Demographic Change and the UK Savings Rate

David Demery and Nigel W. Duck

February 2003

Discussion Paper No 03/550

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

Demographic change and the uk savings rate

David Demery
Nigel W. Duck[1]

February 2003

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

Corresponding Author: Nigel W. Duck
E-mail: N.W.Duck@bristol.ac.uk
Telephone: +44 (0)117 928 8406
Fax: +44 (0)117 928 8577

Abstract

We use *microeconomic* data to explore the effects of a changing age-structure on the UK's *aggregate* personal savings rate. Our findings suggest that changes to the population's age structure age have had detectable, sustained, but, relative to the yearly changes observed in the savings rate over the previous century, modest effects on aggregate personal sector savings. We estimate that the projected changes to the UK's age structure over the next 40 years are likely to raise the UK's savings rate but by no more than 2 percentage points. We find no basis for the view that the aggregate savings rate will decline as a result of the anticipated ageing of the UK population.

Classification Code: G1, G10, G11
Key Words: Saving, ageing population

[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 Maceoeconomy.* 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 (ONS) through The Data Archive; it has been used by permission. Neither the ONS nor The Data Archive bear any responsibilty 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

The UK population is projected to age noticeably over the next 30 to 40 years: the Government Actuary's Department (GAD), for example, estimates that the proportion of the total UK population over the age of 65 will rise from 15.6% in 2000 to 20.23% in 2025 and to 24.9% in 2040.[1] Such a substantial change in the UK's age-structure could, in principle, have significant effects on many features of the UK economy: for example on the structure of total spending; on the housing market; and on labour supply. In this paper we focus on another potentially important effect - the impact on aggregate savings. Specifically, using data from the UK Family Expenditure Survey (FES)[2] over the period 1969 to 1998, we first derive estimates of how consumption and income vary with age and date of birth. These estimates, together with data on the UK's age-structure in the past allow us to assess how the UK savings rate - the ratio of savings to income - has moved in response to changes in age-structure in the past. We compare these movements with the actual behaviour of an independently constructed measure of the savings rate over the period 1855-2001. We also use our estimates, and projected changes in the UK's future age-structure, to assess the likely response of the UK's savings rate to these changes over the next 40 years.

Broadly, we find that a changing age-structure has had effects on the savings rate which, as one would expect, are slow-moving but which are nevertheless detectable, though they have been obscured by the more dramatic changes in the savings rate caused by events which are specific to particular years. Our estimates suggest that the projected changes in the UK's age-structure over the next 30 to 40 years will cause changes in the personal savings rate that are similarly slow-moving and modest.

The paper is in three main sections sections. In a preliminary first section we discuss our measure of savings and some problems associated with it. In the second we explain how we use *microeconomic* data to assess the influence of age-structure on the *aggregate* savings rate. In the third we use those estimates to explore the effects that a changing age-structure has had in the past and is likely to have over the next 40 years. We end with a set of conclusions.

[1] These figures are based on the 2000-based Principal Projection, which is GAD's central projection.

[2] The FES is an annual cross-section survey of around 7,000 households (or around 20,000 individuals). It is a voluntary survey with a response rate 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.

2 Savings and the Family Expenditure Survey

Our investigation of the effects of demographic change on UK savings will focus on *personal sector* savings.[3] There are two reasons for this: first, personal sector savings are likely to be particularly sensitive to demographic changes; and secondly our investigation is based primarily on the FES which covers individuals and households. In this section we address two issues. The first concerns the reliability of FES-based aggregate savings estimates. To do this we compare the aggregate personal sector savings rate published by National Statistics (NS) with an estimate based on the FES. The second issue is the possibility that, since they relate to *persons* and because of the way pension contributions and benefits are treated, the age-savings profiles we estimate from FES data may give a distorted picture of the effects on savings of changes in age-structure.

2.1 The FES and National Accounts Savings Rates

To address the first of these issues we compare the personal sector savings rates implied by FES data with those published in the NS personal sector accounts. For each single-year age group of adults in the FES sample (i.e. individuals over 15) and for each calendar year, we computed means of disposable income and consumption expenditures.[4] *Aggregate* income and consumption estimates were computed by weighting these means by the known population shares in each cell. The FES aggregate personal sector savings rate was computed from these aggregations.

The NS estimates of aggregate disposable income and consumption are derived from the income and capital accounts for household and non-profit institutions.[5] To allow a direct comparison with the FES estimates, two adjustments were required. First, the household savings rates reported by NS (code: NRJS) includes in income an adjustment for changes in net equity of households in pension funds, an item not included in the FES income definition. For our NS income figure we adopted gross disposable income (code: RPHQ) before the addition of the pension fund item. The second adjustment involved the treatment of imputed rents from owner occupation, an item that was discontinued in the FES income and consumption definitions from 1992.[6] The NS income and consumption estimates (codes: RPHQ and RPQM respectively) were reduced by the owner-occupier imputation (code:

[3] National savings includes personal and corporate sector savings (the private sector), and saving by the government and the rest of the world.

[4] Income data for *individuals* are explicitly reported in the FES but to derive individual consumption expenditure we adopted a regression approach we explain below. Our reasons for aggregating over the *adult* population (over 15) is also explained below.

[5] The FES covers households only and does not include non-profit institutions.

[6] In our FES analysis, imputed rent from owner-occupation was removed from the pre-1992 income and consumption data to ensure consistency over the whole data period.

ADFU).

In Figure 1 we graph the resulting NS and FES personal sector savings rates for the years 1971-1998. The starting point is 1971 rather than 1969 - the year of the first FES in our data set - because single-year estimates of age-structure were available only from 1971. Overall, the NS and FES savings rates show an acceptable degree of correspondence except for the period 1987 to 1990 when the FES-based estimates fail to pick up the decline in the savings rate in the NS accounts. The NS mean savings rate over the period 1971-1998 is 5.485% and the mean FES rate is 5.612%. For our purposes - to explain and forecast *longer run* trends in the aggregate savings rate due to a changing age-structure (rather than its year to year changes) - the fact that our micro-data yields similar savings rates *on average* to those in the national accounts, and the close correspondence of the two series over much of the period, suggests that our projections will provide a good guide on average to the NS's estimate of the UK's aggregate personal savings rate, or at least to the influence that a changing age-structure exerts on it.

2.2 Age-Savings Profiles

The age-savings profiles which are central to the results in this paper are estimated from data on the incomes and expenditures of households or individuals. Such age-savings profiles may give a distorted picture of the effects of changes in age-structure on *national* savings because of the treatment of pension contributions and benefits. First, the savings of those who contribute to pension funds whilst working are likely to be underestimated. In household surveys such as the FES, employer contributions to an employee's pension fund are not recorded and so are not included in the calculation of an individual's income (or, therefore, her savings).[7] Moreover, the interest earned 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 working-age individuals who are contributing to private pensions. Consequently the age-savings profiles derived using personal sector data will underestimate the full savings 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 funds, and so we can do little but recognize that our data are likely to underestimate the savings rate of the working population.

The second distortion to age-savings profiles based on personal sector data arises from the treatment of pension benefits. Income in the FES (and in the personal sector national accounts) treats the entire pension benefit or annuity payment as income. As Deaton and Paxson (2000) point out,

[7] In the National Statistics personal sector accounts, personal savings are defined as income net of taxes and social contributions - the latter including funded and unfunded pension contributions by employees and employers - minus personal consumption.

this 'misclassifies the component of annuity disbursement that is not income but comes from running down the underlying asset' (p212). Any associated decline in the value of the asset backing that annuity will affect the institution paying the annuity - whose accounts do not of course appear in the FES. Thus, for example, two otherwise identical individuals - one of whom has built up his own stock of financial assets and consumes the constant amount in retirement which will reduce his financial assets to zero at his (assumed known) date of death, and another who, in effect, gets a pension fund to pay him an annuity[8] so that he can consume exactly the same constant amount - will be measured as having different savings: the former will be measured as dissaving, the latter as not. Age-savings profiles derived using personal sector income and expenditure data are likely to reveal inflated savings rates for the elderly, especially if private pensions are widespread. Later in the paper we attempt to correct for this distortion by identifying the component of pension income that can be attributed to the running down of the underlying asset.

3 Estimating the relationship between the aggregate savings rate and age-structure from microeconomic data

Our framework for examining the relationship between the UK's age-structure and its savings is the following definition of the economy's personal sector savings rate in period t,

$$\left(\frac{S}{Y}\right)_t = \frac{\sum_{a=1}^{A} \left(y_{ab} - c_{ab}\right) n_{at}}{\sum_{a=1}^{A} y_{ab} n_{at}} \tag{1}$$

where y_{ab} is the average income of individuals aged a in period t, i.e. of people who were born in year $b = t - a$; c_{ab} is the average consumption of those aged a who were born in year b; n_{at} is the proportion of the population aged a in period t; A is the oldest age group; S is aggregate personal savings; Y is aggregate personal income.

Broadly, our approach is to use a simple lifecycle framework which views y_{ab} and c_{ab} as functions solely of age and date of birth - an age effect and a cohort effect. We first derive estimates of these cohort and age effects - and hence estimates of y_{ab} and c_{ab} - for the cohorts and age-groups we observe. We then use these results to make estimates of the cohort and age effects - and hence of y_{ab} and c_{ab} - for the earlier and later cohorts that we do not observe but which we need if we are to 'forecast' a savings rate for the years after 1998 and before 1969. Finally, we combine all our estimates of y_{ab} and c_{ab}, and past and projected population share data to assess what influence a changing age-structure has had on the savings rate in the past and what influence it is likely to exert in the future.

[8] We treat annuities and pensions as identical financial assets. For simplicity we refer to income derived from both as pensions.

Before we can carry out this procedure we address three technical issues. The first is how to obtain estimates of *individual* income and consumption from a survey whose unit of observation is the *household*. The second is the appropriate treatment of pension income. The third is how to estimate age - and (especially) cohort - effects for *unobserved* cohorts.

3.1 Estimating individual effects from household data

Consumer theory focuses on the individual whereas, in surveys such as the FES, the unit of observation is the household. Many empirical studies deal with this problem by treating the household as an individual whose age is that of the designated 'head of household'. This is unsatisfactory on a number of grounds: first, the concept of head of household is contrived and anachronistic; secondly, if individuals live in multi-generational households the true relationship between age, consumption and income will be obscured; and, thirdly, to the extent that the poorer elderly might be obliged to live with their children whereas the better-off elderly continue to exist as independent households, the income and consumption patterns of the elderly might be systematically biased.

In fact, from the FES we can *directly* extract information on the *income* of individual members of the household but we cannot directly extract information about the *consumption* of each individual. To extract that information indirectly - and avoid the problems caused by treating the household as an individual - we employ the technique adopted in Deaton and Paxson (2000).[9] As a check on the accuracy of this technique we also apply it to income despite being able to extract individuals' income directly, though in all the results we report below, we use the direct measure of individual income.

We extract estimates of *individual* income and consumption by regressing *household* income in year t, y_{ht}, and consumption, c_{ht}, on n_{aht}, the number of people in the household in age group in year t, where a indexes (single) age-groups 1 to A.[10] We interpret the coefficient estimated on n_{aht} 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:

$$y_{ht} \equiv \sum_{a=1}^{A} n_{aht}\overline{\beta}_{aht} \tag{2}$$

where $\overline{\beta}_{aht}$ is the *average* contribution (for this particular household) to this household's income of people aged a born in year $t-a$. Write $\overline{\beta}_{aht} =$

[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] For a more detailed discusssion of this procedure see Demery and Duck (2001).

$\overline{\beta}_{at} + \varepsilon_{aht}$, 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:

$$y_{ht} = \sum_{a=1}^{A} n_{aht}\overline{\beta}_{at} + \eta_{ht} \tag{3}$$

where $\eta_{ht} \equiv \sum_{a=1}^{A} n_{aht}\varepsilon_{aht}$. Estimating equation (3) across households in the FES, the average contribution of those aged a in period t is identified as the coefficient estimated on n_{aht}. This estimate would be unbiased if there were no correlation between n_{aht} and the error term, which requires that $E\left(n_{jht}\varepsilon_{iht}\right) = 0$ for all j and i. We discuss the possible failure of this assumption below.

In a simple life-cycle model, an individual's consumption at age a will be proportional to his or her lifetime resources. 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[11] - and the process by which assets evolve allow the individual's total income to be expressed as an age-determined proportion of lifetime resources. Formally,

$$c_{iat} = f_i(a)W_i \tag{4}$$

$$y_{iat} = h_i(a)W_i \tag{5}$$

where c_{iat} is the consumption of individual i aged a in period t; y_{iat} is the income of individual i aged a in period t; W_i is the lifetime resources of the individual; and $f_i(a)$ and $h_i(a)$ are respectively the age-determined proportions of lifetime resources which the individual consumes and receives as income when aged a.

Within this framework the estimated value of each $\overline{\beta}_{at}$ can be interpreted as an estimate of $h\left(a\right)W_{t-a}$, the *average* contribution to income of someone aged a born in period $t-a$, which we write as y_{ab}. A similar regression involving consumption would provide an estimate of $f(a)W_{t-a}$, the *average* contribution to consumption of someone aged a born in period $t-a$, c_{ab}. Clearly both combine the average age effects, $h\left(a\right)$ or $f(a)$, and the average 'cohort effect' W_{t-a} . To disentangle the two in the case of income one can take logs of the estimates of $\overline{\beta}_{at}$ and regress them on age and cohort (year of birth) dummies.[12] A similar technique can be used to derive an estimate

[11] Ignoring, for the moment, bequests.

[12] There is a possibility that some of the estimates of $\overline{\beta}_{at}$ may be negative. For this reason Deaton and Paxson (2000) employ a different technique to extract the age and cohort effects. This problem does not arise in our case once we restrict age-groups - as we do for reasons explained below - to those aged 16 and above.

of the average contribution each age group and cohort makes to household consumption. Formally, we carry out the regressions[13]

$$ln(c_{ab}) = D^a\alpha_c + D^b\gamma_c + u_c \tag{6}$$

$$ln(y_{ab}) = D^a\alpha_y + D^b\gamma_y + u_y \tag{7}$$

where c_{ab} and y_{ab} are our estimates of the combined average age and cohort effects on consumption and income respectively; D^a and D^b are matrices of age and cohort (year of birth) dummies respectively; α_c, and α_y are each a vector of coefficients showing the effects of age on consumption and income respectively; γ_c, and γ_y are each a vector of coefficients showing the effects of date of birth or cohort membership on consumption and income respectively; and u_c and u_y are error terms.

The cohort dummies pick up both the effects of changes in wealth between cohorts and any other factors that cause one cohort to behave differently from another. Under simplifying assumptions - that bequests are a fixed proportion of lifetime resources and that lifetime consumption and bequests together exhaust lifetime resources - the vectors of coefficients, γ_c and γ_y, should be equal.

An attractive feature of this approach is that it should reduce the sample-selection bias which might result from the poorer elderly 'disappearing' from the sample as they cease to be household heads. To the extent that they have not been institutionalised but live with their children or in other households, this sub-group's behaviour will still contribute to the estimate of the relevant $\overline{\beta}_{at}$ and hence this term will not be so dominated by those who, by the standards of their peers, are high earners and consumers.

In fact, 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. So, at least in the case of income, the approach we have just described is unnecessary. However, the fact that we can *directly* observe an individual's income does allow us to check the accuracy of the regression approach by comparing the regression coefficients $\overline{\beta}_{at}$ with the sample means of individual incomes. Demery and Duck (2001) estimated equation (3) across households in the FES for each year between 1969 and 1998 and compared the estimates with the *known* mean incomes of individuals by age. They found that: (a) the regression method produced negative estimates for the mean incomes of individuals less than 15 years old, whereas in fact their incomes were zero (or very close

[13] The regressions also include year dummies to pick up cyclical influences on income and consumption, influences that are common to all ages and cohorts. 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)). The year dummy effects sum to zero by construction.

to zero);[14] (b) the regression method over-stated the incomes of individuals in the child-bearing ages, 20-45 years; (c) the regression method under-estimated 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 hence savings.

This bias may be because the number of individuals within the household by age is not independent of the residual in equation (3). For example, the presence of children in the household may prevent one parent from working and therefore 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 reasons for the downward bias in the estimates of the incomes of elderly individuals is more difficult to explain. One possibility is that if richer households are more likely to place their elderly relatives in institutional care we may observe more elderly members of poorer households and this may explain the observed bias.

Demery and Duck (2001) found that the accuracy of the regression approach improved markedly if the first 15 age-groups were omitted from equation (3). In particular, the first two problems of the regression approach mentioned above were effectively overcome. However, the incomes of the over-60s were still seriously underestimated.

Since we directly observe actual individual income, the problems associated with the regression method are not directly relevant in the case of income, but they will apply to our estimates of individual consumption where we are obliged to use the regression approach. Since we believe it highly likely that the residual in the consumption equation is also likely to be correlated with the number of children, we follow Demery and Duck (2001) and estimate the consumption equivalent of equation (3) by including only those aged 16 and above. The possibility that, by using the regression approach, we may underestimate the consumption of the elderly - and therefore over-estimate their savings - remains. To address it 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. The results we report in this paper are based on these corrected consumption-age profiles.

[14] 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 number (see Deaton and Paxson (2000), p. 220).

3.2 Pension income adjustment

In the previous section we discussed a number of possible distortions to the age-savings profile caused by the treatment of pension income. Some of these we can do little about, but the distortion caused by the misclassification of annuity income highlighted by Deaton and Paxson (2000) can be addressed.[15]

To correct for this potentially serious overestimate of retirement income we make an adjustment to reported non-state pension income based on the idea that interest earned on the underlying asset should be treated as income, whereas the decline in the value of the underlying asset should not. We derive an adjustment factor which attempts to identify empirical counterparts to those conceptually different components of pension benefits. In effect we replace the reported figure with our estimate of the pure interest income from non-state pensions.

To do this we first calculate the present value of an annuity stream to an individual j years into retirement with given 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 $P\left[1+\phi\pi\right]^{i}$ at the end of period $t+i$ for $i \geq 0$, where P is some given nominal sum, π is the (assumed constant) rate of inflation, and ϕ 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:

$$\begin{aligned} NPV_{t+j} &= \frac{P\left[1+\phi\pi\right]^{j+1} f_{t+j}}{1+r^{n}} + \frac{P\left[1+\phi\pi\right]^{j+2} f_{t+j}.f_{t+j+1}}{\left[1+r^{n}\right]^{2}} + \\ & \quad \frac{P\left[1+\phi\pi\right]^{j+3} f_{t+j}.f_{t+j+1}.f_{t+j+2}}{\left[1+r^{n}\right]^{3}} + \ldots\ldots \end{aligned} \tag{8}$$

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 $P\left[1+\phi\pi\right]^{j+1}$. Hence the *fraction* of reported income that we will attribute to interest income in period $t+j$

[15] Jappelli and Modligliani (1998) attempt to overcome this problem by omitting pension income altogether, treating all pension benefits as a 'decumulation of the stock of pension wealth' (p.11). Using Italian cohort data, they find, on making this adjustment a more hump-shaped pattern to the savings-age profile and negative savings for the elderly. Bosworth *et al.* (1991) also treat all pension benefits as dissaving. Miles (1999) also suggests that part of the measured high savings rates amongst the elderly is due to the treatment of pension benefits.

for someone who retired in period t will be:

$$F_{t+j} = \frac{r^n.NPV_{t+j}}{P\left[1+\phi\pi\right]^{j+1}} = r^n \left[\begin{array}{c} \frac{f_{t+j}}{1+r^n} + \frac{[1+\phi\pi]f_{t+j}.f_{t+j+1}}{[1+r^n]^2} + \\ \frac{[1+\phi\pi]^2 f_{t+j}.f_{t+j+1}.f_{t+j+2}}{[1+r^n]^3} + ... \end{array} \right] \quad (9)$$

For some of the results we report in the next section, we apply this adjustment factor to reported non-state pension income.[16] 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. Survivorship rates were derived as follows: for the years 1980-1999 they were provided by the Government Actuary's Department;[17] for the earlier years they were computed from mid-year single-age population estimates.[18] For each year we assumed that no-one aged below 50 received pension or annuity income.[19] Because of the difficulty (see Finkelstein and Poterba (2000)) in obtaining data on the extent to which non-state pensions are indexed we considered two cases: the first assumes that no pension income is indexed; the second that 25% of pension income is fully indexed. Demery and Duck (2001) report that the two sets of assumptions gave very similar results and in what follows we report results based on those which assume 25% indexation.[20]

3.3 Estimating age and cohort effects for unobserved cohorts

Since one of our main aims is to assess the likely effects of projected population changes on the *future* aggregate UK savings rate, we shall need some estimate of the age and cohort effects for cohorts who are as yet unborn or, at least, who do not yet figure in our data set. To assess the influence of age structure on *past* savings rates we also require estimates of earlier cohort effects for cohorts which do not appear in our data set.

[16] We treat the state pension as a transfer from one generation to the next and do not apply the adjustment factor. 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.

[17] We are grateful to Steve Smallwood of the Government Actuary's Department for providing us with mortality date for the period 1980-1999.

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

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

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

For the age effects, the solution we shall adopt is straightforward: we shall assume that the age effects for future cohorts are the same as those estimated from the observed cohorts. Our estimates of these age effects, i.e. our estimates of α_y and α_c, are presented in Figures 2A and 2B. Those in Figure 2A are based on the standard measure of income (i.e. no adjustment for the treatment of pension income) and consumption corrected in the way described above. We refer to these as Model A results. Those presented in Figure 2B are based on adjusted income and corrected consumption. We refer to these as Model B results. Recall that Model B incorporates an adjustment aimed at securing a more accurate estimate of the relationship between age-structure and personal income (and hence saving) *for the economy as a whole*; Model A makes no such adjustment and will therefore provide a less accurate guide to this relationship the more the adjustment is required, i.e. the more wide-spread are non-state pension funds and the greater the proportion of older individuals in the population.[21]

Both sets of estimates of α_y and α_c show signs of a hump-shape. Income in particular rises to a peak in the mid 40-50 age-groups though both age-income profiles also show a tendency for income to rise with age beyond the age of 70 - a result that may be due to positive correlation between longevity and wealth (see Attanasio and Hoynes (1995)). This feature is, of course, less pronounced in the case of Model B where we have adjusted income downward to allow for the misallocation of part of any annuity disbursement. Figure 3, which plots the relationship between age and the savings ratio (defined here as $\ln y_{ab} - \ln c_{ab}$) for the cohort born in 1953, shows that the adjustment to income clearly has a significant qualitative impact on the predicted relationship between an economy's age-structure and its savings. Without any adjustment to income, savings are predicted to be *higher* for an economy with a high proportion of elderly people; with the adjustment savings are predicted to be *lower*.

The problem of estimating unobserved cohort effects is more difficult and is highlighted in Figure 4 which shows our estimates of the income and consumption cohort effects, γ_y and γ_c respectively, for Model B.[22] Because our data are from the FES over the period 1969-1998, and since we use age-groups from 16 years to 80 (and above),[23] we can *directly* estimate cohort effects only for those born between the years 1889 and 1982. Furthermore,

[21] The savings rate implied in Model A corresponds to the personal sector savings rate published in the national accounts (ignoring the equity adjustment discussed above). The implied aggregate savings rates from Model B would be closer to a savings rate derived from a *consolidation* of the accounts of the personal and financial sectors. However because Model B underestimates the savings of working individuals, it will underestimate the true 'consolidated' savings rate.

[22] The cohort estimates for Model A are virtually identical to those estimated for Model B.

[23] We assume that individuals over 80 years old have the same age and cohort characteristics of those aged 80.

our estimates of the very earliest and most recent cohort effects are based on fewer and fewer observations. For example, the cohort born in 1982 appear only once in our data - as 16 year-olds in the 1998 FES; the cohort born in 1981 appear twice - as 16 year-olds in the 1997 FES and as 17 year-olds in the 1998 FES. Since they are based on such a small number of observations, the very earliest and most recent cohort effects are likely to be erratic. Figure 4 suggests that this is indeed the case. The estimated 'extreme' cohort effects behave quite differently from those in the 'middle'. Their erratic behaviour suggests that they are not only dubious in themselves but are likely to provide a poor foundation for projecting into the future (or the past).[24]

One method suggested by Deaton and Paxson (2000) for estimating the future savings rate which bypasses this problem is to assume that the income and consumption cohort effects are identical and primarily the result of lifetime wealth differences. On the assumption that each cohort's life-time resources equal those of its immediate predecessor times a constant growth rate, one can write the aggregate savings rate as

$$\left(\frac{S}{Y}\right)_t = \frac{\sum_{a=1}^{A} [1+g]^{t-a} \left(e^{\gamma_{ay}} - e^{\gamma_{ac}}\right) n_{at}}{\sum_{a=1}^{A} [1+g]^{t-a} \left(e^{\gamma_{ay}}\right) n_{at}} \tag{10}$$

These assumptions would allow us to make forward (and backward) predictions of the savings rate for any period for which we have population proportions.

The estimates of the 'middle' (and more reliable) cohort effects strongly suggest a more or less constant growth in the cohort effect of around 2% per annum, which is, of course, close to the long-run growth rate of the UK economy, and is consistent with the Deaton and Paxson (2000) view. However, we have not adopted their approach because of the clear evidence from Figure 4 and from the tests reported in Demery and Duck (2001) that whilst their *growth* rates might be very similar, the *levels* of the income and consumption cohort effects are quite different.[25]

As an alternative we have adopted the following approach. We take our original estimates of the 1936 cohort effects for income and consumption as our base. This is the cohort which is at the centre of those cohorts appearing in *each* FES over the period 1969-98. We then calculate the average growth (or change in the log) of these estimated cohort effects over the n years before and after 1936. We considered a range of values for n but always

[24] The same argument does not apply to the various *age* groups since each age group appears in each FES data set.

[25] Demery and Duck (2001) estimated age effects after imposing the restriction that cohort effects for income and consumption are the same. They found that the estimates of age effects were largely unaffected but that the restriction itself could be formally rejected. One possible reason for this is that the assumptions about the bequest motive - that bequests are a fixed proportion of lifetime resources - is invalid.

one such that the first and last cohorts used to calculate the average growth appear in all 30 FES surveys. So, n never exceeds 17 since the 1953 cohort is the last cohort to appear in all FES data-sets. In fact the choice of n made little difference to the estimated growth rates, which ranged from 1.8% to 2% per annum. We then applied our preferred estimated growth rates of 1.9% forwards and backwards to generate our new estimates of γ_y and γ_c. In Figures 5 and 6 we show the relationship between these (re-)estimates of γ_y and γ_c (with an assumed growth rate of 1.9% per annum) and our original estimates. Clearly the re-estimates track the behaviour of the originals well until the last (and possibly erratic) twenty or so cohorts.

By projecting these growth rates forward (or backward) we can derive estimates of y_{ab} and c_{ab} for any cohorts in the future and the past.[26] We can then use equation (1) to explore the implications for the aggregate savings rate of any past or future changes in the age-structure of the population.[27]

This procedure maintains the strong assumption of the lifecycle model that there is no interaction between age and cohort effects, i.e. that the age profiles for different cohorts are parallel. The fact that the more-recent estimated cohort effects are not only erratic but are also always below our estimated trend may signal that this assumption is incorrect. For example, if, because more of them go to university, the more-recent cohorts have income age profiles with less weight at the start of their life cycles, we would expect to find declining cohort effects when, in our estimation procedure, we restrict the age effects to be cohort-independent. The cohort effect will be 'picking up' a distortion to the age profile for the more recent cohorts.

However, the size of the dip below trend, and the complete absence of any such dip throughout the 'middle period', suggests to us that this is implausible. An alternative possibility is that the more recent cohorts were coming onto the labour market at a time of recession: the estimated coefficients on the year dummies in the income equation - which sum to zero overall and are assumed to be independent of both cohort and age effects - declined over the period 1990-1994 and are all negative for the period 1992-1997. If new labour-market entrants are *disproportionately* affected by recession then our estimates of their cohort effects may be biased downwards because the cohort estimate is picking up this distortion. So this distortion - together with the small number of observations on which to estimate their particular cohort effect - may explain why the estimates for the later cohorts

[26] One practical restriction that this method imposes is that the growth rates of the income and consumption cohort effects must be the same. If they are not, the difference between them dominates the projection of the savings rate. In practice, the estimated growth rates were very close.

[27] The choice of the year of birth on which to base these re-estimates does have implications for the savings rate we predict. As a close examination of Figures 5 and 6 reveal, the 1935 cohort effect is low. We did carry out the estimates reported below taking the 1935 cohort as our base. The general behaviour of the agggregate savings rate was much the same as that reported below but was approximately 1 percentage point lower.

are both erratic and low. And, on this latter interpretation, the assumption of independence between age and cohort effects, which we use to estimate future cohort and age affects, can be maintained.

4 The savings rate and changes in the UK's age-structure

Having explained our approach we now present our estimates of the extent to which changes in its age-structure have affected the UK's aggregate savings rate in the past and are likely to in the future.

4.1 Age structure and the savings rate 1855-2001

Figure 7 shows the actual behaviour over the period 1855-2001 of four broad age-proportions: the young, 15-29 year olds; the early middle-aged, 30-44 year olds; the late middle aged, 45-64 year olds; and the elderly, those aged 65 and above.[28] They clearly show that, over the last 100 years, the population has noticeably aged - in fact to an even greater extent than it is projected to over the period 2001 to 2040. The proportion of the population aged over 65 has increased almost monotonically, with a marked increase in trend after 1900. It rose from 7% in 1930 to 14% in 1975. The trend decrease in the proportion of the population in the 15-29 age group has only been interrupted by the 'baby boom' after the second World War.

The thinner line in Figure 8 shows the behaviour of the aggregate personal savings rate - constructed by the authors from *macroeconomic* data sources and in its later years equal to the NS personal sector savings rate - over the same period. Its behaviour has clearly been dominated by the two World Wars and by short term fluctuations. The thicker and dotted lines show the savings rate 'predicted' by applying the techniques explained in the previous sections of this paper to Model A and Model B definitions of income and consumption respectively.[29] Clearly, this figure suggests that the effects of changes in age-structure are slow-moving but are nevertheless detectable but are liable to be dwarfed and obscured by much sharper shorter-term movements, especially the effects of the two World Wars. The most obvious effect is the substantial long-term rise over the first half of the

[28] The data for the population and savings series were derived from published sources by the authors and are explained fully in Demery and Duck (2003). The early population figures involved a certain amount of interpolation and so the proportions appear a little smoother than they probably really are.

[29] Of course for all the pre-1936 cohorts we we have applied the growth factor -1.9% rather than +1.9%. The projections shown for the 1969-1998 period ignore the estimated effects of the year dummies for those years.

For the early part of this period the population proportions were only available in 5-yearly groups. To construct an aggregate savings rate we had then to apply these proportions to the *average* income and consumption of each 5-year age group predicted using our FES-based estimates of the yearly age and cohort effects.

20th century, a slight fall over the next 20 years, and a rise until the end of the century.

Given the shape of the savings-age profiles shown in Figure 3 and the historical behaviour of the age-structure shown in Figure 7, our interpretation of these movements is that they were primarily the result of the prolonged rise in the proportion of the middle-aged from 1900 to the late 1950s, the fall in that proportion from then until the late 1980s, and its subsequent rise. The opposite movements in the proportion of the population who are young reinforced these effects on the aggregate savings rate. The somewhat smaller rise in the savings rate associated with Model B from 1900 to 1955, the sharper fall from the late 1950s until the late 1980s, and the more muted rise since then are mainly the result of the key difference between Models A and B: that in Model B an increase in the proportion of the elderly leads to a larger fall in the economy's overall savings rate.

4.2 The savings rate under alternative population projections: 2001-2040

In this section we investigate how the aggregate savings rate might change over the next 30 to 40 years in response to a variety of different projections about the age-structure of the population. The projections are the 2000-based projections of the UK Government Actuary's Department:[30] they cover the years (2000 2001) 2002-2026, 2031, and 2036. To obtain an appropriate projected age-structure we then express each age group's projected population as a proportion of the total projected population 16 years old and above.

We begin with a brief discussion of the assumptions behind the projections we use. The population projections are based on *cohort* fertility rates, which are more stable than those for calendar years. So the fertility rates used in the GAD projections are derived from assumptions relating to the year in which women were born. For the United Kingdom as a whole, completed family size has been falling steadily from an average of around 2.45 children for women born in the mid 1930s. The family sizes to be achieved by younger cohorts are highly conjectural, but for the Principal Projection GAD assumed that average completed family size, for the United Kingdom as a whole, will continue to decline until around the 1985 cohort and eventually level off at 1.74 children per woman. In its 'High Fertility' and 'Low Fertility' variants, GAD assumes a levelling off at 1.94 and 1.54 children respectively. The Principal Projection is based on continued increases in life expectancy at birth, which is projected to rise to 79.7 years in 2040 for males and to 83.9 years for females. The respective 'High Life Expectancy'

[30] GAD have also recently produced a 2001-based set of projections. As we wished to consider the 'Very High Life Expectancy' variant and as this was not available for the 2001-based projections, we use the 2000-based set of projections.

and 'Low Life Expectancy' variants assume male life expectancy at birth to rise to 81.9 and 77.5 years in 2040 and the corresponding figures for females are 85.5 and 82.2 years. GAD also produced a 'Very High Life Expectancy' variant which assumed that life expectancy at birth rises to 83.5 years for males and 87.1 years for females in the year 2040. This variant is assumed to be broadly consistent with the high sensitivity variant produced by Eurostat in 2000. To cover the broadest range of possible demographic projections we consider four cases: the Principal Projection (i.e. GAD's 'central' projection), a 'Young' Projection (combining the High Fertility and Low Life Expectancy variants), an 'Old' Projection (combining the Low Fertility and High Life Expectancy variants) and the 'Very High Life Expectancy' variant (which combines the Eurostat very high life expectancy case with fertility assumptions used in the Principal projection).[31] These are scenarios that are likely to lead to the widest range of population age-structures we can expect over the coming decades.

In Table 1 we present the main features of the Principal Projection and the variants we consider. The UK population rises to 65.8 million in 2040 under the assumptions used in the Principal Projection. The projected population in 2040 is 8% higher (71.4 million) under the Young variant and 8% lower under the Old variant. The proportion over 65 years old projected in 2040 ranges from 21% (Young variant) to nearly 28% (Old variant). To illustrate the broad implications of these projections we graph in Figure 9 the age-structure implied by the Principal Projection.[32] We have aggregated the projections into four broad age-groups: the young, 15-29 year olds; the early middle-aged, 30-44 year olds; the late middle aged, 45-64 year olds; and the elderly, those aged 65 and above. The central feature of the figure is that the proportion who are elderly will increase from below 20% in 2000 to almost 25% in 2040; the increase is marked up to the mid-2030s when it slows down somewhat. The proportion who are late middle aged is projected to increase for almost 20 years and then decrease and the proportions who are young or early middle-aged are expected to fall.

In Figure 10 we plot the savings rate implied by combining the Principal Projections and the Model A and Model B age and cohort effects over the period 1971 to 2040. The predicted series up to 2000 are those implied by our estimates of the age and cohort effects and the actual population proportions. (We ignore the effects of the year dummies for the purposes of these projections, though they were included in the regressions from which the age and cohort effects were derived.) We also include in the figure the

[31] The Principal Projection assumes a long term annual net migration of +135,000 from 2002-03 onwards and this assumption is also used in the Very High Life Expectancy variant. The Young projection is based on a higher long-term net migration figure of +195,000 and the Old variant assumes a lower net migration figure of +75,000.

[32] The projections are annual up to 2026 and then for the individual years 2031, 2036 and 2040. The graph linearly interpolates the shares over the period 2026-40.

relevant National Statistics personal sector savings rate.[33]

As the Model A savings rate is calculated on the assumption that all pensions are treated as income, it is always higher than that of Model B, where only a part of the pension is treated as income. As is clear from the figure, in both models demographic changes anticipated in the Principal Projection imply changes in the aggregate savings rate that are quite modest in comparison to the actual year-to-year fluctuations that occurred from 1971-1998. In Model A, the personal sector (long-run) savings rate is predicted to rise gradually from 5.8% in 2000 to 6.5% in 2025 and 7.0% in 2040.

We interpret these predictions as follows. Our estimates of the age effects, as illustrated in Figure 3, suggest that increases in the proportion of the population who are elderly will reduce Model B's aggregate savings rate but will either leave Model A's rate unchanged - if it is offset by a fall in the proportion who are late middle-aged and if the elderly themselves are not mainly over 70 - or will raise it. A rise in the proportion who are in their late middle age will generally raise the aggregate savings rate in both cases. Variations in the other age proportions are likely to have only modest effects on aggregate savings rates - unless changes in the young are due to changes in the very young. In fact the proportion who are elderly is predicted to rise continuously over the next 40 years whereas the proportion who are middle-aged is predicted to rise for the first 20 years and then fall back. The predicted rise in the proportion who are middle-aged is sufficient to induce a rise in the predicted savings rate for both models, but, when this is reversed, the continued rise in the proportion who are elderly is sufficient to reduce the predicted aggregate savings rate in the case of Model B but not in the case of Model A.

Figures 11A and 11B show the savings rates we predict on the basis of the variant population projections (see Table 1): the rates in Figure 11A are based on Model A and those in Figure 11B on Model B. A striking feature of the predictions using Model A is the higher savings rates predicted from the Old and Very High Life Expectancy variants: they are 0.5 percentage points higher in 2040 than those from the Principal Projection. And the savings rate predicted using the demographic structure of the Young variant is around 0.5 percentage points lower than the Principal Projection. This result is unsurprising given the savings-age profiles for Model A shown in Figure 3 where the very old are predicted to have higher savings rates and the young to have very low savings rates.

The predicted savings rates in Figure 11B are based on Model B and they suggest that the different demographic structures implied by the three variants lead to virtually no change in the savings rates. The stretching of the vertical axis somewhat exaggerates the differences, but in the year

[33] The NS rate is adjusted in the way described in Section 2.1.

2040 the savings rates are all in the range 2.8-3.0%. The Young variant anticipates a smaller proportion of old (dissavers) but a higher proportion of 16-29 year-olds, whose savings rates are also low (and negative in the age range 16-19 years). The net effect is to leave aggregate savings largely unaffected. In the Old variant the proportion of the elderly dissavers is higher, but the proportion of the low-saving 16-29 year group is lower, again leaving the aggregate savings rate roughly the same.

Overall, our results suggest that changes to the age structure that are projected to occur over the next 40 years are likely to have modest effects on the personal sector savings rate. Because of the way pension contributions are treated in our data sources, our estimates probably *underestimate* the savings of working-age individuals. For this reason, the overall conclusion we would draw from our results is that *both* our measures of the savings rate are likely to rise modestly over the next 40 years in response to the projected changes in age structure. This is especially likely for the first half of that period when the proportion of the late middle-aged is projected to rise to a peak of approximately 27%, and especially true of the measure of the savings rate which does not make any adjustment for the treatment of pension income. However we would not expect any rise to be higher than 1 or 2 percentage points in either case. And we find no basis for the view that the aggregate savings rate will decline because of the anticipated ageing of the UK population.

5 Conclusions

From micro-economic data we have derived estimates of the relationship between the UK personal savings rate and the age-structure of the population. Our results strongly suggest that changes in age-structure have had detectable, sustained but modest effects on the savings rate over the last 150 years and they are likely to continue to do so. The source of this influence is mainly the effect that one would expect from a simple lifecycle model: the middle aged, and especially the later middle aged, tend to save a higher proportion of their income than the young and the elderly, and so shifts between the middle-aged and the young or elderly affect the aggregate savings rate. However, our estimates suggests that even these effects are quite modest, at least by the standards of the fluctuations seen in the actual savings rate over the previous 150 years. They also suggest that the projected changes in the UK's age-structure over the next 20-40 years will lead to a modest increase in the savings rate, especially in the first half of that period of no more than 2 percentage points.

Data Appendix

Our data are derived from the Family Expenditure Survey (FES), an annual cross-section survey of around 7,000 UK households (or around 20,000 individuals) which has been extensively used in the analysis of consumption behaviour.[34] It collects detailed information on household and individual consumption and income, and on household and individual characteristics. Until 1993 the FES organised the data on a calendar year basis; since then the data are organised on a financial year basis. 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 and benefits but it excludes contributions to the state-funded pension scheme.

Importantly, 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) 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 in the main body of the paper.

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.

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 all household members.

[34]Notable recent examples are Banks et al. (1998), Paxson (1996) and Deaton and Paxson (1997). The FES 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) suggest 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).

References

Atkinson, A. and John Micklewright (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 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. Turner (1998), 'Is there a retirement-savings puzzle?', *American Economic Review*, 88, 769-788.

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

Chesher, A. (1997), 'Diet revealed? Semiparametric estimation of nutrient 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.

Demery, D. and N.W. Duck (2001) 'Savings-age profiles in the UK' *University of Bristol Working Paper No. 01/518.*

Demery, D. and N.W. Duck (2003) 'Income, Consumption, Savings and Age Structure in the UK 1855-2001'. *University of Bristol Mimeo.*

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

Jappelli, T. and F. Modigliani (1998), 'The age-saving profile and the life-cycle hypothesis', *Centre for Studies in Economics and Finance, University of Salerno, Working Paper No. 9.*

Mankiw, G. N. 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.

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

Table 1
2000-Based Population Projections
(Government Actuaries Department)

		Principal		Young		Old		Very High Life Expectancy	
	2000	2025	2040	2025	2040	2025	2040	2025	2040
Total Population (000)	59,756	64,836	65,837	68,011	71,367	61,808	60,686	66,155	68,225
Proportion Age 16-29	17.82%	16.08%	16.06%	16.49%	17.32%	15.68%	14.64%	15.78%	15.53%
Proportion Age 30-44	23.01%	19.91%	18.15%	19.88%	18.55%	19.90%	17.66%	19.57%	17.58%
Proportion Age 45-64	23.37%	26.50%	24.88%	25.75%	24.05%	27.23%	25.65%	26.25%	24.35%
Proportion Age 65 and above	15.59%	20.23%	24.49%	18.78%	21.42%	21.77%	27.88%	21.45%	26.66%

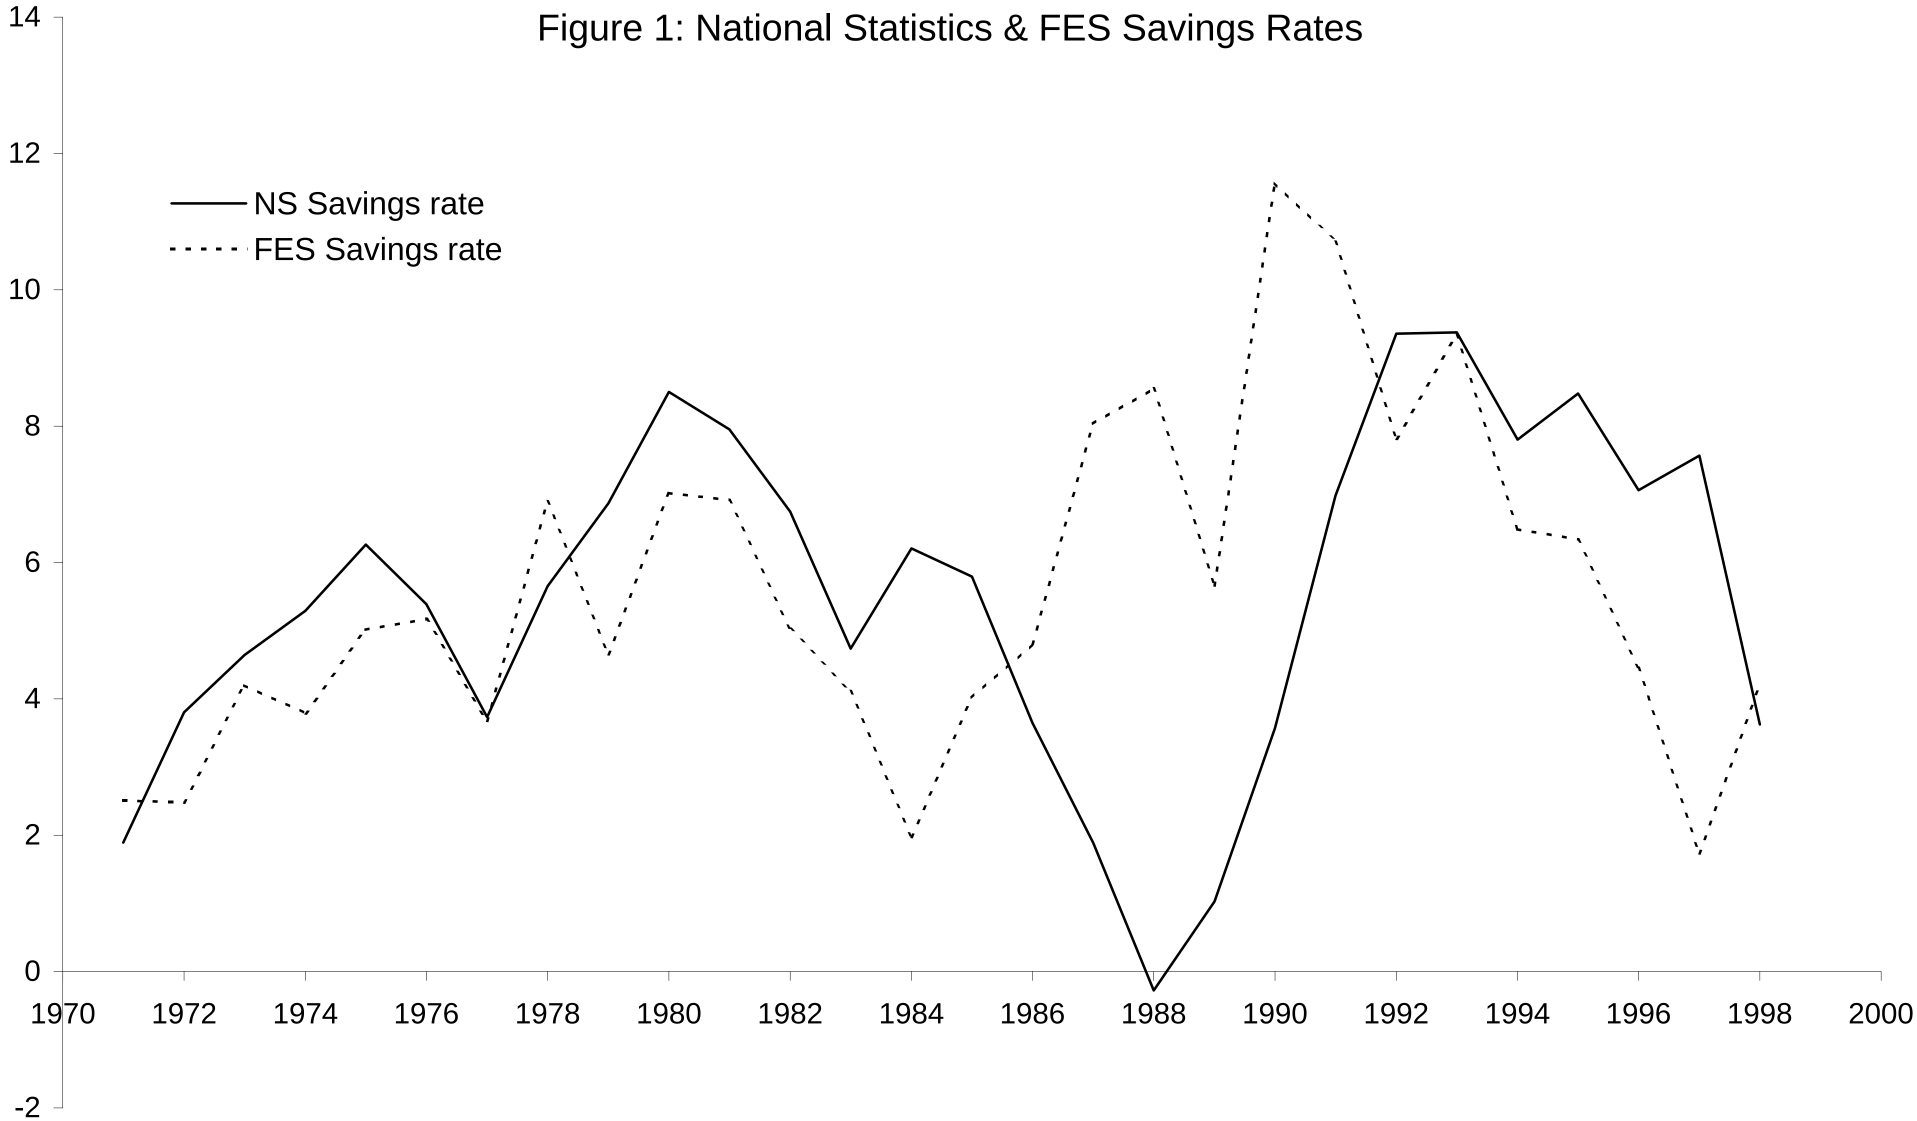
Figure 1: National Statistics & FES Savings Rates
NS Savings rate
FES Savings rate
14
12
10
8
6
4
2
0
-2
1970
1972
1974
1976
1978
1980
1982
1984
1986
1988
1990
1992
1994
1996
1998
2000

Figure 2A: Income and Consumption Age Profiles
Model A, 1953 Cohort.

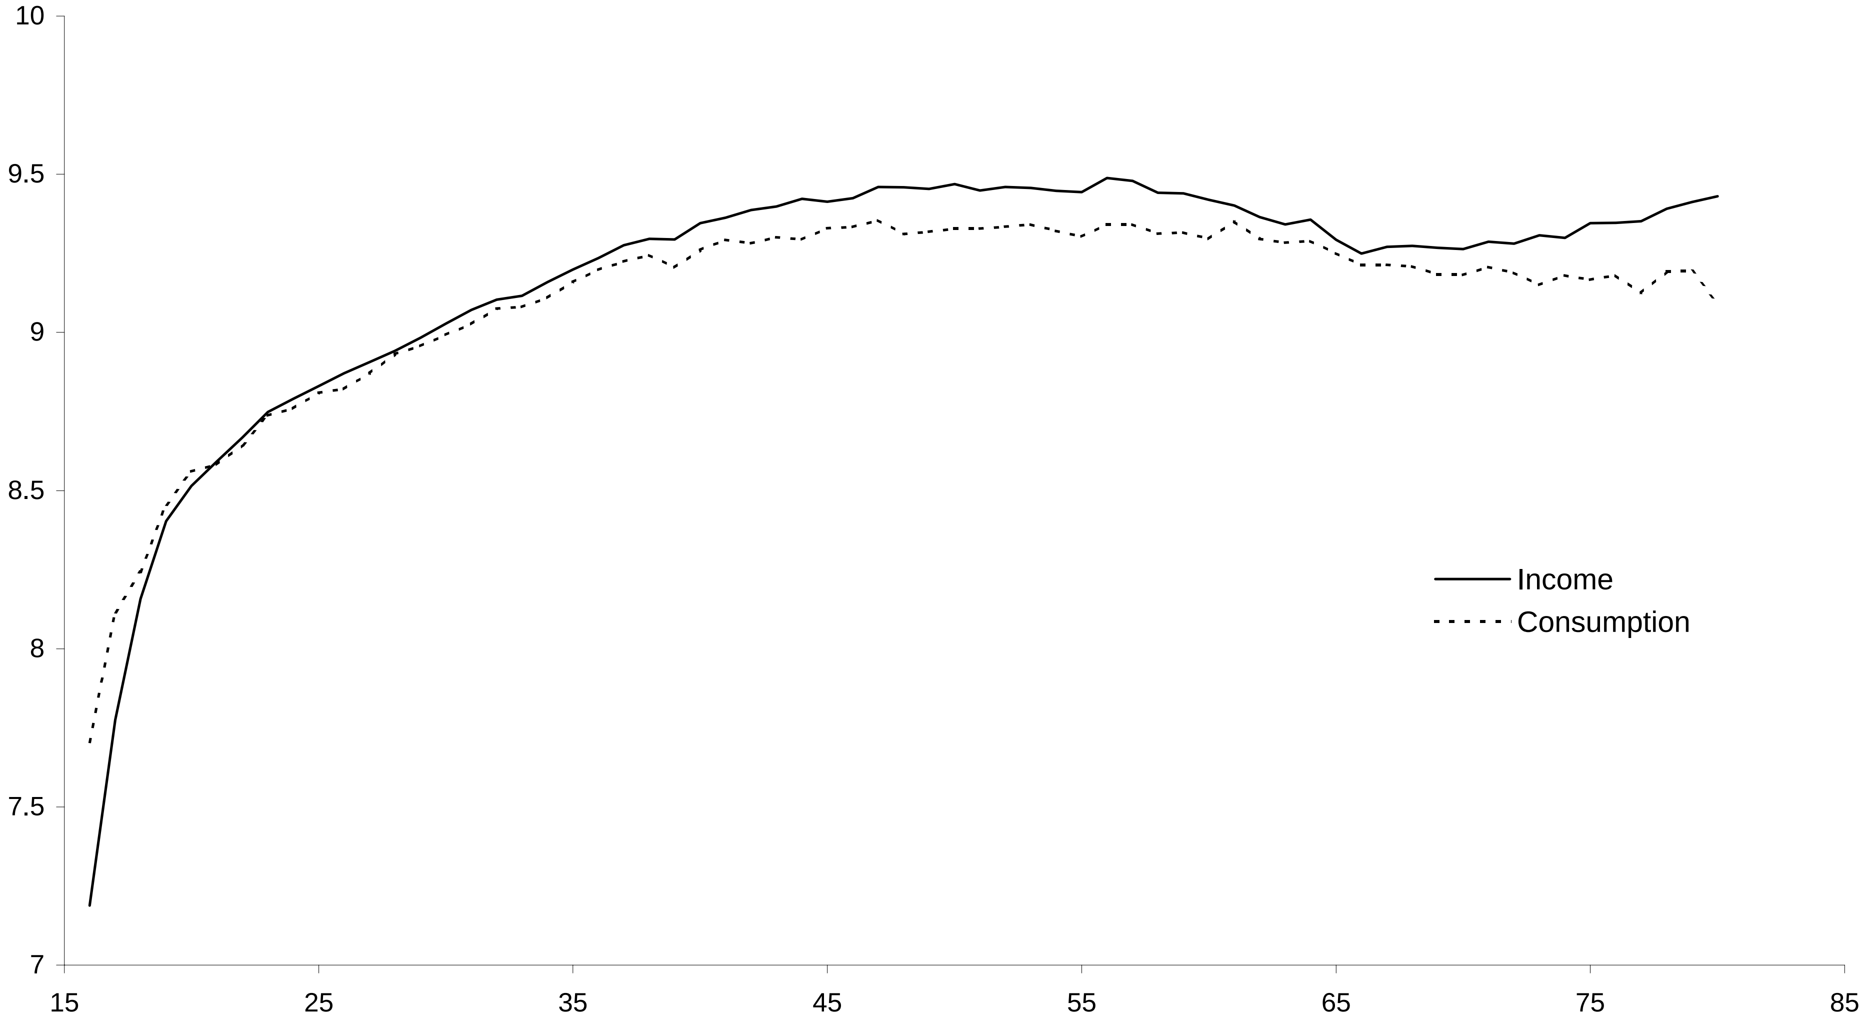

Figure 2B: Income and Consumption Age Profiles
Model B, 1953 Cohort.

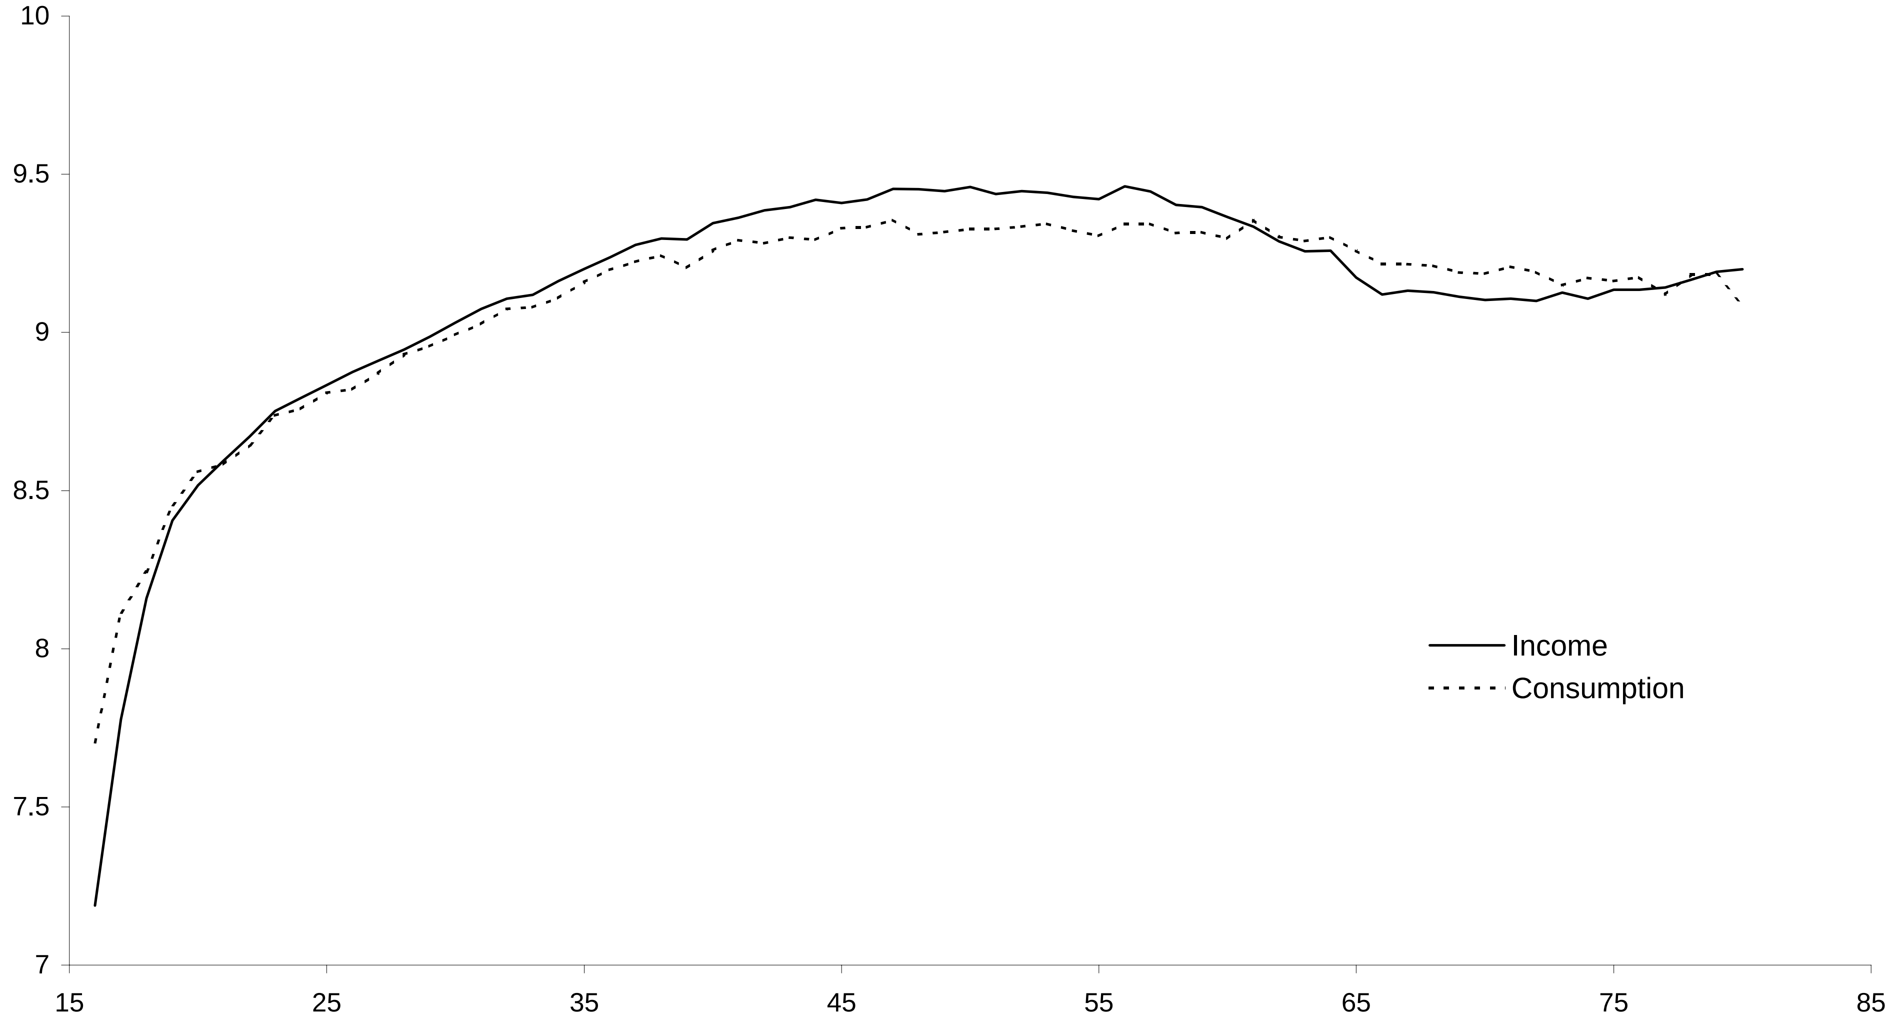

Figure 3: Savings Age Profiles
1953 Cohort.

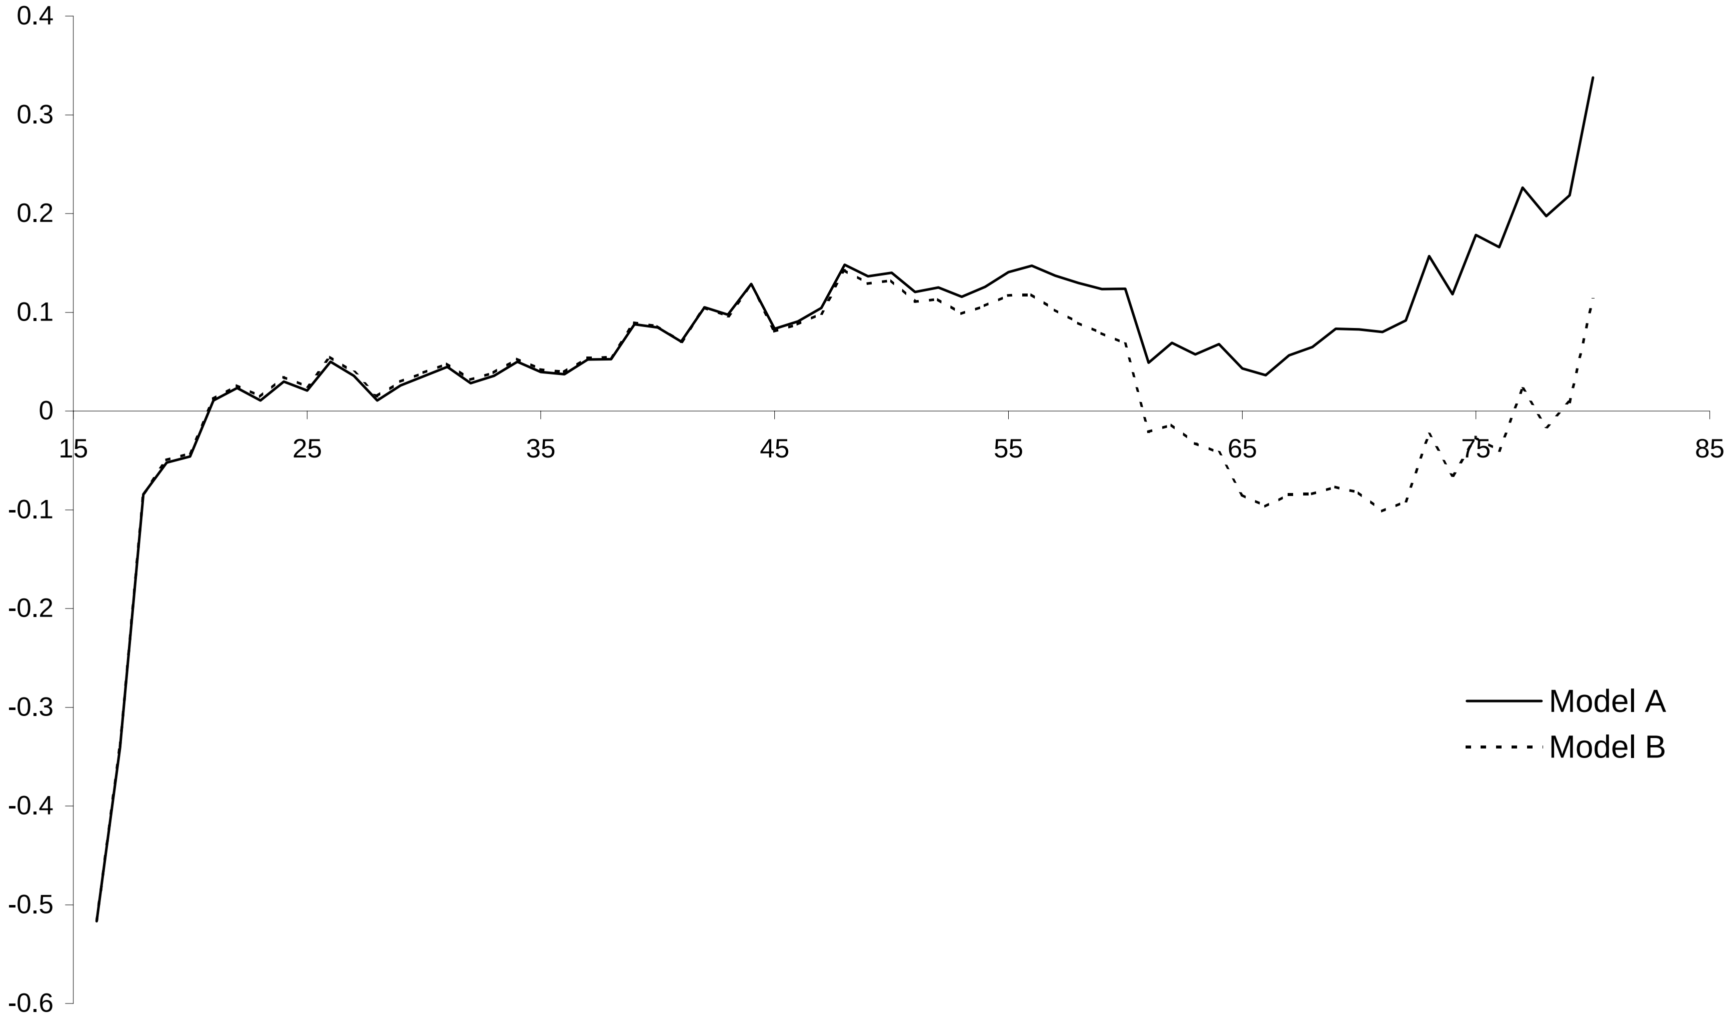

Figure 4: Cohort Effects for Income and Consumption
Model B

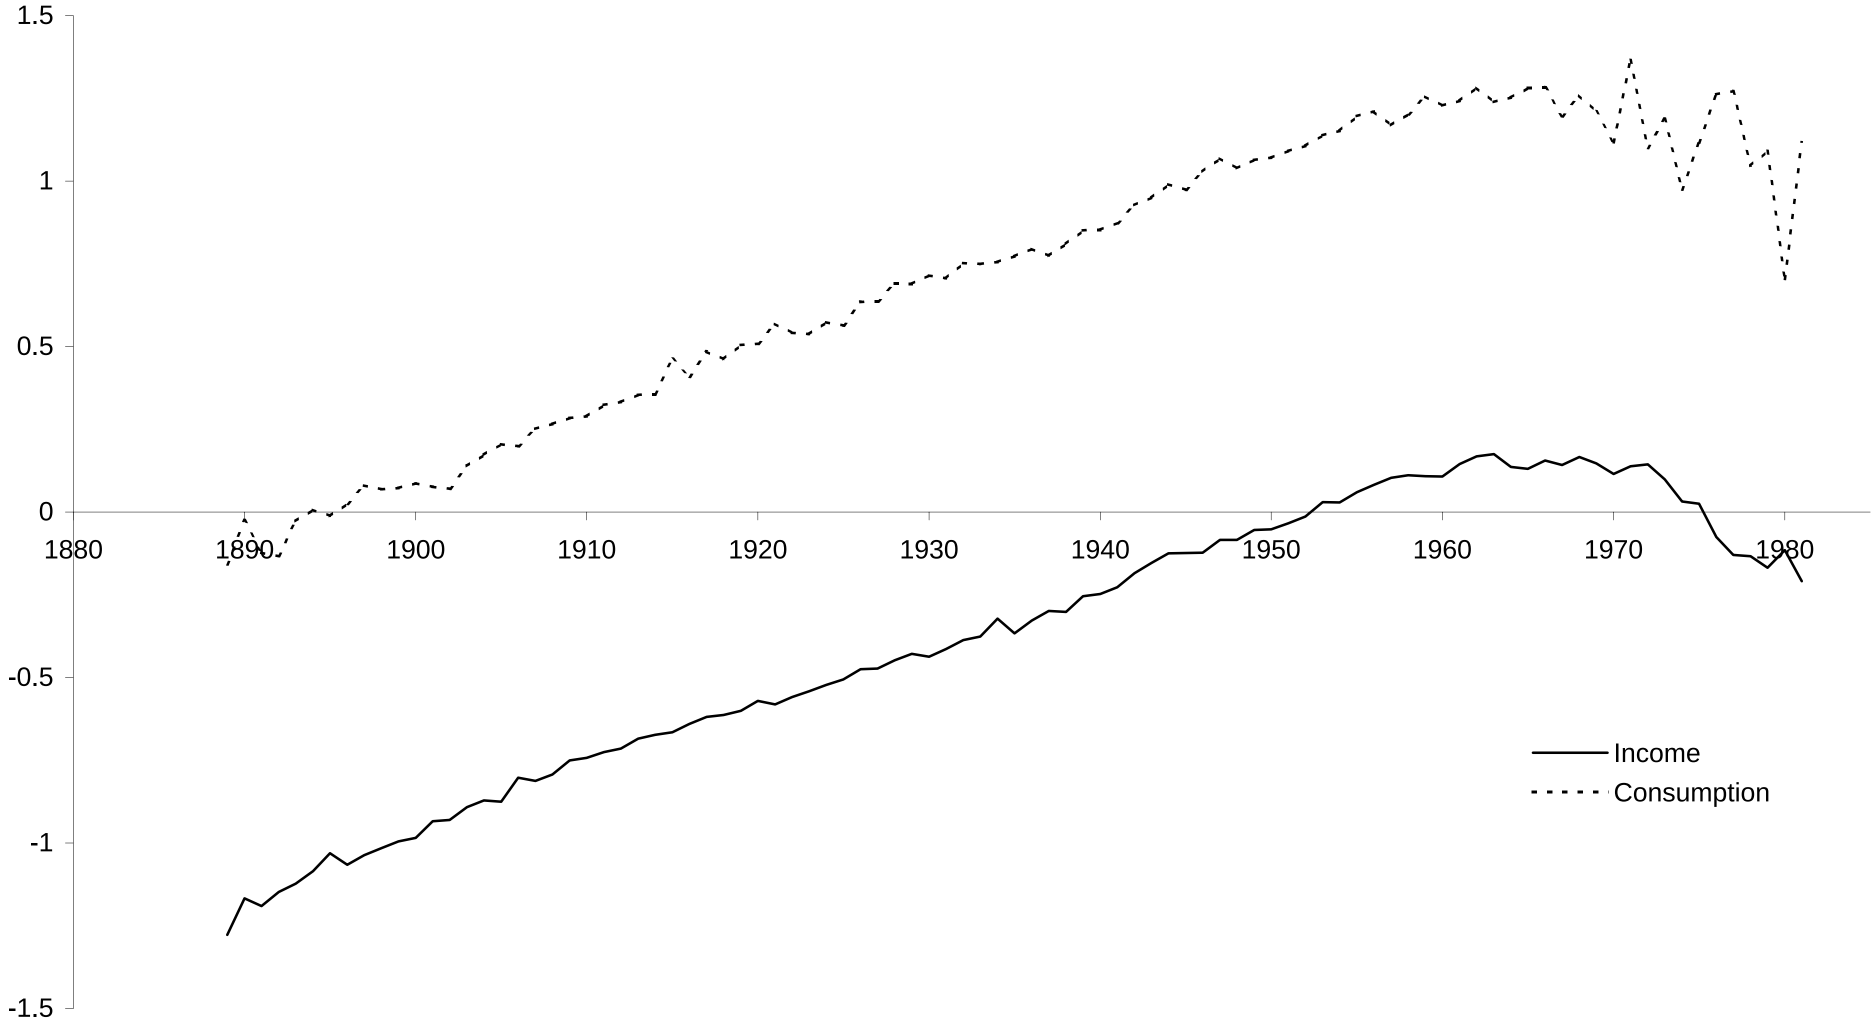

Figure 5: Income Cohort Effects: Estimated and 'Predicted'
Model A. Mean growth rate: 0.019

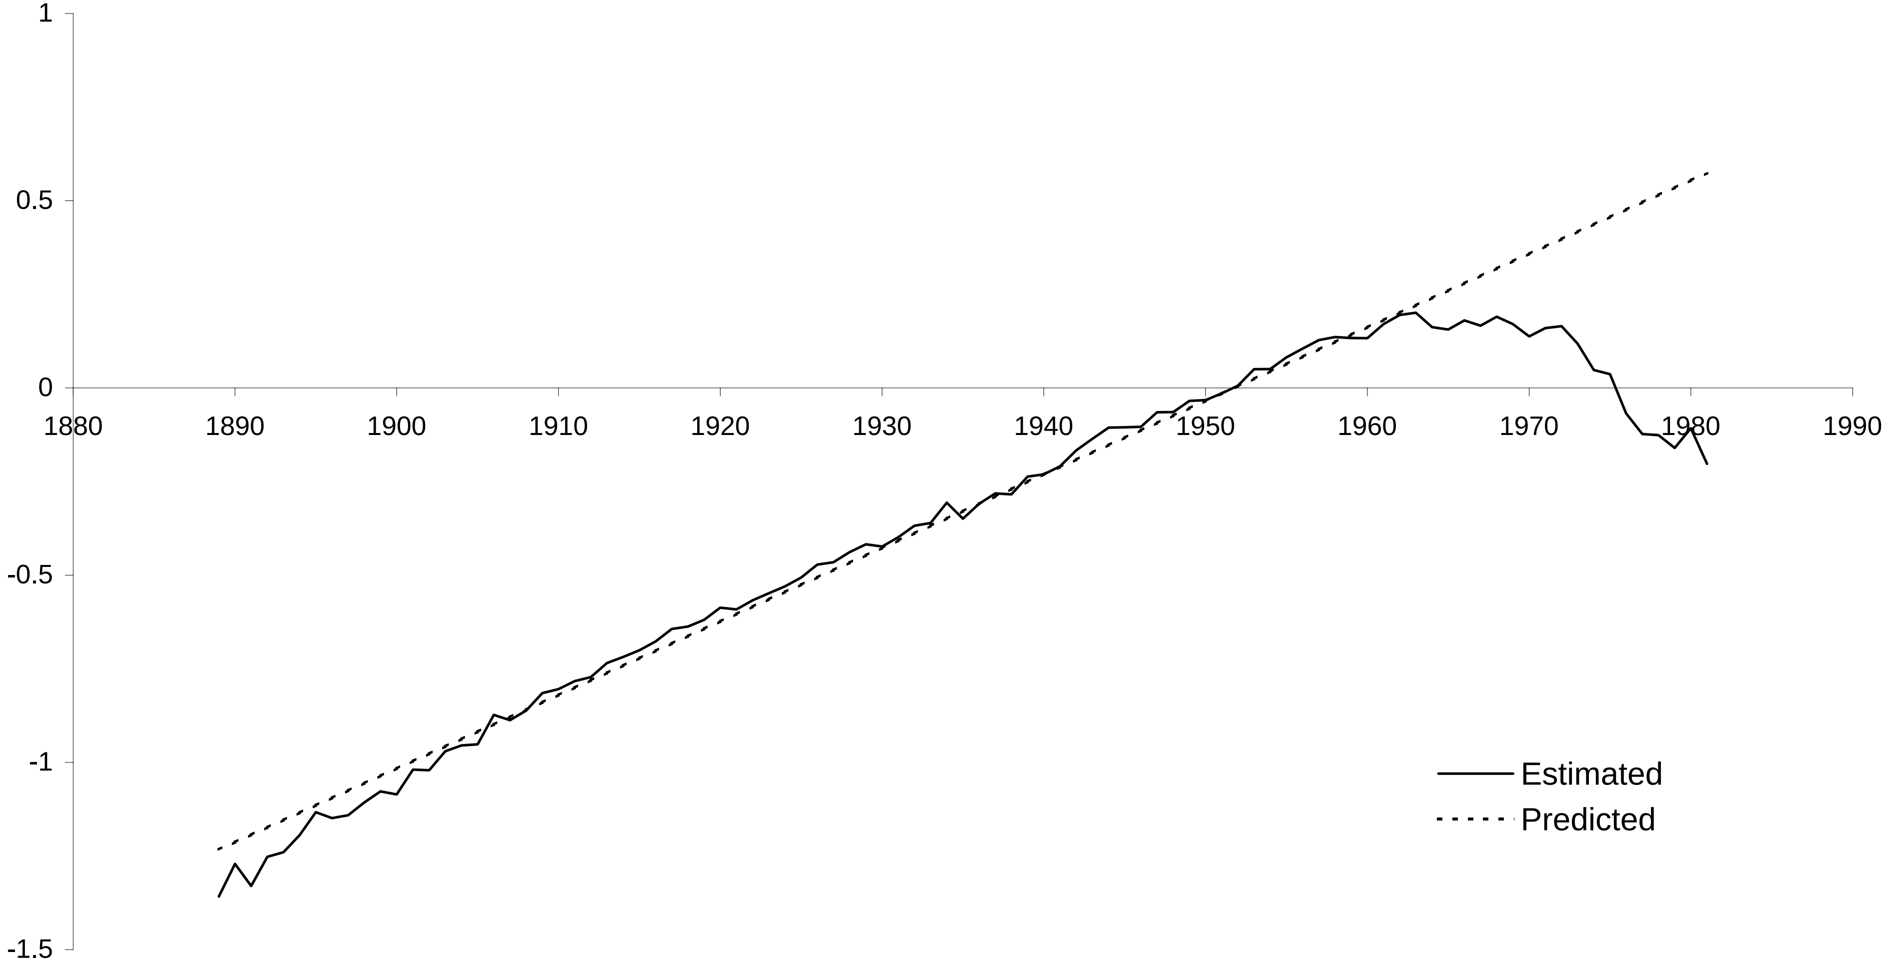

Figure 6: Consumption Cohort Effects: Estimated and 'Predicted'
Model A. Mean growth rate: 0.019

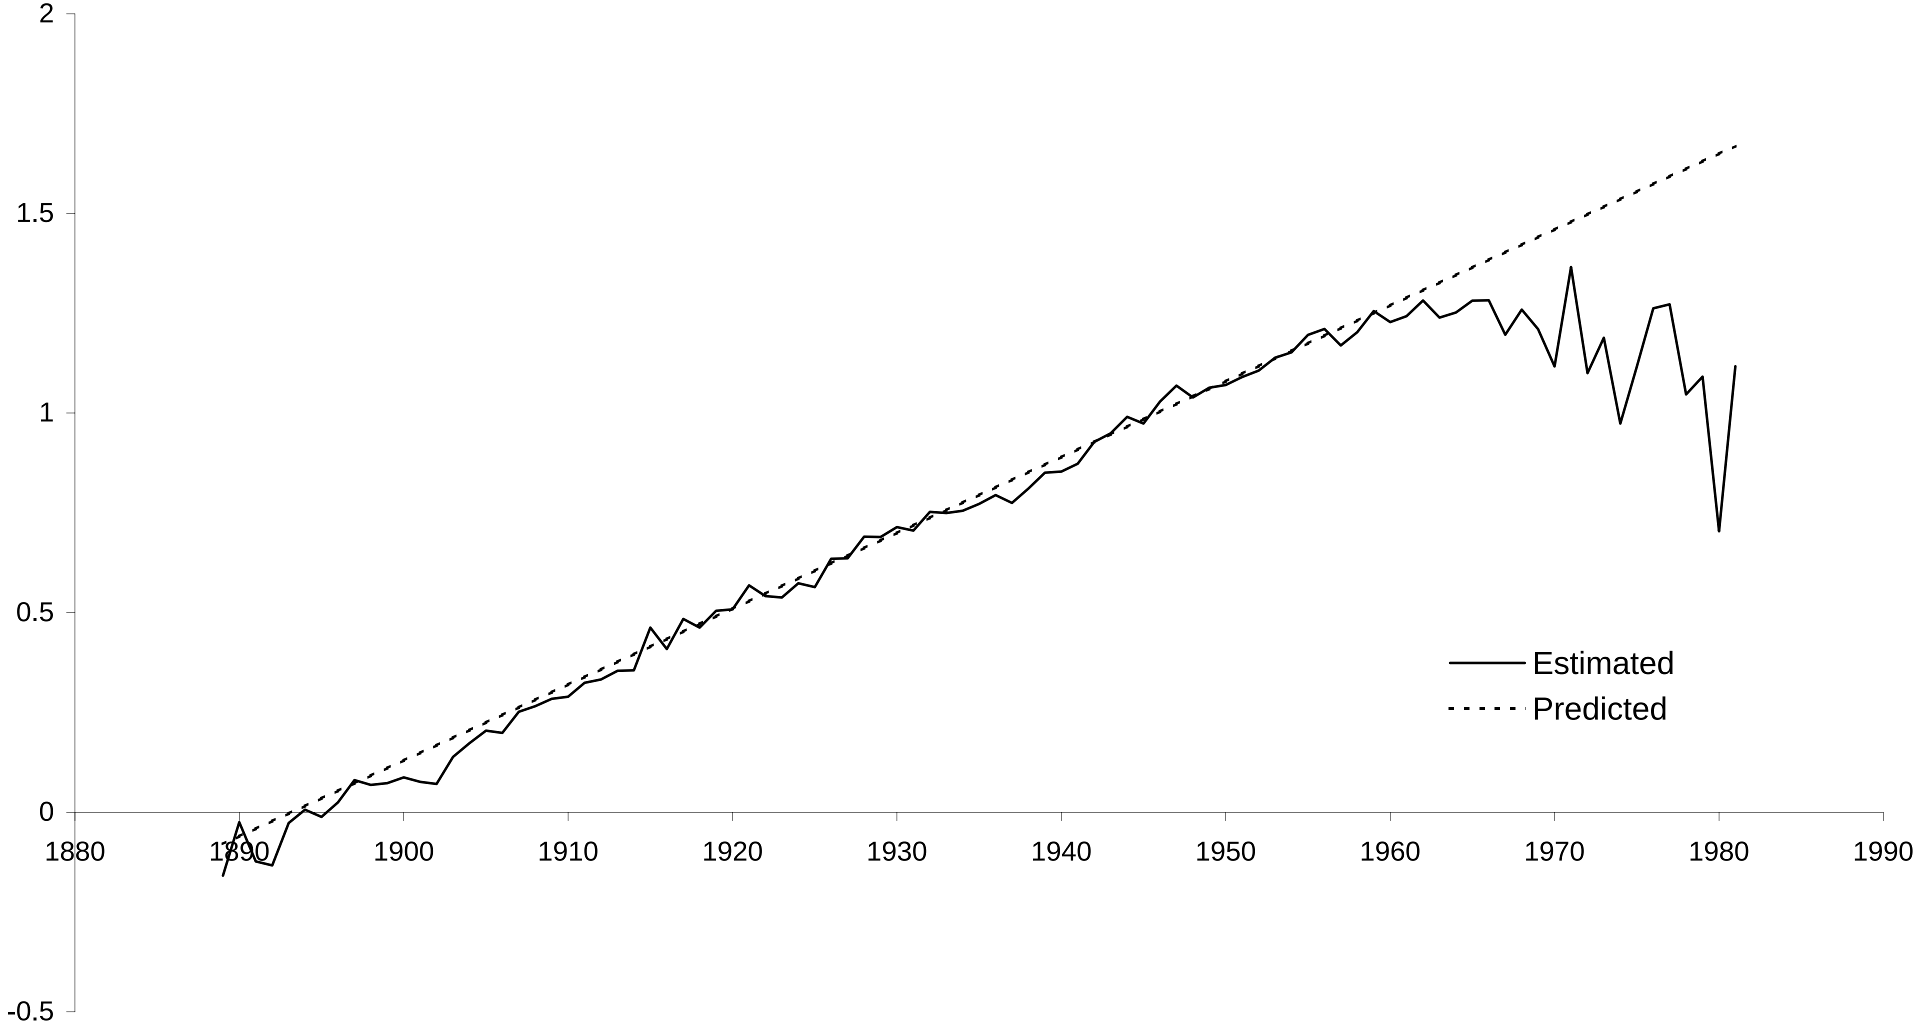

Figure 7: Age Structure of UK Population, 1855-2000

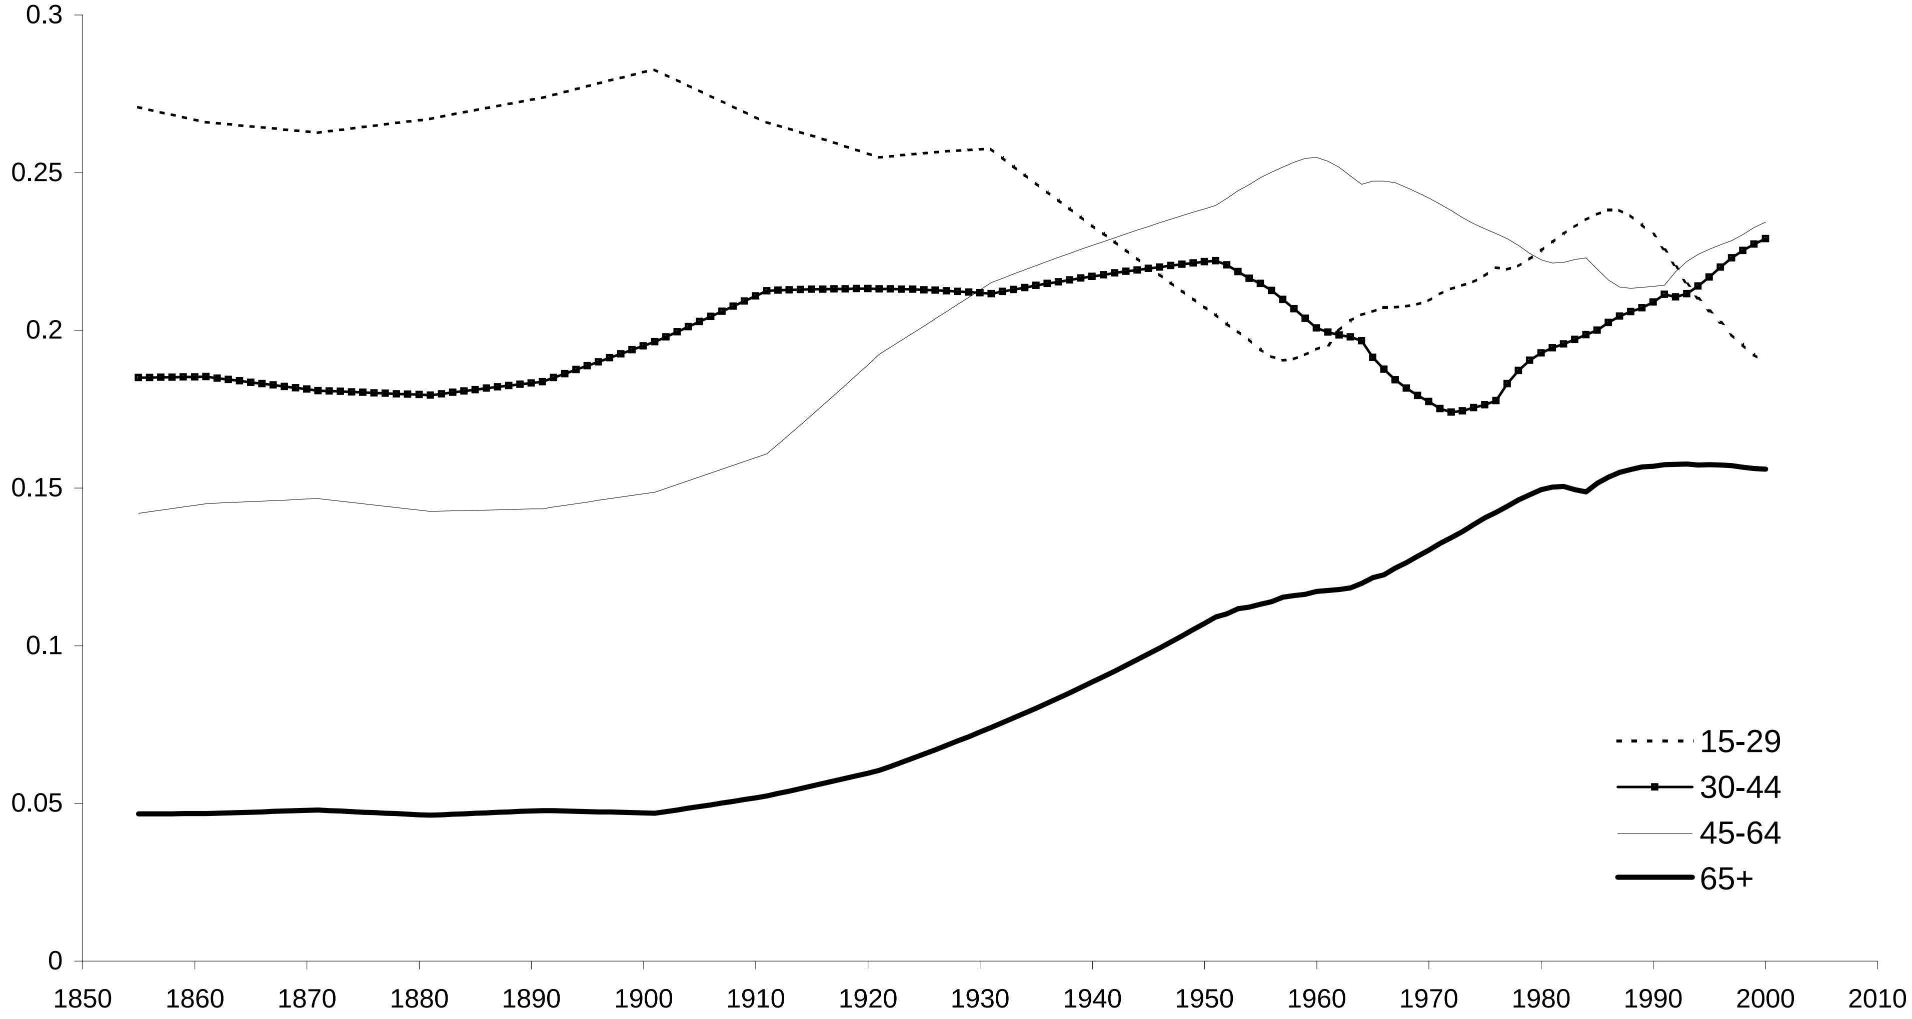

Figure 8: Actual and Simulated Savings Rates: 1855-2001

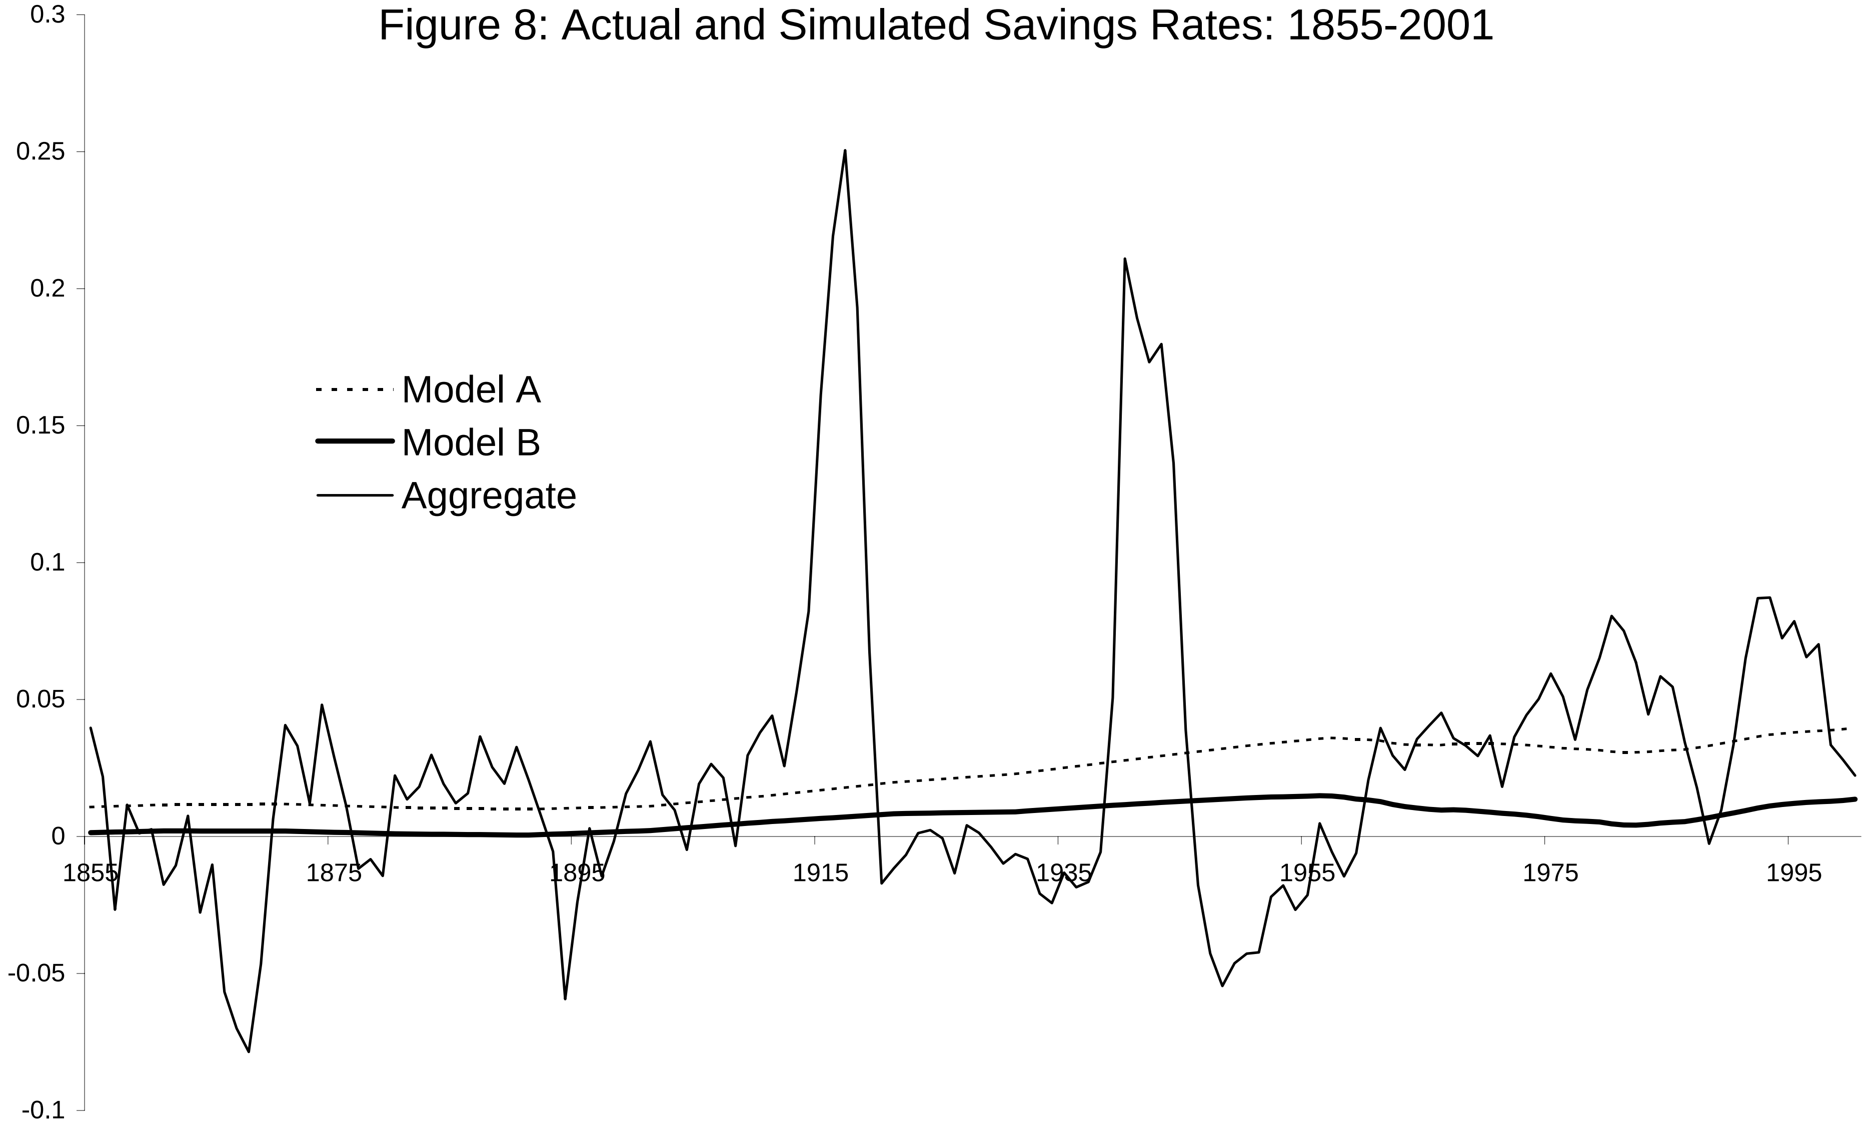

Figure 9: Age Structure of Principal Projection

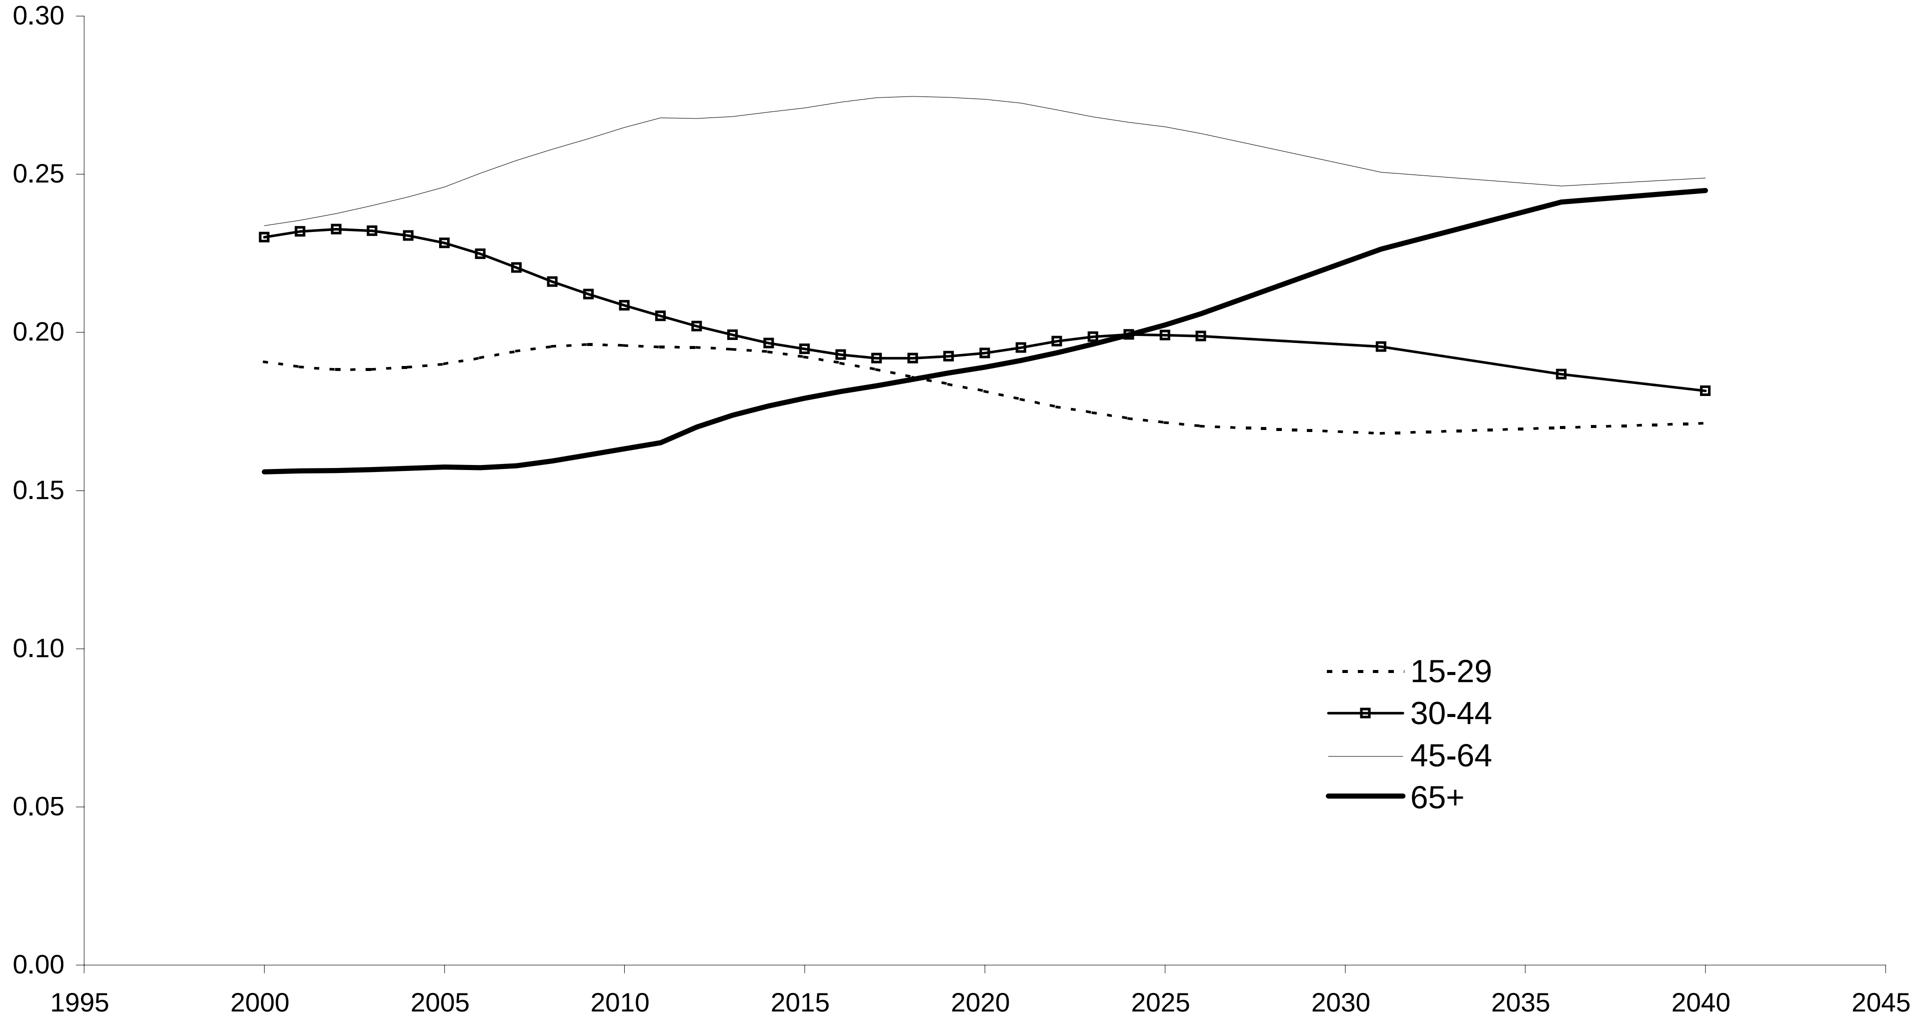

Figure 10: Actual and Predicted Savings Rate
Principal Projection 2000-2040

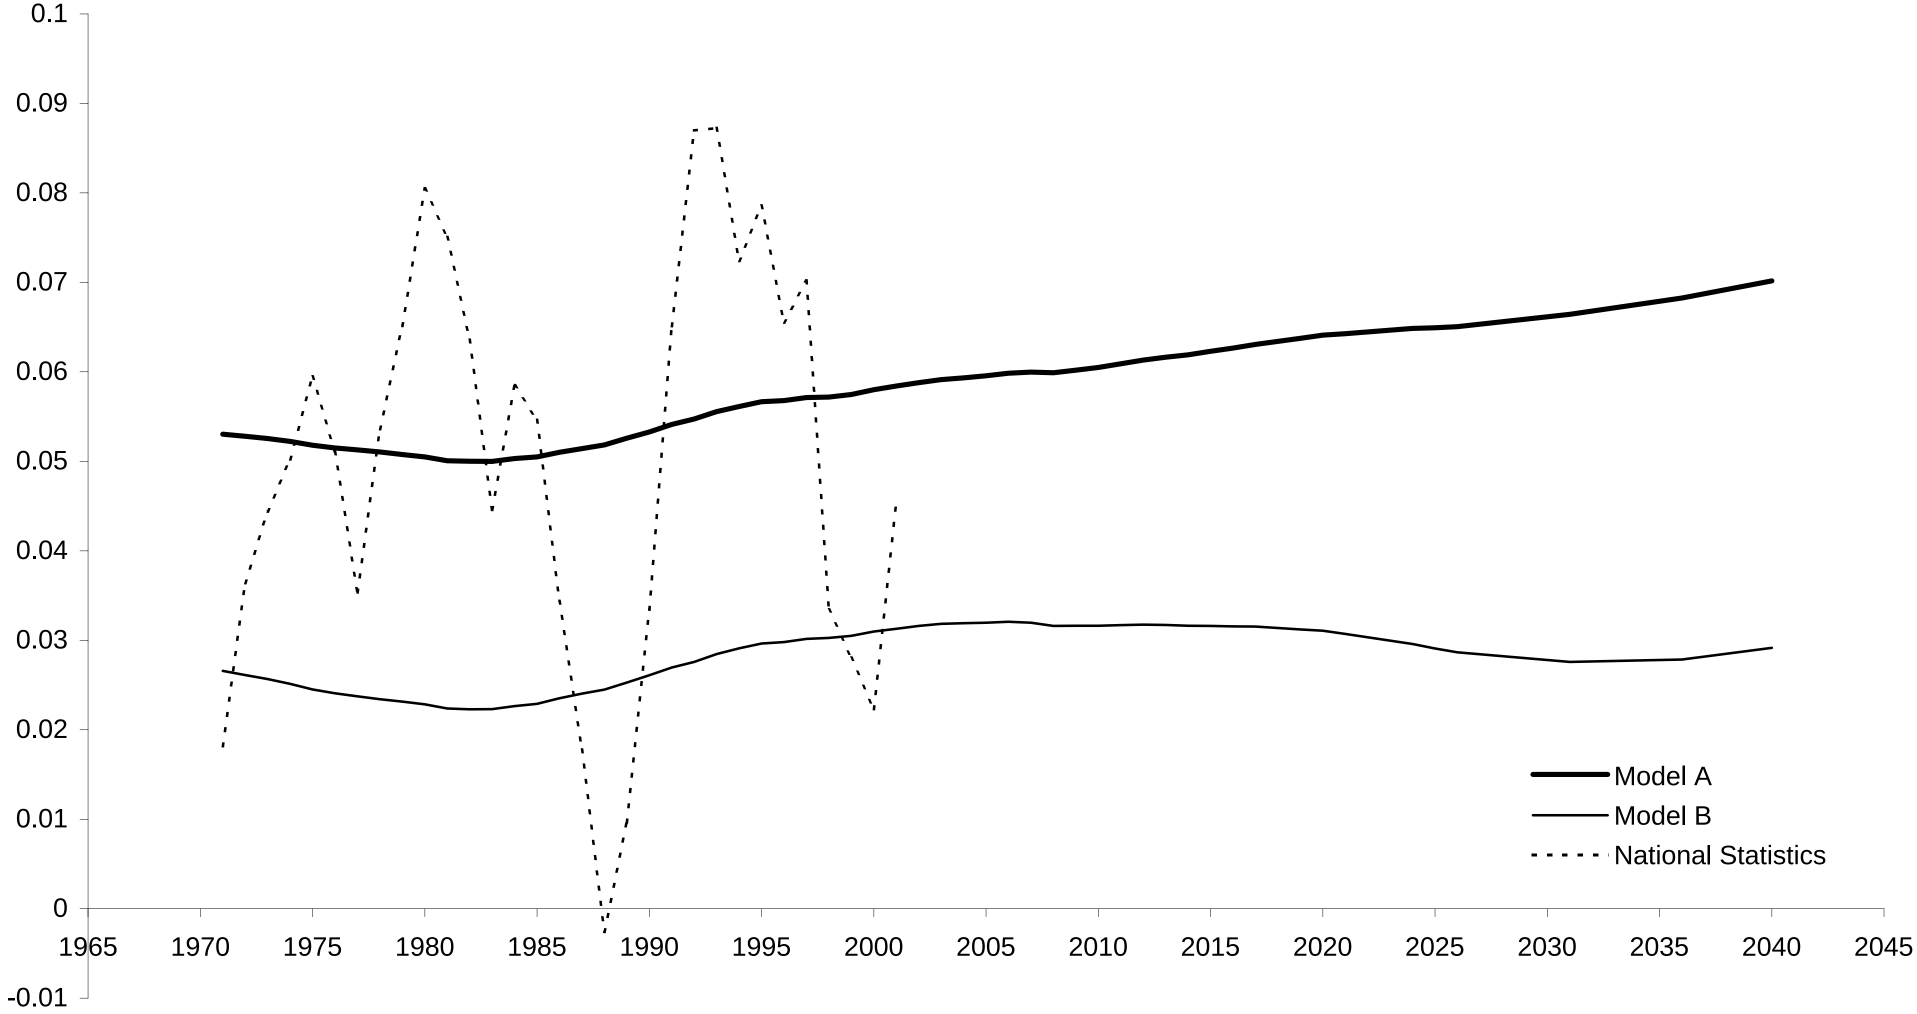

Figure 11A: Variant Savings Rate Forecasts
Model A

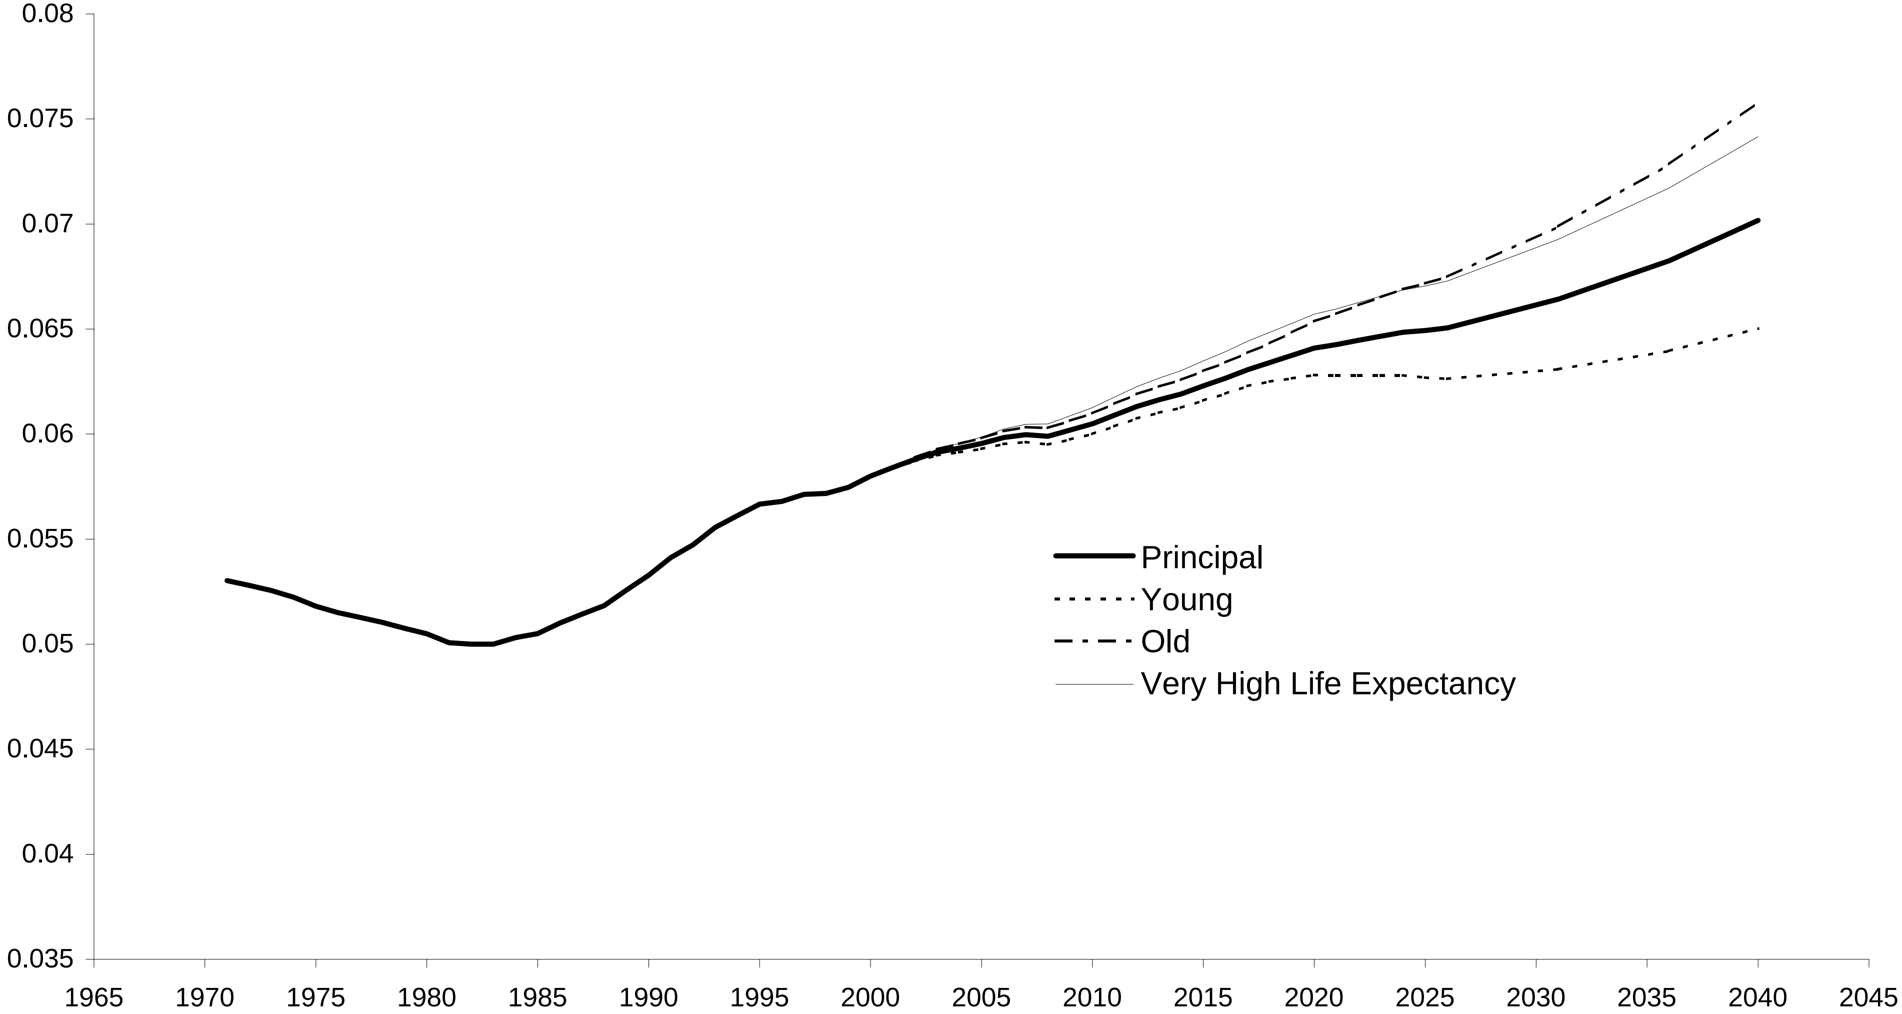

Figure 11B: Variant Savings Rate Forecasts

Model B