuaries Department;[20] for the earlier years they were computed from mid-year single-age population estimates.[21] For each year we assumed that no-one aged below 50 received pension or annuity income.[22] Because of the difficulty (see Finkelstein and Poterba (2000)) in obtaining data on the extent to which non-state pensions and annuities are indexed, we estimated the savings-age relationship for two cases: the first assumes that no pension income is indexed; the second that 25% of pension income is fully indexed. In fact the

[18] We treat state pensions as transfers among generations and do not apply the adjustment factor. In our data the proportion of income attributed to non-state pensions for households with heads aged 65 and above rose from approximately 12% in 1969 to about 30% in 1998.

[19] In an alternative approach, we subtracted from the nominal interest rate an *estimated* expected inflation series in order to derive a time-varying real interest rate, which we then assumed to be constant over the individual's planning horizon. Our results were largely unaffected by this alternative procedure.

[20] We are grateful to Steve Smallwood of the Government Actuaries Department for providing us with mortality data for the period 1980-1999.

[21] This latter approach ignores the effects of migration on the grounds that these effects are small over the relevant age group.

[22] In fact there are a small number of people in each year who receive annuity income who are below the age of 50.

two sets of assumptions gave very similar results and so we report below only those which assume 25% indexation.[23]

The proportions of reported pension income that we attribute to interest income are reported in Table 1 for selected ages and years. The proportion declines with age as the life expectancy and planning horizons reduce and it also shows significant variation across years, as inflation and nominal interest rates vary. The estimated proportions for 1979 and 1989 are just above 1, which is quite possible for fully indexed pensions in periods of high interest rates and inflation. During periods of low inflation and low interest rates (as in 1998), the proportion is significantly lower, falling to 0.19 for those aged 90.

In Figure 5 we present savings-age profiles based on pension-adjusted individual incomes. These profiles present a picture much more in sympathy with the simple life-cycle model: there is a marked inverted-U shape to the age profile of the savings rate;[24] the savings rate is negative for the very early age group; and most obviously the savings rate becomes negative for most age groups over 65.[25] From the perspective of the simple life-cycle model however, two problems still remain. First, and most importantly for the focus of this paper, there is a marked tendency for savings rates to increase for the very elderly. Several possibilities suggest themselves. One is that ill-health and growing infirmity can make the act of consumption itself increasingly difficult.[26] Another is that health and income are positively correlated, and that at progressively higher ages poor health leads to death. Hence in the higher age groups the FES sample is drawn increasingly from the wealthy[27] who may have a higher propensity to save. The marked rise in the savings rate for the 80-year age group could also be due to the fact that the 80-year group is open ended, covering all individuals aged 80 and over, so the observed effect of age on savings will be somewhat amplified for this group. Secondly, the version of the life-cycle model set out in section

[23] Findings reported by Finkelstein and Poterba (2000) and by Murthi, Orzsag and Orzsag (1999) suggest this 25% indexation may be an over-estimate (at least for the private sector) and hence the proportion we attribute to interest income may be on the high side.

[24] Recall that the omission of employer pension contributions and imputed interest earned by the pension fund mean that we are likely to be underestimating the savings rate of those in employment.

[25] Figure 5 reports results for the cohort born in 1953. The negative savings rates during retirement for this cohort are representative of most other cohorts. However those born before 1918 and those born in 1957 and 1970 have higher estimated savings rates which generally dip to around zero at retirement. Of course they have the same inverted U-shape as those shown for the 1953 cohort.

[26] The finding of Banks et al (1998) that consumption growth is negative over retirement could offer some support for this explanation.

[27] Banks et al (1998) report some evidence in support of this, at least for households with heads over 75. The correlation of wealth and life expectancy in the UK is analysed in Attanasio and Emmerson (2001).

2 suggests the absence of significant cohort effects in the savings equation. We find cohort and year dummies to be highly signficant both on their own and jointly. However, tests that we carried out on sub-periods suggest that, for all definitions of income and consumption, cohort dummies are significant only in the period after around 1980.[28] A detailed explanation of this feature should be the focus of future research but a number of factors may be responsible: there have been, for example, since 1980 substantial changes in the provision of state pensions, such as the indexation of pensions to prices rather than wages. These changes may well have had diffferential cohort effects.

6 Conclusions

In this paper we have derived age-savings profiles for the UK using Family Expenditure Survey data over the period 1969-1998. We address two unresolved problems with cohort methods of estimating age-saving profiles: the first is that of sample selection bias, as high-saving households tend to survive longer than low-savers; and the second is the mismeasurement of pension income, due to the failure to remove that component of pension income associated with the running down of the underlying asset. We find that these problems are serious in the UK case. If we treat the household as if it were an individual whose age is that of the household head and if we treat as income the entire pension, the typical age-savings profile we estimate shows little conformity with the conventional life-cycle model: there is little evidence of a hump-shape to the savings rate; there is no evidence of negative savings rates after retirement; and there is a tendency for the savings rate to rise strongly with age after retirement. The savings age-profiles derived from estimates of *individual* savings rates and *after* adjusting reported pension income are much more in line with the life-cycle model: there is a noticeably more pronounced hump-shape to the savings rate over the working life, with a peak around the age of 50; there is a sharp drop in the savings rate around the age of retirement; the savings rate is generally negative for those who are past retirement age; and there is a much less marked tendency for the savings rate to rise with age after retirement.

Three further findings suggest some caution in interpreting our results. First, cohort effects are consistently and strongly significant in every savings rate equation. We have, as yet, no explanation for this feature. Secondly, we find that the method we (and others) use to extract individual behaviour from household data can be misleading. Thirdly, our estimate of the

[28] For example, estimating the model using corrected consumption and pension-adjusted income over the period 1969 to 1980, we derive an F-test for the null hypothesis of zero coefficients on the cohort dummies of 1.275, which is distributed as $F(75, 630)$ and with a p-value of 0.0673. The same test statistic using data over the period 1981 to 1998 was 2.087, which is distributed as $F(81, 1008)$.

savings rate of individuals during their working life takes no account of the employers' pension contribution nor the interest earned on the individual's pension fund. These features merit further investigation.

Table 1
Pension Adjustment Ratios

Proportion of pension income imputed as interest income (25% indexation)

	1969	1979	1989	1998
Age				
60	0.710	1.083	1.095	0.679
65	0.614	0.967	0.983	0.593
70	0.515	0.836	0.864	0.504
75	0.421	0.700	0.736	0.415
80	0.333	0.570	0.607	0.331
85	0.255	0.454	0.481	0.253
90	0.206	0.384	0.367	0.191

References

Atkinson, A. and J. Mickelwright (1983), 'On the reliability of income data in the Family Expenditure Survey 1970-1977', *Journal of the Royal Statistical Society* 146, 33-61.

Attanasio, O. and J. Banks (1998), 'Trend in household saving: a tale of two countries', Institute for Fiscal Studies mimeo.

Attanasio, O. and C. Emmerson (2001), 'Differential Mortality in the UK', National Bureau of Economic Research, Working Paper No. W8241, (April).

Attanasio, O and H. Hoynes (1995), 'Differential mortality and wealth accumulation', National Bureau of Economic Research Working Paper.

Attanasio, O and G. Weber (1993), 'Consumption growth, the interest rate and aggregation', *The Review of Economic Studies*, 60, 631-649.

Banks, J., R. Blundell and S. Tanner (1998), 'Is there a retirement-savings puzzle?', *American Economic Review*, 88, 769-788.

Blundell, R, M. Browning and C. Meghir (1994), 'Consumer demand and the life cycle allocation of household expenditures', *Review of Economic Studies,* 61(1), 57-80.

Börsh-Supan, A and K. Stahl (1991), 'Life-cycle savings and consumption constraints: theory, empirical evidence and fiscal implications', *Journal of Population Economics*, 4, 233-255.

Bosworth, B., B. Burtless and J. Sabelhaus (1991), 'The decline in saving: some mincroeconomic evidence', *Brookings Papers on Economic Activity*, 183-241.

Carroll, C. and L. Summers (1991), 'Consumption growth parallels income growth: some new evidence', in B. Bernheim and J. Shoven eds, *National Saving and Economic Performance* (University of Chicago Press: Chicago, IL).

Chesher, A. (1997), 'Diet revealed? Semiparametric estimation of nutrienty intake - age relationships', *Journal of the Royal Statistical Society,* Series A, 160, 389-420.

Chesher, A. (1998), 'Individual demands from household aggregates: time and age variation in the composition of diet', *Journal of Applied Econometrics*, 14, 505-524.

Deaton, A.(1997), *The Analysis of Household Surveys: A Microeconometric Approach to Development Policy*, Johns Hopkins University Press.

Deaton, A and C. Paxson (1997), 'The effects of economic and population growth on national saving and inequality', *Demography*, 34, 97-114.

Deaton, A and C. Paxson (2000), 'Growth and saving among individuals and households', *Review of Economics and Statistics*, 82, 212-225.

Finkelstein, A. and J. Poterba (2000), 'Adverse selection in insurance markets: policyholder evidence from the UK annuity market' MIT mimeo.

Mankiw, G. and D. Weil (1989), 'The baby boom, the baby bust, and the housing market' *Regional Science and Urban Economics*, 9, 235-258

Miles, D. (1999), 'Modelling the impact of demographic change upon he economy', *The Economic Journal*, 109, 1-36.

Murthi, M., M. Orzsag and P. Orzsag (1999), 'The charge ratio on individual accounts: lessons of the UK experience', Birkbeck College, University of London, No. 99-2.

Paxson, C. (1996), 'Saving and growth: evidence from micro data', *European Economic Review*, 40, 255-288.

Presser, H. (1998), 'Decapitating the US Census Bureau's "Head of Household": feminist mobilization in the 1970s', *Feminist Economics*, 4, 145-158.

Weil, D. (1994), 'The saving of the elderly in micro and macro data', *Quarterly Journal of Economics*, 109, 55-82.

Figure 1: Household Savings-Age Profile

Regressions include year dummies

(2 adults and 1 child; 1953 Cohort)

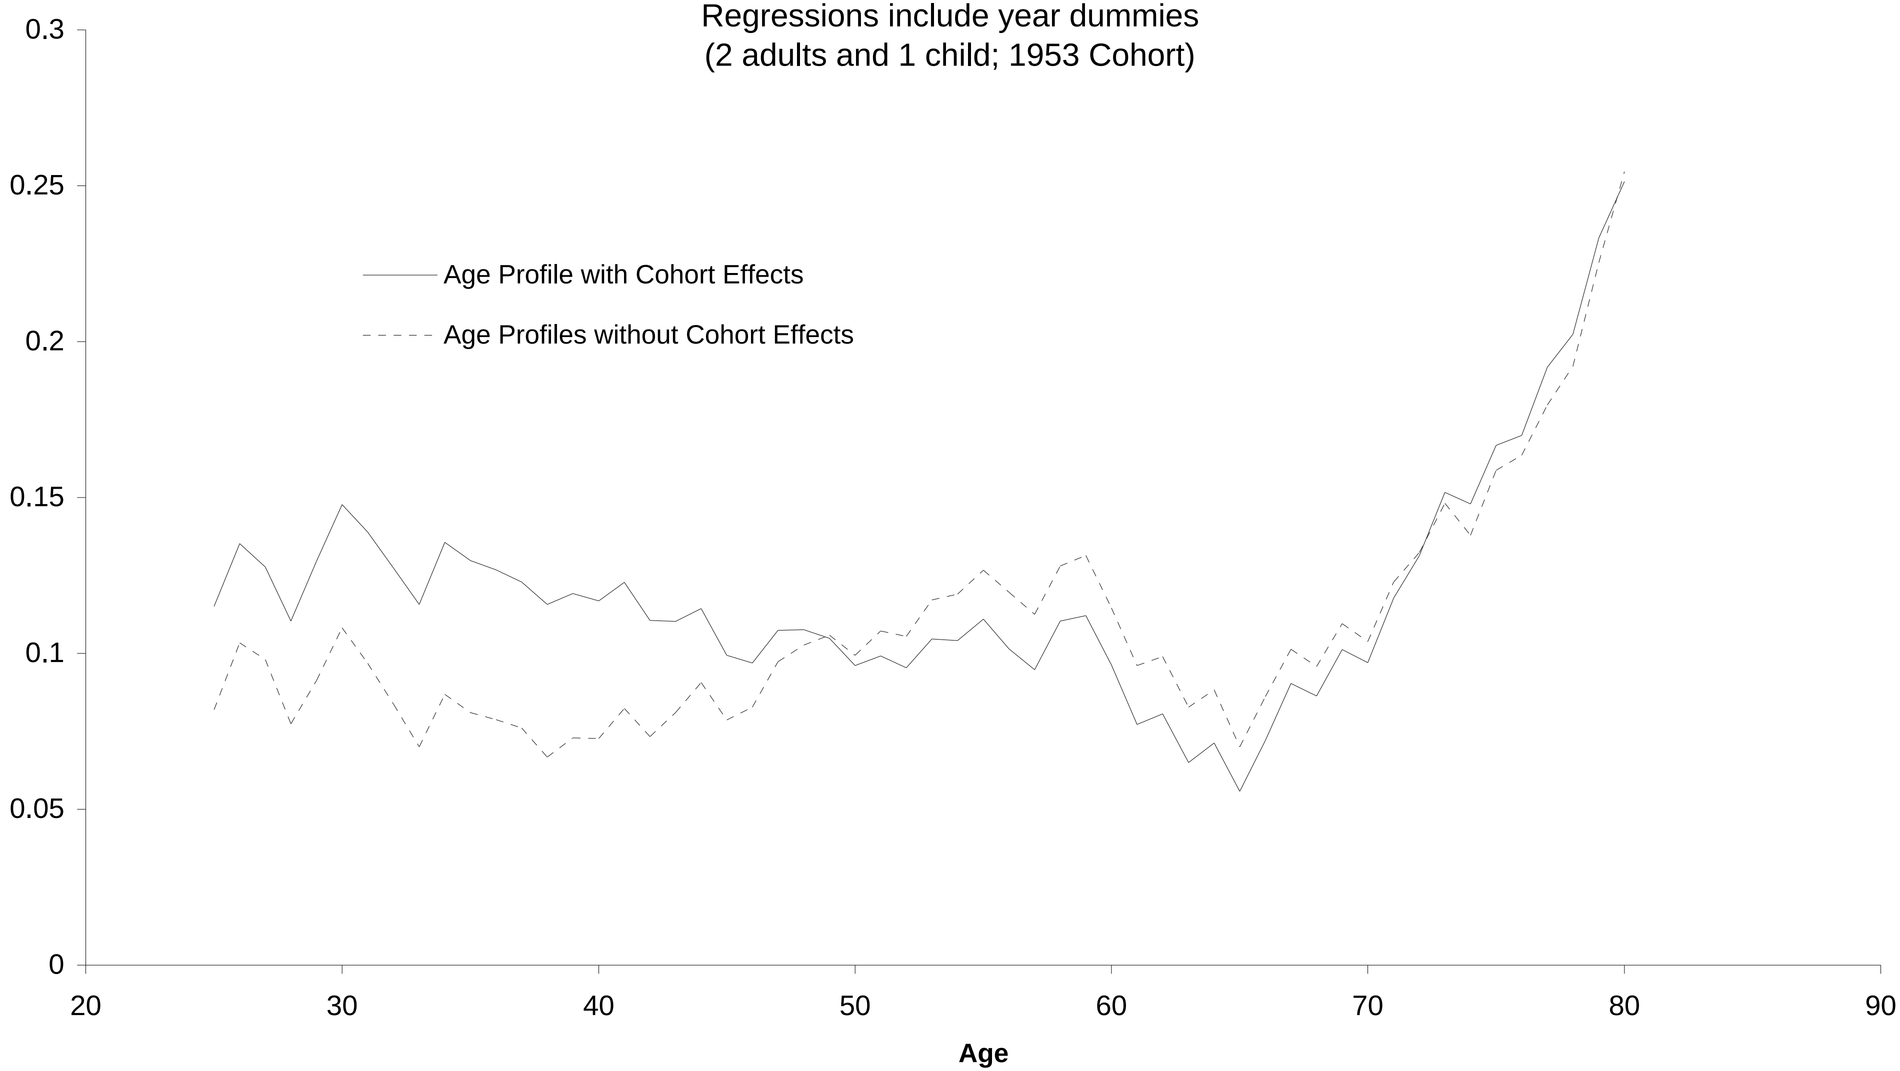

Figure 2: Comparisons of Income-Age profiles
Averages 1969-1998

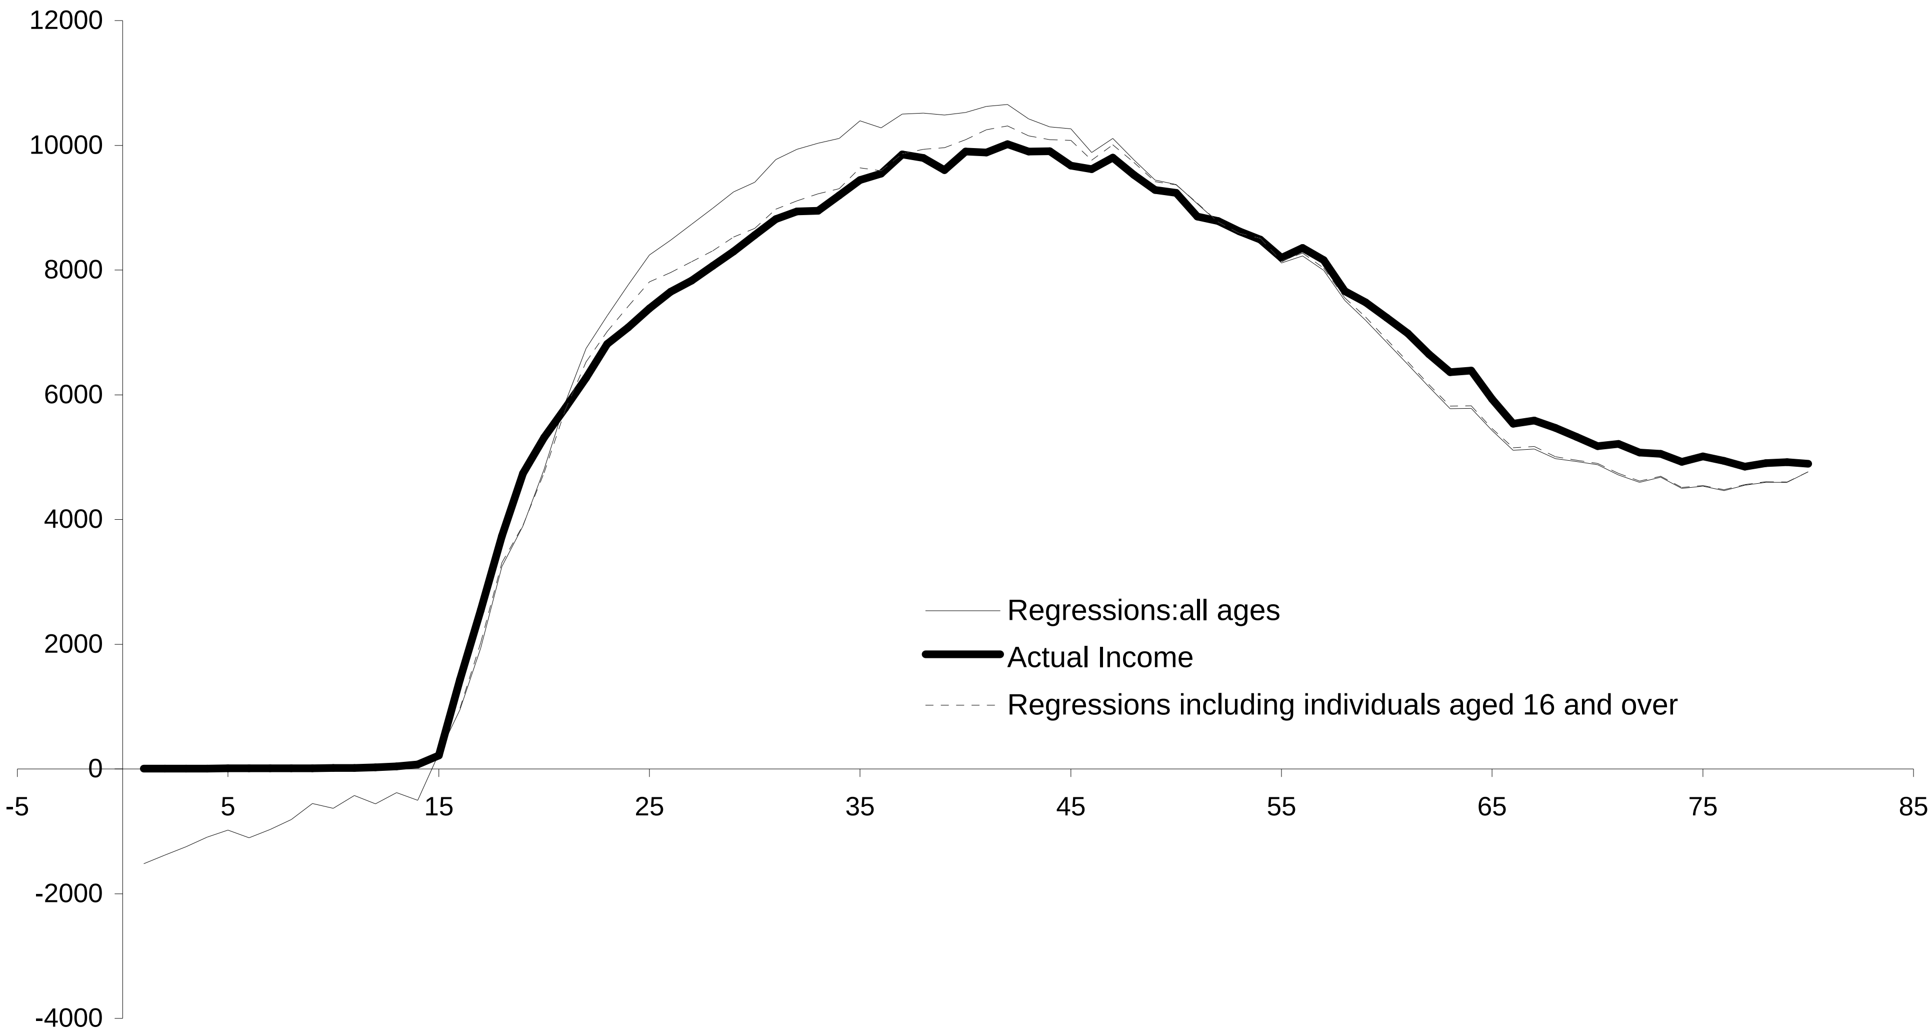

Figure 3: Individual Savings Age Profiles: Income Unadjusted

Regressions include age, cohort and year dummies

(1953 Cohort)

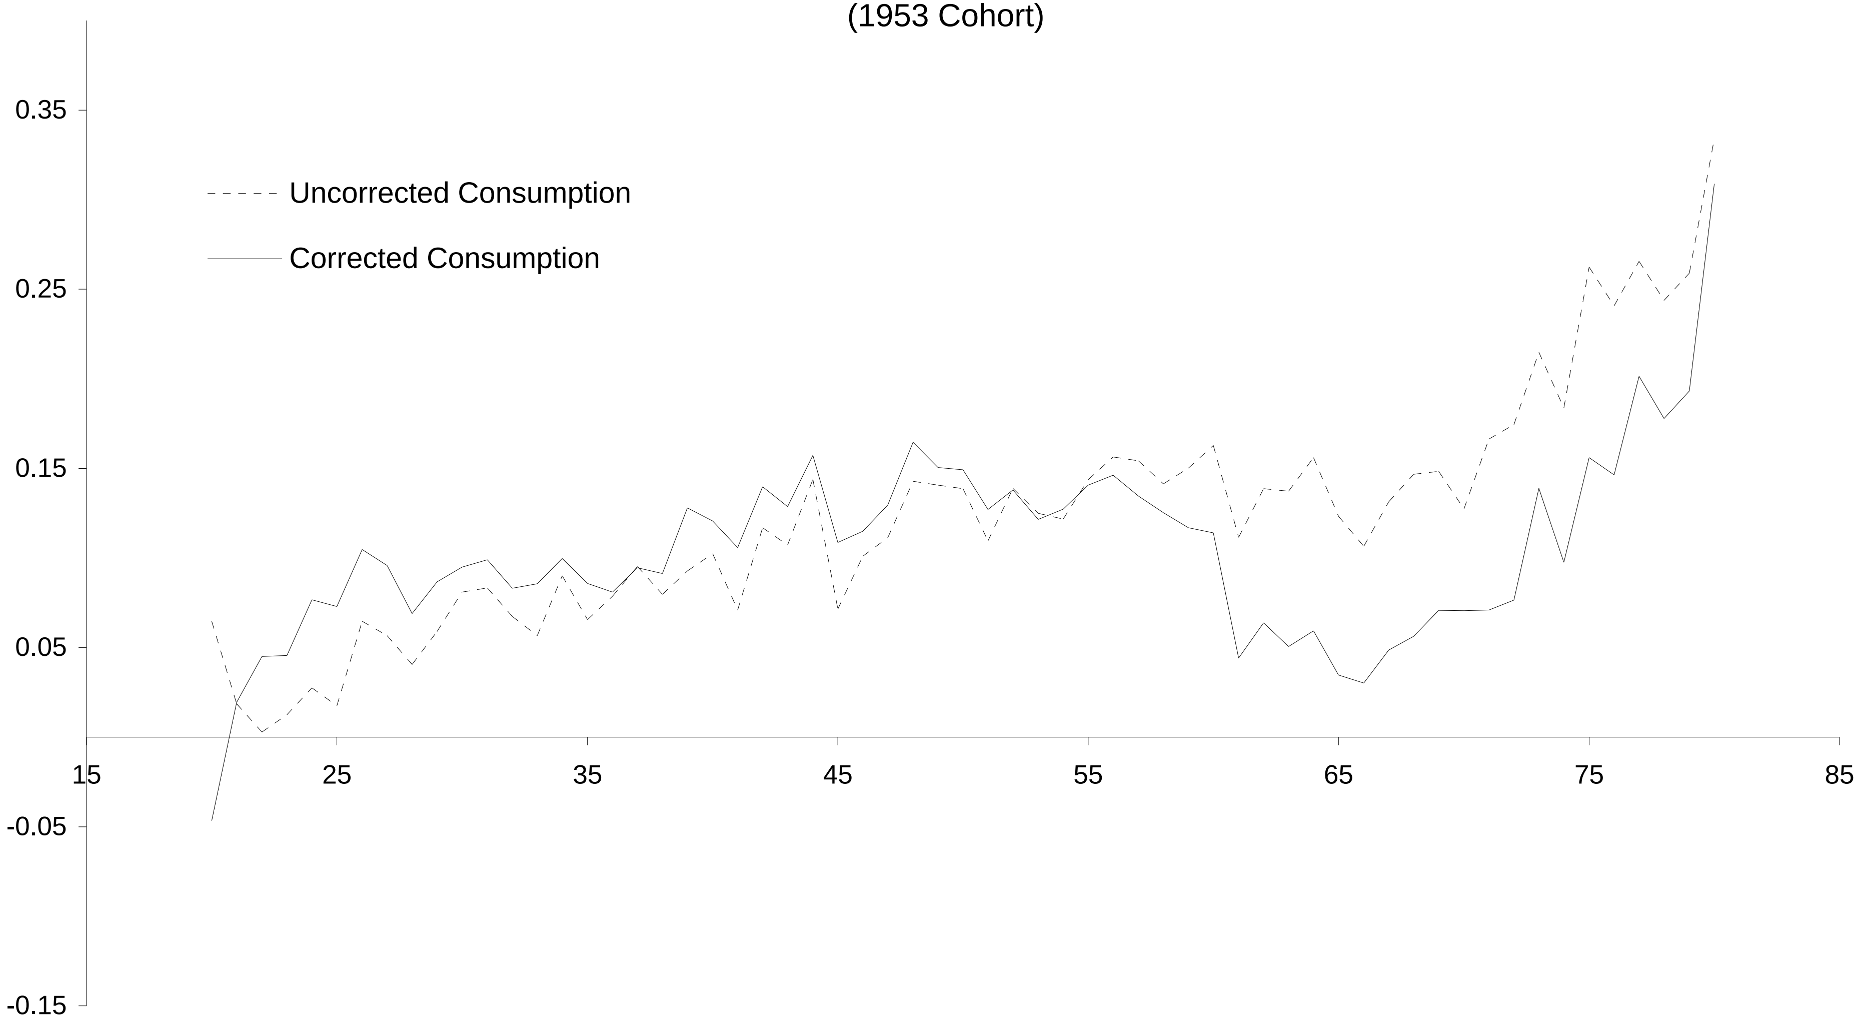

Figure 4: Household and Individual Savings Age Profiles
Regressions include age, cohort and year dummies. Pension income unadjusted
(1953 Cohort)

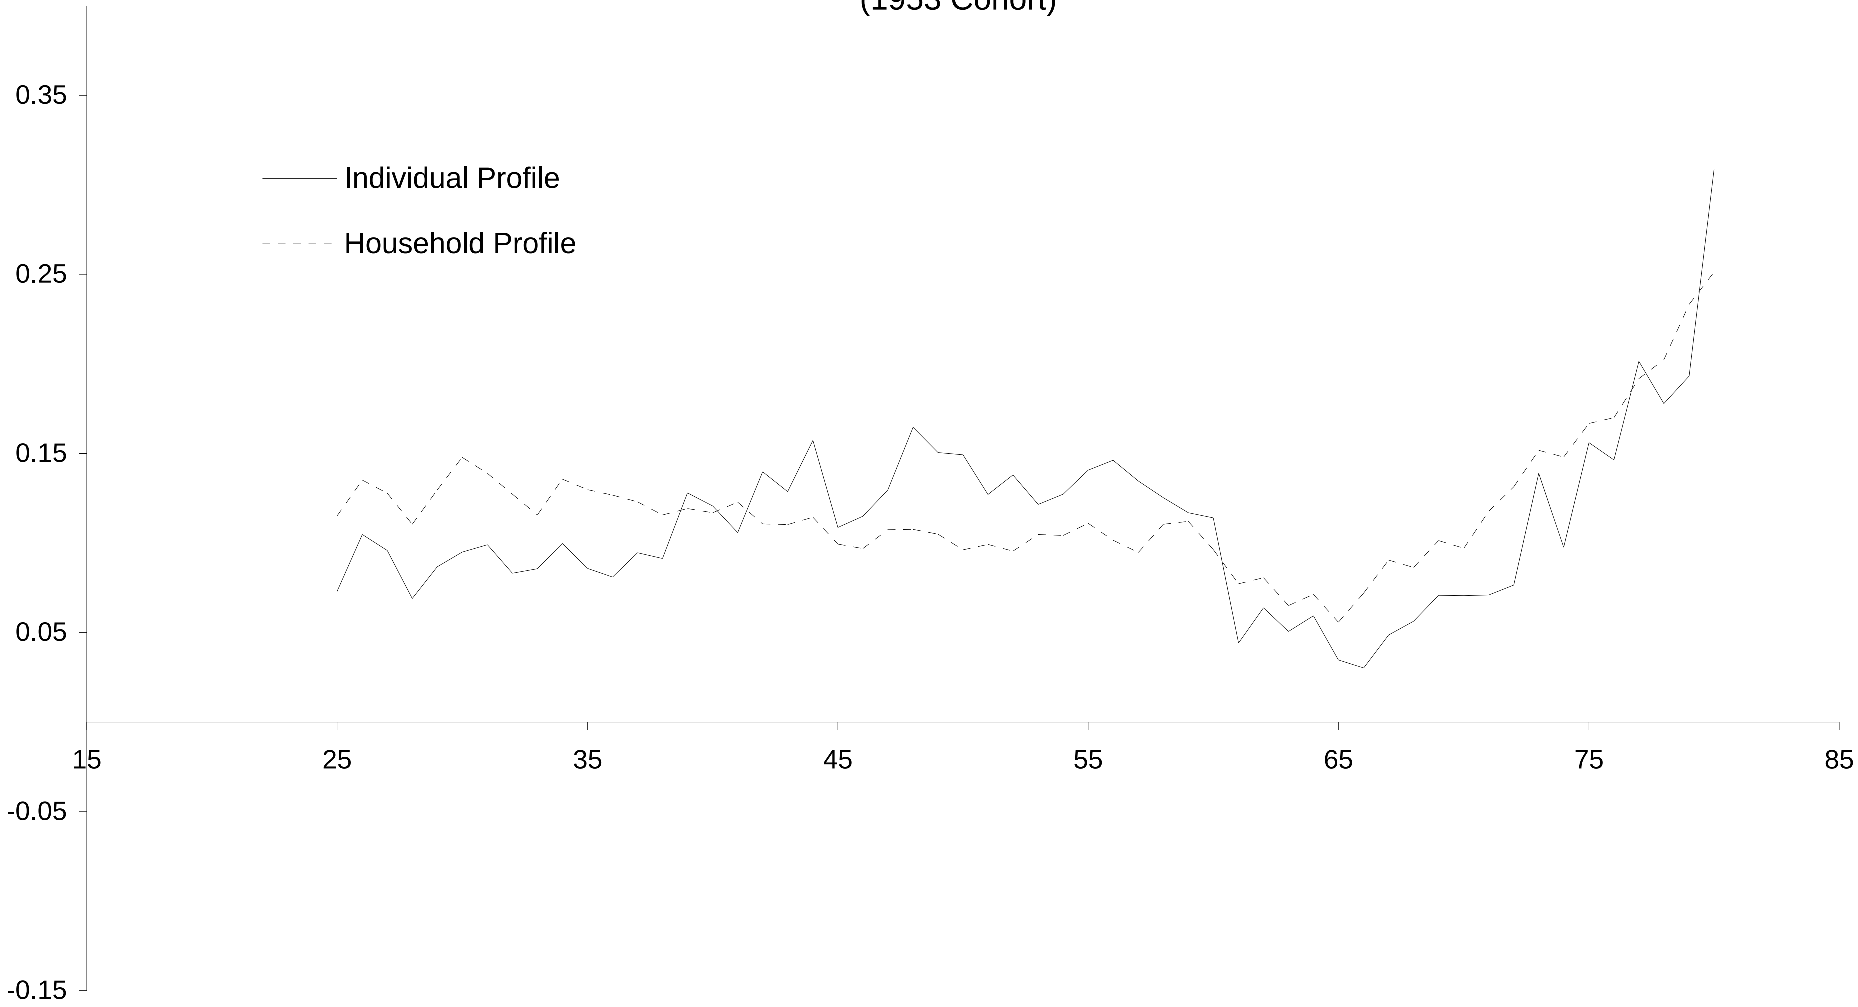

Figure 5: Individual Savings Age Profiles: Income Adjusted
Regressions include age, cohort and year dummies
(1953 Cohort)