Institute for Fiscal Studies

IFS Briefing Note BN326

Rowena Crawford
Laurence O'Brien
David Sturrock

When should individuals save for retirement? Predictions from an economic model of household saving behaviour

Economic and Social Research Council

Nuffield Foundation

When should individuals save for retirement? Predictions from an economic model of household saving behaviour

Rowena Crawford
Laurence O'Brien
David Sturrock

Copy-edited by Rachel Lumpkin

Published by
The Institute for Fiscal Studies

ISBN 978-1-80103-037-3

This briefing note is an output from a programme of research on 'Pension Saving over the Lifecycle' that is funded by the Nuffield Foundation. Co-funding from the ESRC-funded Centre for the Microeconomic Analysis of Public Policy (ES/T014334/1) is also gratefully acknowledged.

Preface

The authors would like to thank Alex Beer, Jonathan Cribb, Carl Emmerson, Cormac O'Dea and members of the project advisory board for helpful feedback and comments.

The Nuffield Foundation is an independent charitable trust with a mission to advance social well-being. It funds research that informs social policy, primarily in Education, Welfare and Justice. It also funds student programmes that provide opportunities for young people to develop skills in quantitative and scientific methods. The Nuffield Foundation is the founder and co-funder of the Nuffield Council on Bioethics and the Ada Lovelace Institute. The Foundation has funded this project, but the views expressed are those of the authors and not necessarily the Foundation. Visit http://www.nuffieldfoundation.org/.

Executive summary

Most individuals need to save privately for retirement if they are to maintain their living standards when they stop working. There has been lots of research and discussion on how much individuals need to save, and how this compares with the saving being encouraged through automatic enrolment into workplace pensions. However, there has been little discussion of *when* individuals should save for retirement and the appropriateness of a single default contribution rate for all.

In this briefing note, we use a life-cycle economic model to illustrate that there are good reasons for saving rates not to be constant over working life, due to predictable factors that change with age. The model is a simple approximation to real life, in that individuals face little uncertainty and can only save in one asset that has a known rate of return. Individuals choose how much to spend each year, and how much to save, with the objective of smoothing their living standards over their life cycles. While necessarily simple to be computationally tractable, this model yields important conclusions with implications for the design of real-world policies.

Key findings

1. Most individuals expect some earnings growth over their working lives. If an individual is aiming to smooth their spending over their lifetime, then they should save a greater proportion of their earnings for retirement at later ages when earnings are higher, rather than saving at a constant rate throughout working life.

2. Households with children are typically assumed to require higher spending to achieve the same standards of living as those without. Given this, most parents aiming to smooth their living standards over their lifetime should save relatively more for retirement before their children arrive and/or after they have left home.

3 Many recent graduates hold student loans that will be written off after a certain period of time (30 years after graduation for those entering higher education from 2012). Graduates aiming to smooth their living standards over time should increase their pension saving by the amount of their previous loan repayments when loans are written off.

4 Employer pension contributions that are only made if the employee also contributes incentivise individuals to contribute throughout working life, even in years when earnings are relatively low or there are children in the household (or if, for some other reason, living expenses are greater). However, individuals are likely to want to contribute only the minimum required to receive the employer contribution for the first half of working life. When earnings increase and/or children leave home, they should then markedly increase their saving rate for retirement.

5 This profile of the appropriate saving rate over working life in the presence of contingent employer pension contributions – flat at the minimum required employee contribution, and then increasing markedly when children leave home – is robust to a range of plausible assumptions about the rate of return on saving.

6 Uncertainty over the future path of earnings deters individuals from leaving all their retirement saving to a short period of time at the end of working life. In the presence of uncertainty about earnings and employment, individuals should save more at younger ages, in particular in years when earnings are high. However, the general pattern remains – that is, many would be expected to save the minimum amount early in working life, and then increase their saving rate substantially when children leave home.

These findings demonstrate that there are good reasons why individuals should not want to save at a constant contribution rate over their entire working life.

This does not imply that the UK's current automatic enrolment policy – which strongly nudges many employees to save in a pension at all working ages, and at a single default minimum contribution rate – makes them worse off than they would otherwise be. Much of the justification for automatic enrolment is the argument that people are *not* rational and/or well-informed, and that encouraging people to save in this way therefore has a beneficial impact on retirement outcomes. In other words, automatic enrolment in its current guise may overall still be better than leaving people to have to choose to join a pension scheme.

However, going forwards, government and industry policymakers should carefully consider these life-cycle factors when debating and developing policies to increase retirement saving. Adjustments to automatic enrolment that result in non-constant saving rates over the life cycle – in particular, higher rates of saving at later ages – are likely to have advantages over an across-the-board increase in default contribution rates. Examples of such policies that should be considered include default employee contribution rates that depend on age, increases in employee contribution rates that are triggered by earnings increases, and nudges to encourage individuals to increase their pension saving when their children leave home or when they finish debt repayments such as student loans or mortgages.

These findings also illustrate an important downside of defined benefit pension arrangements: contributions cannot be easily varied over working life to suit an individual's circumstances or the timing of their capacity to save (in contrast to defined contribution arrangements, which can come with a lot of flexibility). Some employees therefore face a difficult choice about whether to join a defined benefit pension scheme with contributions at a higher level than they would ideally make at that stage in their lives, or not to join the pension and forgo the generous employer contribution. This lack of flexibility is an oft overlooked cost of public-sector pensions that are often viewed enviously for their high generosity.

1. Introduction

The need for private pension saving

Most people need to save for retirement if they want to be able to maintain their standard of living when they stop working. The new state pension currently provides an income of just over £9,300 per year to those with a full entitlement, which is around 30% of median net earnings.[1] Those on middle and higher incomes in particular are likely to be used to spending more than this, in which case they will need to supplement their state pension income with private saving in order to maintain their standard of living. This saving does not necessarily need to be in the form of a private pension – but often this will be the most financially attractive way of saving for retirement, given the tax incentives involved and the contributions often also made by employers.

The role of automatic enrolment

It is now the case that most employees earning over the equivalent of £10,000 a year are automatically enrolled into a private pension by their employer, with a default contribution rate of at least 8% of 'qualifying earnings'.[2] This policy of automatic enrolment, phased in between 2012 and 2018, has dramatically increased the proportion of employees saving in a private pension, particularly among younger individuals and other groups who were historically less likely to save in a pension.[3] Concerns remain, however, that the levels of saving being accumulated through automatic enrolment are not enough, and that many need to save more in order to have an 'adequate' standard of living in retirement.[4]

[1] See the Department for Work and Pensions publication, 'Abstract of DWP benefit rate statistics 2020', https://www.gov.uk/government/statistics/abstract-of-dwp-benefit-rate-statistics-2020.

[2] In 2021–22, 'qualifying earnings' are earnings in excess of £6,240 up to a £50,270. To be eligible for automatic enrolment, an employee must also be aged between 22 and the state pension age, and have been with their employer for at least three months.

[3] See Cribb and Emmerson (2020).

[4] See Finch and Gardiner (2017) and Pensions Policy Institute (2019a, 2019b).

How much is enough?

Exactly how much individuals need to save for retirement is a difficult question to answer. It depends on how long they will live and when they want to retire (which, combined, determine how many years of retirement they need to finance), how much they want to be able to spend each year, how much they expect to receive from others (for example, inheritance from their parents) and how much they want to give to others (either during life or as a bequest), their level of income during working life and state pension benefits in retirement, the rate of return they can get on their savings, and on how patient they are.

Assessments of the 'adequacy' of retirement saving typically take a broad-brush approach, and have tended to focus on individuals' ability to replace a certain proportion of their pre-retirement income. However, more recently, there has been a resurgence of assessing adequacy in terms of *levels* of spending that can be achieved, on the basis that these may be more salient for individuals. For example, in 2019, the Pensions and Lifetime Savings Association launched their Retirement Living Standards, which indicate the cost of different retirement life styles.[5] The Resolution Foundation has also focused on a 'Living Pension' standard that would aim to achieve a minimum acceptable standard of living.[6]

When should individuals save for retirement?

While there has been lots of discussion and research around how much people need to save for retirement, and how to increase the salience of that to bring about greater saving, much less attention has been paid to *when* individuals should save for retirement. This is an important omission from the debate. Changing incomes and spending needs over the life cycle mean that, for most individuals, saving a constant proportion of income every year of working life may not be the best strategy for their lifetime living standards. Indeed, even saving anything every year of working life may not be the 'best' thing to do.[7]

[5] https://www.retirementlivingstandards.org.uk/
[6] See Finch and Pacitti (2021),
[7] A point also made by Scott et al. (2021).

Modelling the ‘right’ timing of saving for retirement

In this briefing note, we use a simple economic model to illustrate when and how pension saving might be expected to change over the life cycle, in response to predictable factors that generally change with age. The model is only a stylised representation of reality, and not intended to capture all the nuances of real life that affect pension-saving decisions. However, even using just this simple model – the basic assumption of which is that people want to smooth their standards of living over time, and will save in order to transfer resources from ‘high’ to ‘low’ income periods – we can illustrate the effect that ‘life-cycle factors’, such as earnings growth, children and student loan repayments, would be expected to have on saving for retirement. Specifically, we quantify that most of these factors would lead one to expect saving rates to vary across the life cycle to a significant degree.

The rest of this briefing note is set out as follows. In Section 2, we briefly describe the model we use (it is described in more technical detail in the Appendix). In Section 3, we discuss in turn the implications for the profile of retirement saving of earnings growth, children, student loan repayments and employer contributions. We conclude in Section 4 with a discussion of the results and the implications, particularly for the government’s automatic enrolment policy, going forwards.

2. A simple model of retirement saving

Using an economic model of spending and saving, we illustrate quantitatively how individuals 'should' change their retirement saving over their working lives. The level of saving individuals 'should' achieve is that which best keeps their standard of living constant across the life cycle. This section provides a non-technical summary of the model, and a description of the outputs generated.

Overview of the basic model

We model the behaviour of an individual between age 22 and death. Each year, the individual chooses how much of their income to spend on purchasing goods and services, and how much to save for the future. Any saving is in the form of a safe asset, which offers a known rate of return, and cannot be accessed until retirement. When deciding how much to save, the individual's aim is to choose their spending and saving allocations such that their living standards are as smooth as possible over their life cycle. The level of saving individuals 'should' achieve, or the 'best' saving profile, is therefore that which best keeps their standard of living smooth across the life cycle.

During working life, the individual pays tax on their earnings, according to a schedule that approximates the UK tax system. They retire at the state pension age, and once they retire they no longer have any earnings, but they will receive state pension income of £9,000 per year. During retirement, they pay tax on their state pension and any withdrawals from savings (as savings are made out of pre-tax earnings during working life).

The individual faces some uncertainty, in that they do not know when they will die. We calibrate our model using Office for National Statistics 2014-based cohort survival probabilities for men born in 1985.

We start with a baseline model, in which the individual is assumed to be in paid work every year until they retire, to earn the same amount every year until retirement, and to face no uncertainty about this path of earnings. There are no other differences between any years of working life. The solution of this model – the life-cycle profile of saving generated as the 'best' outcome for the individual – provides a benchmark against which we can illustrate the effect of different features of reality on life-cycle saving patterns. Specifically, we separately introduce different features – earnings growth (Section 3.1), children and their associated costs (Section 3.2) and student loans (Section 3.3) – to illustrate the effects of these on the timing of saving for retirement. In Section 3.4, we consider the effect of all these features jointly.

We assume that the rate of return on saving is equal to 3% per year (above inflation) in the baseline model, and that the individual also 'discounts' the future by 3% per year. In other words, they are indifferent between having £1 today, or £1.03 next year. When the rate of return on saving and the individual's discount rate are the same, then the ideal will be to consume the same amount each year. However, if the real interest rate were greater than the individual's discount rate, then the individual would prefer a profile of spending that was increasing with age. (Conversely, if the real interest rate was lower than the individual's discount rate, then they would prefer a spending profile that was falling with age.) The interest rate – and how it compares to an individual's discount rate – will therefore have an important bearing on the timing of an individual's saving decisions, and how they choose to spread their spending over their life cycle. Given this, within each section, we illustrate the sensitivity of our findings to the assumed interest rate (while holding discount rates unchanged).

In the baseline model, saving is made out of pre-tax earnings, but there are no other financial incentives to save (e.g. employer or government contributions that are conditional on an individual's own level of saving). In Section 3.5, we test the sensitivity of our results to the presence of employer pension contributions that are contingent on positive employee contributions.

Finally, in Section 3.6, we illustrate the implications of uncertainty about the path of future earnings for saving behaviour.

Model outputs

The model produces an estimate for how much the individual should consume and save in each year between age 22 and death. This is illustrated in the top panel of Figure 2.1 for an individual who earns a constant amount every year. The individual's income is illustrated by the red line – this is equal to earnings each year of working life (£21,000 for the example individual in this figure), falling at the state pension age of 67 to a state pension income in each year of retirement (£9,000 here). The individual saves during working life, and draws down those savings during retirement, in such a way that spending (shown by the blue line) is smooth over most of working life. Spending declines at later ages due to mortality risk – the individual does not want to hold on to enough assets to fund spending at the same level all the way through to age 100 as they are unlikely to live that long. Instead, they enjoy slightly higher spending throughout their life, but accept that if they do survive until older ages they will have less to spend in those later years.

The bottom panel of Figure 2.1 illustrates the estimated profile of saving over the working life. In this baseline model, with constant earnings over the life cycle and no other differences between years of working life, the individual saves a constant proportion of their earnings every year until retirement (around 8% of gross earnings in this case). After retirement, the savings rate turns sharply negative as individuals draw on their accumulated wealth to top up their state pension income.

Figure 2.1. Baseline model output

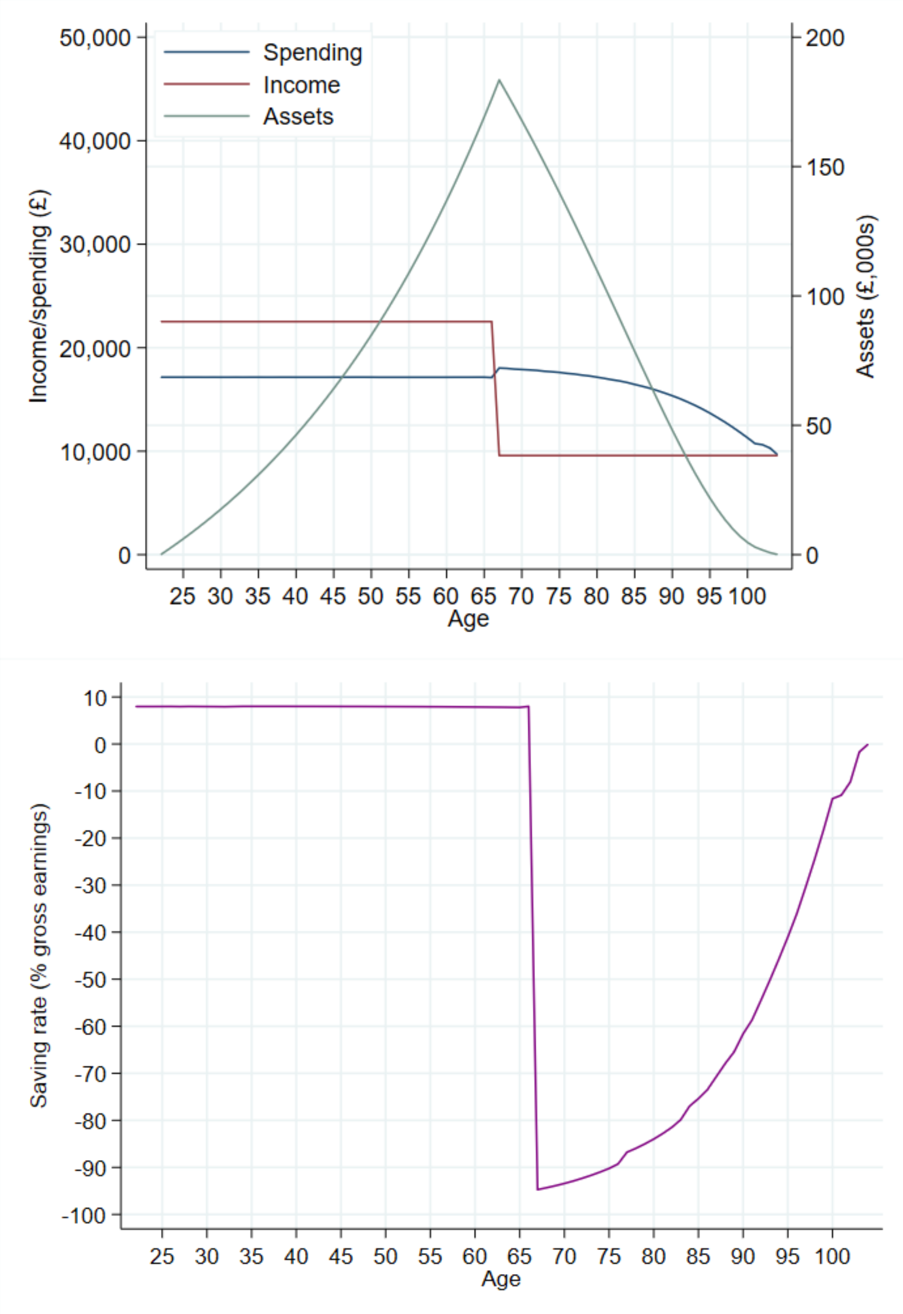

3. Factors affecting the timing of saving

3.1. Earnings growth

Earnings normally grow over an individual's working life. As people get more experience and are more productive, they can typically command higher wages. The rate of earnings growth varies across individuals, depending on their rate of skill acquisition and their occupation. On average, those with higher levels of education enjoy faster earnings growth than those with lower levels of education, and men on average experience greater wage growth than women as the gender wage gap gradually increases after the birth of a first child.[8]

Effects of earnings growth on the timing of saving

We illustrate the effects of earnings growth on the 'best' profile of saving over working life by running our model with three different illustrative lifetime earnings profiles. These are chosen to approximate the average age–earnings profiles (over and above economy-wide earnings growth, and measured across both men and women) for those with GCSEs only, A levels only, and degree or higher qualifications.[9] The three earnings profiles are shown by the three solid lines in Figure 3.1. The dashed line shows the level of constant earnings from the baseline model, which is set at a level such that it yields the same gross resources over the life cycle as the earnings profile for the mid-educated individual. It should be noted that earnings increase more rapidly with age for those with higher education, but even the lowest education group, on average, experiences some earnings growth. At older ages, the illustrative profiles have declining earnings, as a result of the fact that, empirically, those in work at older ages have declining average earnings with age (in part due to a shift towards part-time working).

[8] See, for example, Costa Dias, Joyce and Parodi (2018).
[9] The estimation of these earning profiles is described in the Appendix.

Figure 3.1. Illustrative earnings profiles

Note: Earnings profiles estimated using data from the Family Expenditure Survey and successor surveys. Details are provided in the Appendix.

The top panel of Figure 3.2 illustrates the spending profile that the model simulates that a mid-educated individual would choose, given their assumed earnings profile and all the other assumptions of the model (as set out in Section 2). Spending is broadly flat throughout most of life; by assumption, the individual adjusts their saving to smooth their standard of living over their life cycle. However, early in working life, individuals are constrained by their income level – they would like to be able to spend more given their lifetime earnings, but there is no borrowing allowed in the model and therefore they cannot bring forward resources from the future. Instead, individuals spend everything they can (their full net income) at younger ages and spending increases with age as income increases. For the mid-educated individual illustrated here, this happens until their late-30s.

For brevity, the figures for the spending patterns of low- and high-educated individuals are not shown, but the patterns are similar. The level of spending is higher at any given age for those with greater education, as they have higher incomes.

Figure 3.2. Impact of earnings growth on life-cycle profile of saving

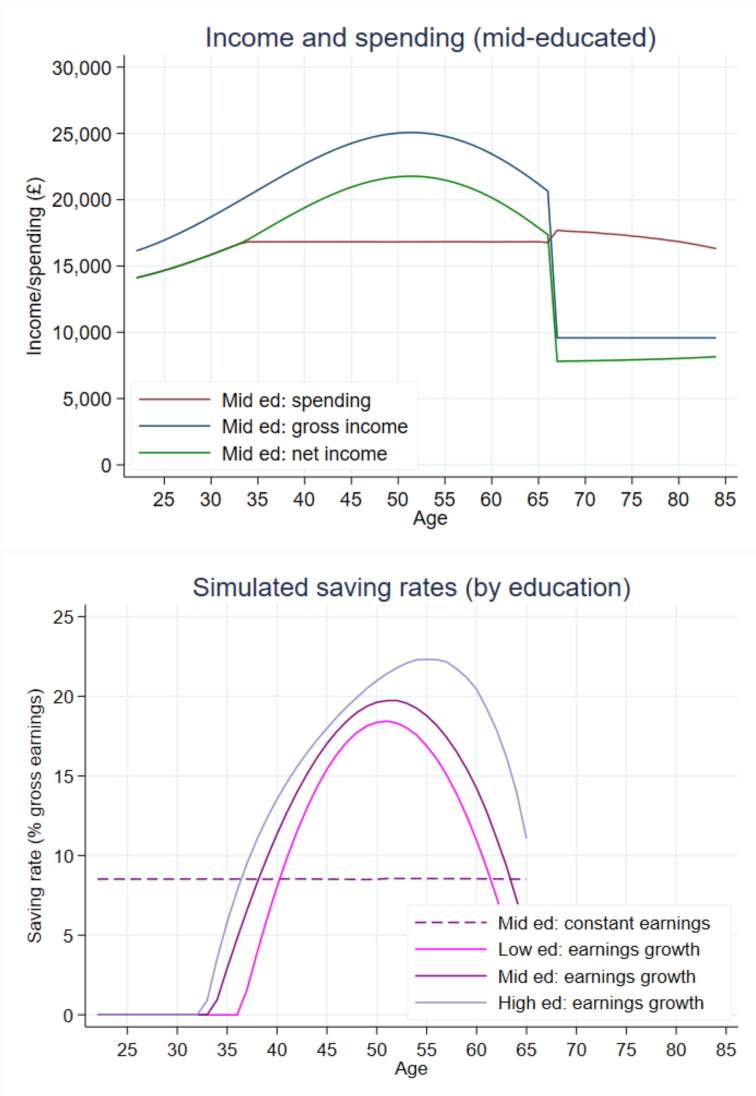

The resulting simulated saving rates in each year of working life for an individual with these alternative earnings profiles are shown in the bottom panel of Figure 3.2. It is striking that, compared with the baseline model (illustrated by the dashed line), saving is no longer uniform over working life. Instead, saving varies dramatically over the life cycle; it is low at younger ages when earnings are lower, and increases substantially at older ages, to what might be considered very high saving rates, when earnings are higher. In fact, saving does not even occur at all ages – in these simulations, low- and mid-educated individuals would only start saving in their late-30s, and a high-educated individual only in their early-30s. And – at the other extreme – at age 50, saving rates are in the region of 20%–25% of gross income.

This profile of saving rates over working life, combined with the profile of annual earnings, implies that a very large proportion of retirement saving is done during the last half of working life. For example, for a mid-educated individual, we simulate that 83% of retirement saving is done between ages 45 and 66, with contributions made between those years contributing 78% of retirement wealth at age 67 when the return to saving is factored in. This compares to around 47% of retirement saving (contributing 31% of retirement wealth) being done during the last half of working life in our baseline model with constant earnings and a simulated constant saving rate.

Sensitivity to the rate of return

The exact profile simulated, and the age at which individuals start saving, is somewhat sensitive to assumptions made about the rate of return on assets and how patient the individual is. In Figure 3.3, we illustrate how the simulated saving profile for low- and high-educated individuals would differ if the real rate of return on savings was 0%, 5% or 7%, (compared to the baseline return of 3%). We leave the individual's degree of impatience unchanged. This shows that with higher returns on assets the individual would choose to start saving earlier, and the peak saving rate at older ages would be lower. However, the general pattern – that an individual who experiences earnings growth through their working life should save less at younger ages and dramatically more at older ages – is robust to these different assumptions.

The proportion of retirement saving done in the last half of working life, and the proportion of retirement wealth arising from those contributions, are set out in Table 3.1 for the different rates of return. As illustrated in Figure 3.3, when returns are higher, more saving is done earlier in working life; a mid-educated individual makes 57% of their lifetime contributions in the first half of working life when returns are 5%, and 24% when returns are 7%, compared to 78% when returns are 3%. Furthermore, the higher the returns, the greater the proportion of retirement wealth that will accumulate from returns made earlier in working life. For example, with returns of 5%, almost half (43%) of retirement wealth for the mid-educated individual arises from contributions made in the first half of working life.

Figure 3.3. Sensitivity of impact of earnings growth to rate of return

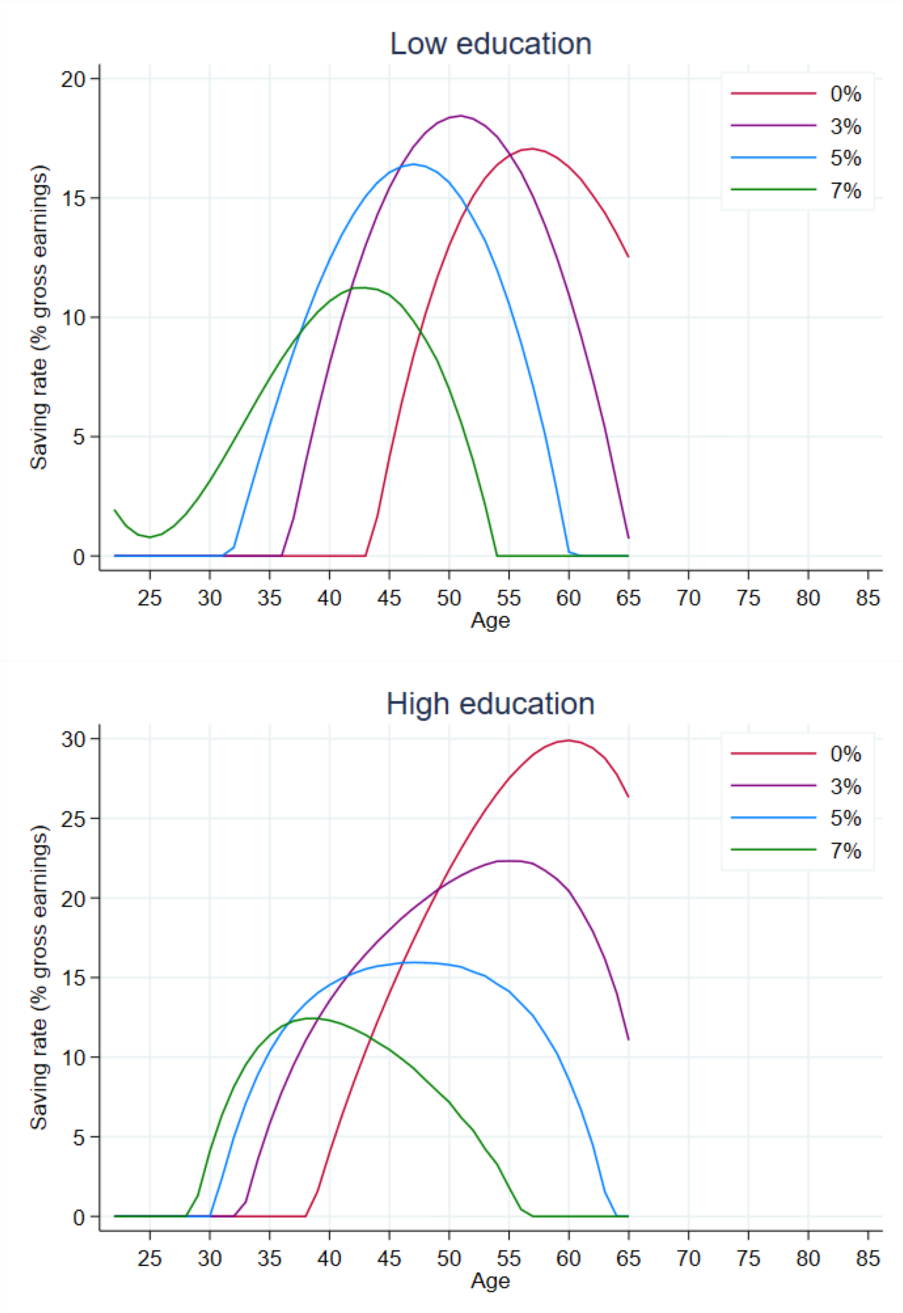

Note: All other parameters and assumptions (aside from the rate of return) are as in the baseline model.

Table 3.1. Relative importance of saving in last half of working life – impact of earnings growth

	Percentage of lifetime contributions made between ages 45 and 66			Percentage of retirement wealth arising from contributions made between ages 45 and 66		
Education:	Low	Mid	High	Low	Mid	High
'Baseline'	49	49	49	33	33	33
With earnings growth	82	78	78	76	71	70
With earnings growth (alternative rates of return):						
0% return	99	97	93	99	97	93
5% return	63	57	62	51	43	47
7% return	37	24	35	22	11	20

Note: 'Baseline' assumes 3% real return. 'Retirement wealth arising from contributions' factors in the real return earned between when the contributions are made and age 67.

3.2. Children

Our basic model assumes that an individual acts to smooth their standard of living over their life cycle, where their standard of living is determined by their level of spending each year. However, individuals are often members of a household, and it is argued that different household types typically have different financial resource requirements to achieve a comparable standard of living. In particular, couples need less than twice the resources of single individuals, as some resources can be shared, while those with children need more resources.

Children have an important effect on the profile of spending and saving decisions as they typically do not place the same financial burden on their parents over the whole of their parents' working lives. One way of illustrating this would be to model the expenses caused by children at different points in their upbringing. For example, childcare expenses for pre-school age children are a significant expense for working parents.[10] However, here we take a more general approach, and 'equivalise' the spending of the modelled individual. Equivalisation is a standard methodology that adjusts household income to account for the different financial resource requirements of different households.[11] Specifically, we assume that for each child aged under 12, and for each child aged 13–17, the modelled parent must have spending 10% and 16⅔% higher (respectively) than the level of spending when there are no children present in the household, in order to achieve the same standard of living. This is half the standard equivalisation amounts, as we assume half the financial cost of any children is borne by the other parent (or another person).

Effects of children on the timing of saving

To illustrate the effects of children on the timing of retirement saving, we adjust our baseline model (which has constant working-life earnings) to give the individual two children, the first born at age 30 and the second at age 32. We assume, for simplicity, that there are no implications of children on the path of earnings. The top panel of Figure 3.4 shows the simulated spending profile that best smooths the individual's standard of living over the life cycle. This is no longer constant across working life as it was in the baseline model. Spending is higher between the ages of 30 and 50 when there are children in the household. Equivalised spending is constant between ages 29 and 31 because spending is increased only to compensate for the presence of the first child. With the arrival of the second child, equivalised spending falls, as spending can only be increased up to the level of net income; the parent would like to borrow, but this is not allowed in the model. Similarly, later in working life when the children are assumed to be more expensive, equivalised spending drops further as the individual is still constrained by their income.[12]

[10] The average price of 25 hours of childcare a week for a child under two in a nursery in 2019 was £131.61 across Great Britain, or £6,800 a year (Coleman, Dali-Chaouch and Harding, 2020).

[11] In line with official income statistics in the UK, we use the modified OECD equivalence scale. See Anyaegbu (2010) for a discussion.

[12] Net income is not constant over working life because, in the model, pension contributions are tax deductible. Net income will therefore be higher when the pension saving rate is higher.

Figure 3.4. Effects of children on life-cycle profile of saving

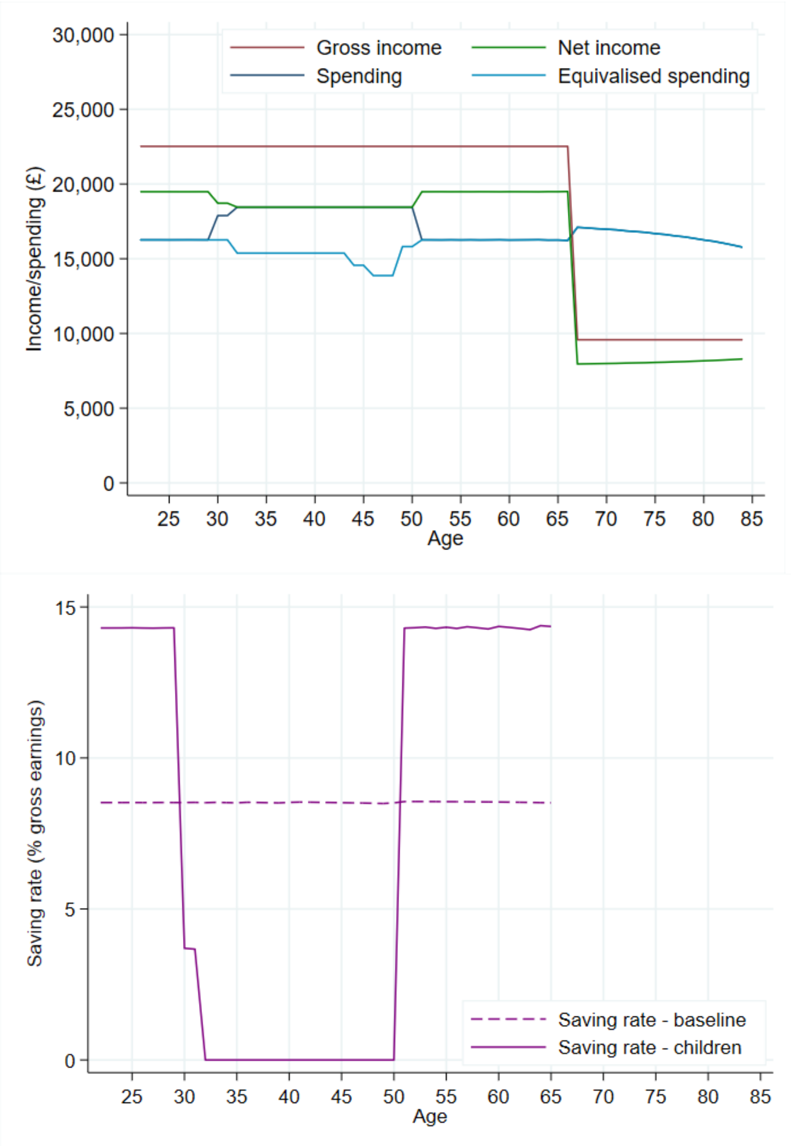

Note: Baseline model with and without children.

The effect on the profile of saving for retirement is illustrated in the bottom panel of Figure 3.4. Rather than being constant over working life, saving for retirement in the presence of children is compressed into the years when children are not present in the household – before they arrive and after they are assumed to leave home. In other words, rather than saving 8% of gross earnings each year, the individual is simulated to save around 13% before the first child is born, nothing while there are two children in the household, and around 13% again once both children have left the home.

For the baseline individual with two children, the proportion of retirement saving done in the last half of working life and the proportion of retirement wealth arising from those contributions are simulated to be 66% and 42%, respectively (compared to 47% and 31% for the baseline individual with no children).

Sensitivity to the rate of return

The simulated saving profile, and whether saving is concentrated in the years before children arrive or the years after they leave, is again sensitive to the assumed rate of return on assets and how patient the individual is. This is illustrated in Figure 3.5. However, in all cases illustrated here, it remains the case that the individual would not want to save for retirement when they have two children in the house. What changes is that when the rate of return is high, individuals should do most or all of their saving in the years before children arrive, whereas when rates of return are low, they should do all of their saving after their children have left home.

Sensitivity to the availability of other assets

The result that individuals should save for retirement before the arrival of children is sensitive to the simplification in the model that there is only one saving asset available to people – the illiquid retirement asset. The modelled parent is spending all their income in each year when they have two children; ideally they would like to spend more in those years and less in other years, but they cannot save in a form that they can then access again before retirement nor can they borrow against future income. If the model allowed saving in a liquid asset then individuals would choose to save in that before the arrival of children, and then spend down those assets in years in which they have two children at home. This would crowd out some or all retirement saving early in working life (depending on relative rates of return), pushing some or all saving for retirement until after children had left home.

Figure 3.5. Sensitivity of the effects of children to assumed rate of return

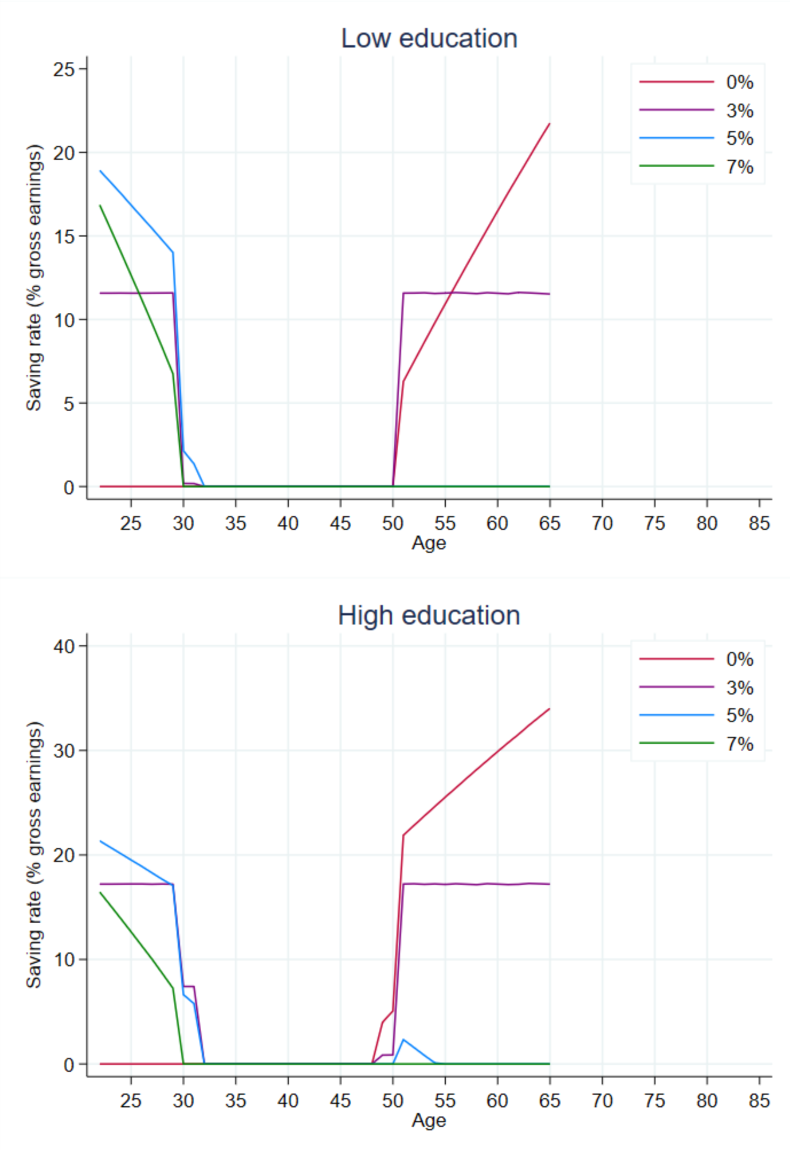

Note: All other parameters and assumptions (aside from the rate of return) are as in the baseline model.

Summary

Children are an important life-cycle factor, and they can be sufficiently expensive to cause individuals to want to save for retirement in years when they do not have children at home. However, the simulated best timing for retirement saving depends on the rate of return, and on how expensive children are, and when they are expensive. For example, our simulations are based on the assumption that children impose costs on their parents only until age 18. Should children remain expensive for the duration of parents' working life, then parents would not be able to postpone retirement saving until after those costs have ended.

3.3. Borrowing and student loans

When earnings grow over working life, or if spending needs change throughout life (for example, because of children), then individuals may find themselves wanting to transfer resources within working life between periods with higher income relative to needs to periods with lower income relative to needs. This was the case in the simulations run in Sections 3.1 and 3.2, when at different stages of life the simulated individual was sometimes spending all of their net income but still having lower spending (or equivalised spending) than during other periods of working life.

It is computationally too challenging here to model both liquid saving or borrowing at the same time as illiquid retirement saving. However, in general it is likely to be the case that individuals might want to borrow when their resources are low relative to their average lifetime income; and if the interest payable on that borrowing is greater than the return on retirement saving, then they would be better off repaying that borrowing before saving for retirement.

One form of borrowing that it is possible to simulate the effects of, because of the fixed timing of the borrowing and known nature of the repayment schedule, is student loans accumulated as a result of studying for a degree. For example, those entering higher education in England and Wales since 2012 have faced university fees in excess of £9,000 per year. The vast majority take out government-backed subsidised loans to cover these fees (and maintenance costs), which are paid back from income after graduation. Under the terms of these loans, once income exceeds a certain threshold (currently £26,575), repayments of 9% of income above that threshold must be made. Then, 30 years after graduation, any remaining debt is written off. Given the size of the debt that people accumulate over the course of their degree, the interest rates charged on that debt and the repayment rules, only a minority of graduates should expect to repay their student loans before the end of the 30-year repayment term.[13] In other words, the majority of recent and new

[13] See Belfield, Britton and van der Erve (2017).

graduates in effect face an additional tax of 9% of their income in excess of £26,575 until age 52.

Effects of student loan repayments on the timing of saving

We model the effects of a ‘student loan’ tax (9% of income above the threshold until age 50) for our illustrative ‘high-educated’ individual who has the earnings growth and earnings profile described in Section 3.1. We assume that student loan repayments are calculated on the basis of gross income, rather than gross income after pension contributions (in effect, we assume that the individual does not make pension contributions through salary sacrifice). The results are shown in Figure 3.6. From the profiles for net income with and without student loans in the top panel, it is clear that, given the nature of the repayment schedule, student loan repayments can be relatively small. For our illustrative individual who earns around £34,000 at the end of their 40s, annual repayments at that point would amount to less than £700 per year, or 2% of gross earnings.

The bottom panel shows that student loan repayments do affect the desired saving profile. If the individual has to make student loan repayments, then they would save slightly less in their late-40s, and increase their saving more sharply at age 51, compared with their saving rate if they did not have to make repayments. Saving would increase by around 2% of gross earnings at age 51; in other words, all the income that was previously used to make loan repayments would be transferred to pension saving instead (spending is smooth between those ages, as shown in the top panel of Figure 3.6). However, in terms of size, this increase in the saving rate at the point at which the student loan repayments end is dwarfed by the increase in saving that would be expected to occur in the last half of working life anyway as a result of earnings growth.

Figure 3.6. Effects of student loan repayments (SLR) on life-cycle profile of saving

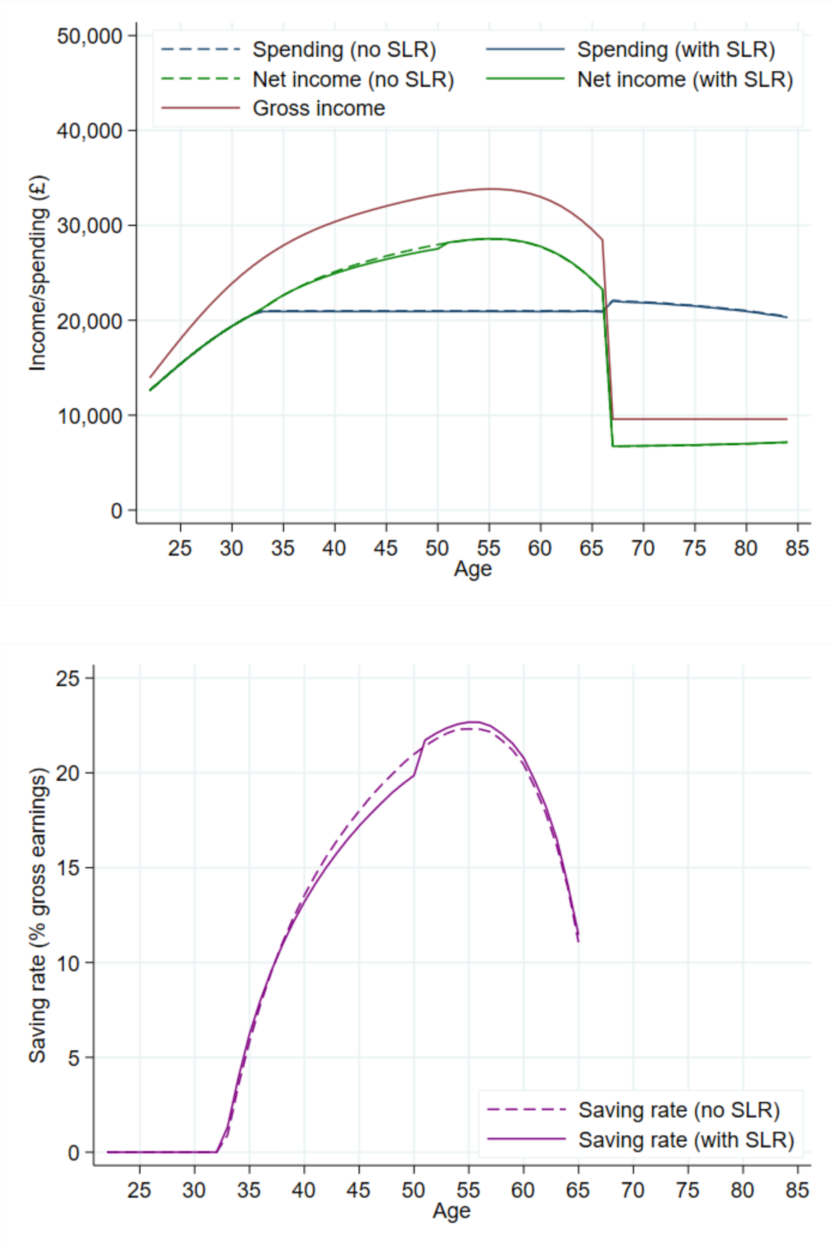

Note: Simulations are for a high-educated individual with earnings growth as in Section 3.1.

If the individual made pension contributions through salary sacrifice, then the simulated saving profile would differ. If the individual does not expect to repay their entire loan before the end of the repayment term, then they have an additional incentive to save in a pension (through salary sacrifice) as this reduces their loan repayments – analogous to reducing their total lifetime tax liability. In fact, our modelled high-educated individual would never make student loan repayments, as their modelled pension contributions would reduce their gross earnings each year to less than the repayment threshold. In this case, there would be no step change in the saving rate when the student loan is written off, as there were never any repayments being made that would cease at this point.

3.4. Composite scenarios

The results of the simulations reported in the previous three sections clearly illustrate that there are good reasons to expect saving for retirement to vary over the life cycle. Earnings growth should lead people to save a greater percentage of their earnings in years in which their earnings are higher, those with children would be expected to save before their children arrive and/or after their children have left home, and those repaying student loans would be expected to have higher rates of saving in years after their loan repayments have ceased.

Here we take all these circumstances together, and simulate the savings rate each year for a low-educated individual with two children born at ages 30 and 32, a mid-educated individual with two children born at ages 30 and 32, and a high-educated individual with two children born at ages 30 and 32 and a student loan. The resulting profiles are illustrated in Figure 3.7 for our baseline rate of return. In all cases, the saving profile is very similar, with virtually no saving for retirement until the individual is in their 50s and then a very high saving rate thereafter. The combination of the costs of children and earnings growth act together to push saving until later in life after children have left home.

Figure 3.7. Joint impact of earnings growth, children and student loan repayments (where applicable) on life-cycle profile of saving

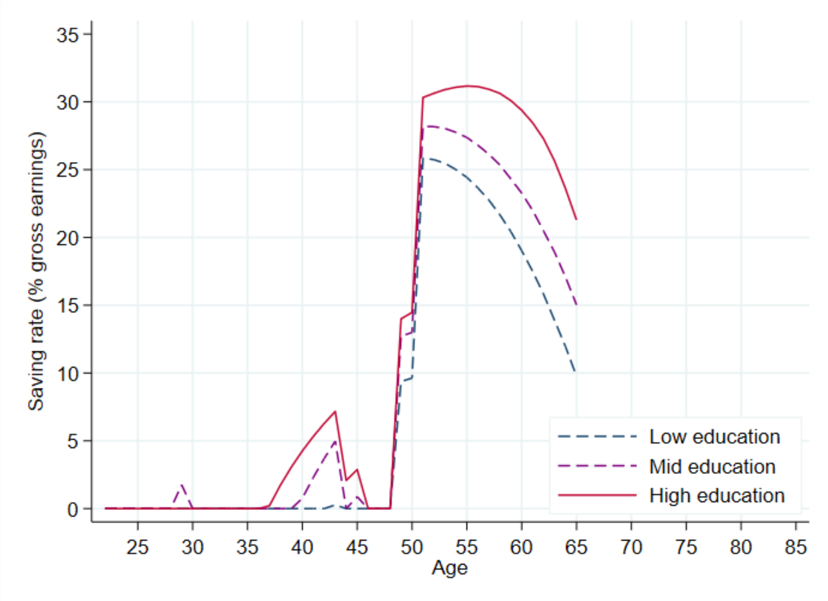

Note: Simulations assume two children, born at ages 30 and 32, as in Section 3.2.

Sensitivity to the assumed rate of return

The sensitivity of the saving rate profile from this 'composite scenario' to different assumed rates of return is illustrated in Figure 3.8 for the low-educated individual (top panel) and the high-educated individual (bottom panel). When the rate of return is higher, more saving would be done earlier in working life. For the high-educated individual, this may even be during years when there are two children living at home. However, as previously discussed, the extent to which retirement saving is brought into the early years of working life will be sensitive to the assumption in our model that no other, more liquid, form of saving is available, and to the relative return on different forms of saving.

Figure 3.8. Sensitivity of the joint impact to assumed rate of return

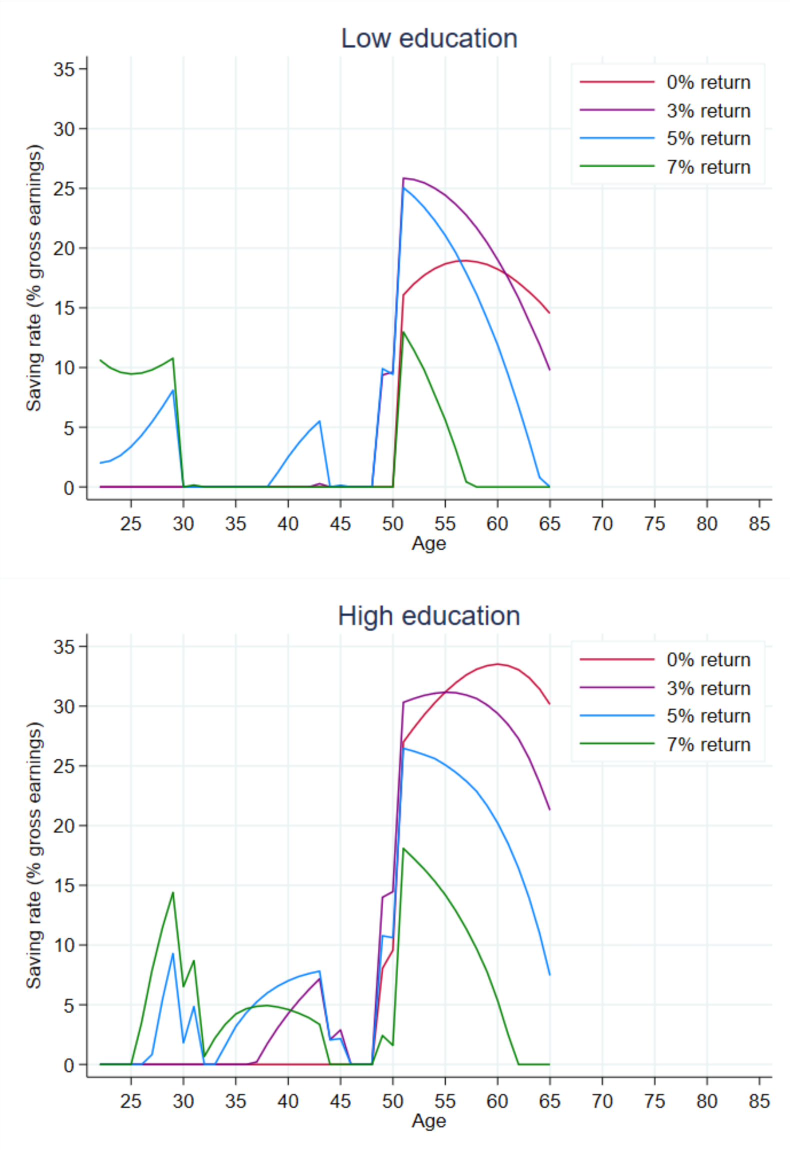

Note: All other parameters and assumptions (aside from the rate of return) are as in the baseline model.

The proportion of retirement saving done in the last half of working life, and the proportion of retirement wealth arising from those contributions, is shown for the different assumed rates of return in Table 3.2. With our baseline 3% rate of return, virtually all of retirement wealth is accumulated in the last half of working life. At higher rates of return, this proportion falls, both because the individual makes more contributions earlier in life, and because those contributions accrue a greater return before retirement.

Table 3.2. Relative importance of saving in last half of working life – joint impact of earnings growth, children and student loan repayments

	Percentage of lifetime contributions made between ages 45 and 66			Percentage of retirement wealth arising from contributions made between ages 45 and 66		
Education:	Low	Mid	High	Low	Mid	High
'Baseline'	49	49	49	33	33	33
With EG, C, SLR	100	97	94	100	95	91
With EG, C, SLR (alternative rates of return):						
0% return	100	100	100	100	100	100
5% return	86	77	83	66	49	66
7% return	50	31	64	14	7	30

Note: Baseline assumes 3% real return. EG = earnings growth, C = children and SLR = student loan repayments (only applicable for high education), "Retirement wealth arising from contributions" factors in the real return earned between when the contributions are made and age 67.

3.5. Contingent employer contributions

One factor not captured in our baseline model, or the permutations in Sections 3.1–3.4, is the possible existence of employer pension contributions that are contingent on an individual making employee contributions. These would give individuals an additional financial incentive to save for retirement.

Contingent employer contributions could take many forms; for example, an employer could match an employee's contributions at a certain rate up to a particular level. Here we simulate the effects of contingent employer contributions that are similar to the automatic enrolment default minimums (where the employer offer is as ungenerous as is allowed). Specifically, we assume that if an individual contributes less than 5% of their gross pay, then they do not receive any employer contribution, but if they contribute 5% or more of gross pay, then they receive a top up to their saving of 3% of gross pay. (For simplicity, we assume that these contributions are a proportion of total earnings, rather than a proportion of ‘qualifying earnings’.)

The effects of this incentive to save in a pension are illustrated in Figure 3.9 for a low-educated individual with two children, and for a high-educated individual with two children and a student loan. The ‘no contingent contributions’ (‘No C.C.’) saving profile is the same as illustrated in Figure 3.7. The effect of this financial incentive to save for retirement is clear – rather than saving nothing until the last half of working life, the individual now chooses to save in all years of working life. However, they choose to save the minimum required to get the employer contribution until their late 40s, and then they increase their saving rate substantially. This same profile holds for different assumptions about the real rate of return (shown in Figure 3.10 for the profile of employee contributions only). With higher rates of return, the peak saving rate is lower, but the same pattern – of the minimum contribution being made until children leave home and then the saving rate increasing substantially – holds throughout.

The presence of contingent employer contributions reduces the proportion of retirement saving that is done in the last half of working life. With our baseline return of 3%, the mid-educated individual is simulated to make 77% of their lifetime contributions between the ages of 45 and 66 when there are contingent employer contributions, compared to 97% when there are not. Given the return to saving, this means that only 64% of retirement wealth arises from contributions made in the last half of working life, compared to 95% (shown in Table 3.3). At higher rates of return, the proportion of retirement wealth accumulated in the last half of working life is even lower.

Figure 3.9. Implications of employer incentive on simulated saving profile

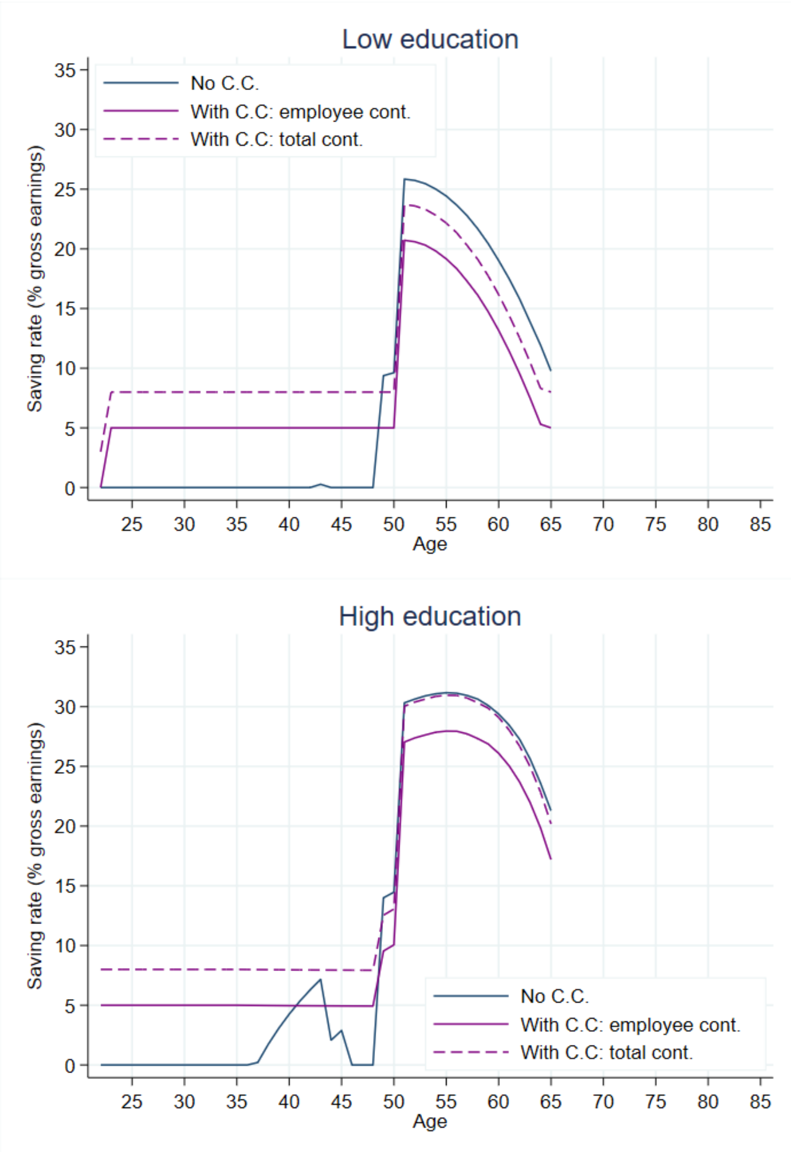

Note: 'No C.C' is the scenario with no contingent employer contributions, while 'With C.C.' is the version of the model including the contingent employer contributions. The savings rate is expressed as a percentage of gross earnings excluding employer contributions.

Figure 3.10. Implications of employer incentive on simulated saving profile

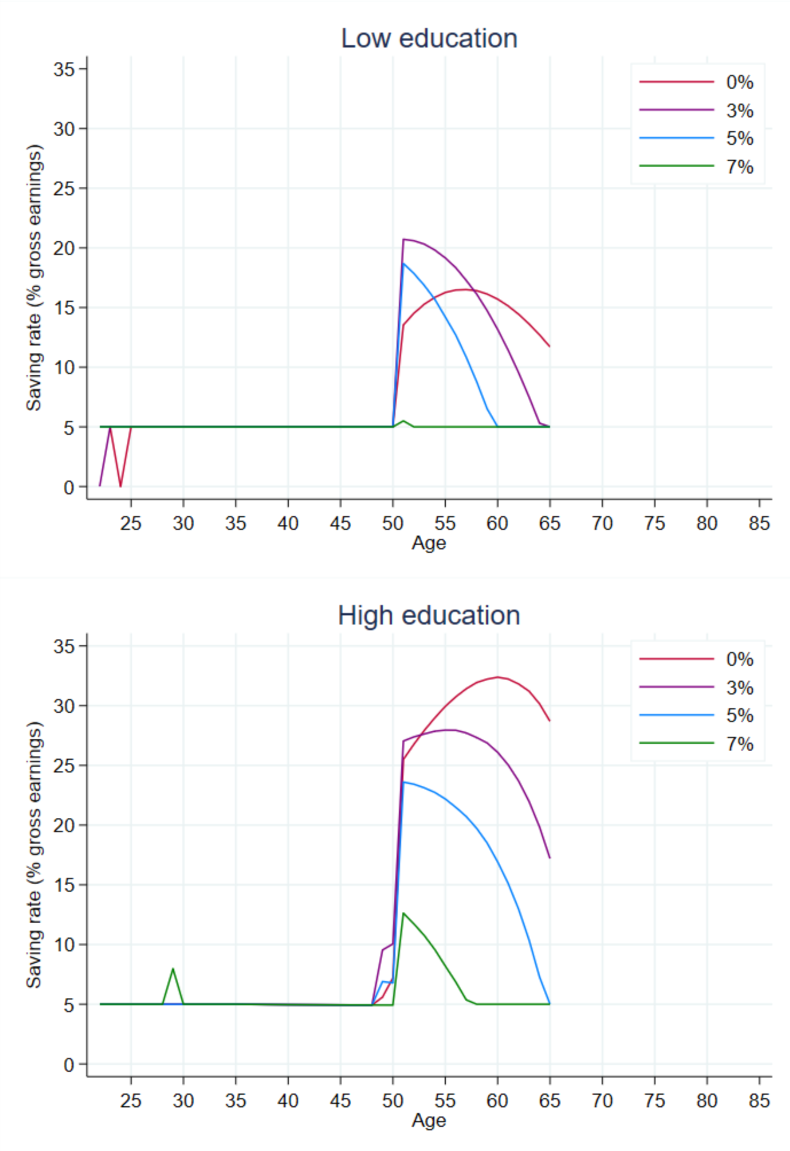

Note: Saving rate excludes the employer contributions, and is expressed as a percentage of gross earnings excluding employer contributions.

Table 3.3. Relative importance of saving in last half of working life – joint impact of earnings growth, children, student loan repayments and contingent employer contributions

	Percentage of lifetime contributions made between ages 45 and 66			Percentage of retirement wealth arising from contributions made between ages 45 and 66		
Education:	Low	Mid	High	Low	Mid	High
'Baseline'	49	49	49	33	33	33
With EG, C, SLR, CEC	74	77	83	61	64	72
With EG, C, SLR, CEC (alternative rates of return):						
0% return	75	79	85	75	79	85
5% return	67	68	78	42	43	55
7% return	54	52	61	22	20	27

Note: Baseline assumes 3% real return. EG = earnings growth, C = children, SLR = student loan repayments (only applicable for high education) and CEC = contingent employer contributions. 'Retirement wealth arising from contributions' factors in the real return earned between when the contributions are made and age 67.

3.6. Earnings and employment uncertainty

In our baseline model, and in the permutations in Sections 3.1–3.5, the modelled individual faces no uncertainty aside from their longevity. They are assumed to be in employment every year, they know exactly how much they will earn each year in future, and those earnings are either constant with age or follow a relatively smooth path.

In reality, individuals do not enjoy such certainty about their future earnings or employment. The profile of lifetime earnings is gradually revealed over time, with expectations about future earnings being revised upwards or downwards as positive and negative shocks come along. This uncertainty can be thought of as having two effects. First, the presence of uncertainty in general will affect saving behaviour. For example, a risk-averse individual will not want to plan to do most of their retirement saving in a short age range, even if they expect earnings to be particularly high then, in case they are hit by a negative shock, which would mean that the period of high earnings does not happen. Second, in any given year, if there is a positive (or negative) earnings shock, then the individual will save relatively more (or less) in that year as that will best smooth their living standards over time.

Here we illustrate the effects on the simulated saving profiles of incorporating earnings uncertainty in our model. Rather than assuming future employment and earnings are certain, we now assume that each year the individual faces a possibility of moving up or down the earnings distribution, including into or out of work. The probabilities of doing so (as well as the level of earnings at different points of the distribution), depend on education, age and the individual's current position in the earnings distribution.[14] The individual is assumed to know these probabilities. We then simulate the behaviour of 1,000 individuals (of each education group), where the saving behaviour of each individual will be different because of their different earnings realisations, and therefore also expected future earnings, at each age.

In Figure 3.11, we illustrate the simulated proportion of individuals saving at each age (top panel) and the median saving rate among those saving (bottom panel) when we introduce earnings and employment uncertainty into our model incorporating earnings growth, two children and student loan repayments (where applicable). In other words, this is analogous to the model run in Section 3.4, but with the addition of earnings and employment uncertainty. One key difference from the simulated saving profiles in Section 3.4 is that here some individuals save at all ages. Around 25%–45% of those in their mid-20s to late-40s save in any given year. The proportion of individuals saving increases sharply when children leave home, but this is not 100% even among those in their 50s. This is in part because those experiencing positive earnings shocks are likely to save, while those experiencing negative shocks are less likely to save, and there are individuals experiencing positive and negative shocks at all ages. Figure 3.12 shows that those saving in a pension, on average, have higher earnings realisations than those not saving in a pension.

[14] These probabilities, and how we estimate them, are described in the Appendix.

Figure 3.11. Proportion of individuals saving and the average saving rate in the presence of earnings uncertainty – no conditional employer contributions

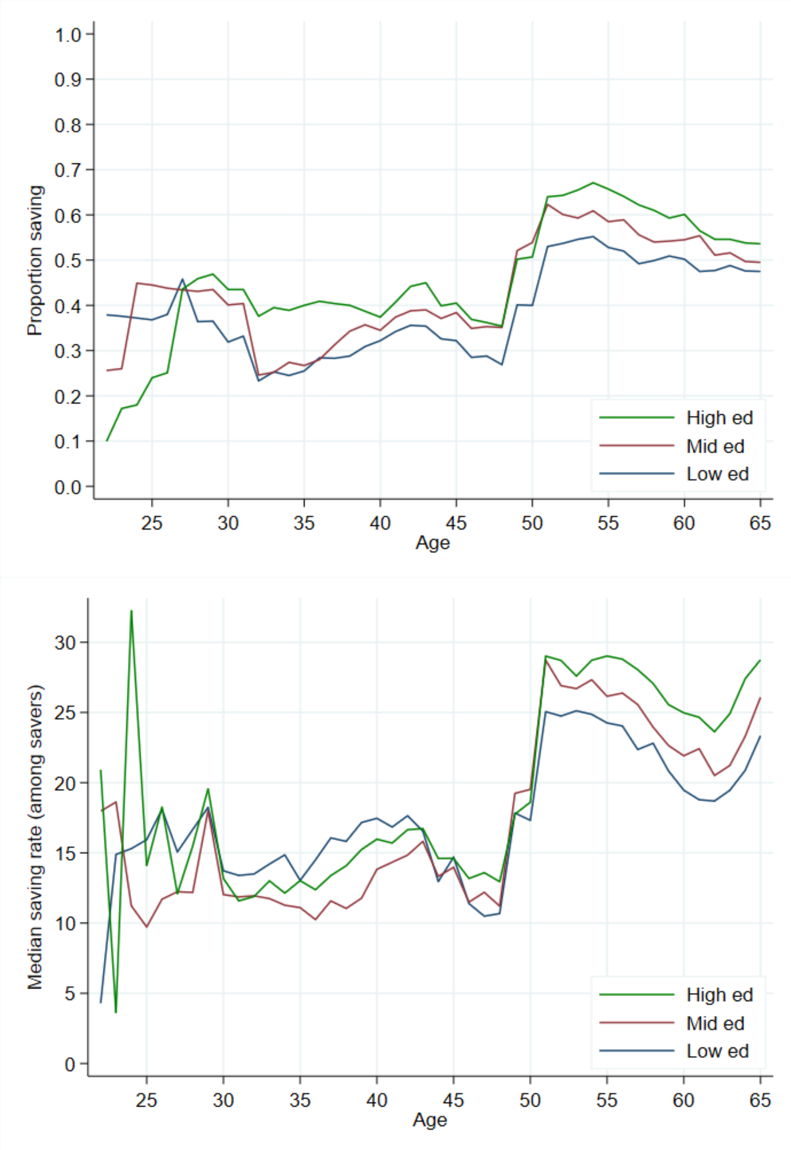

Note: Model includes uncertain earnings and employment, children and student loan repayments.

Figure 3.12. Average earnings among those saving and those not saving in a pension

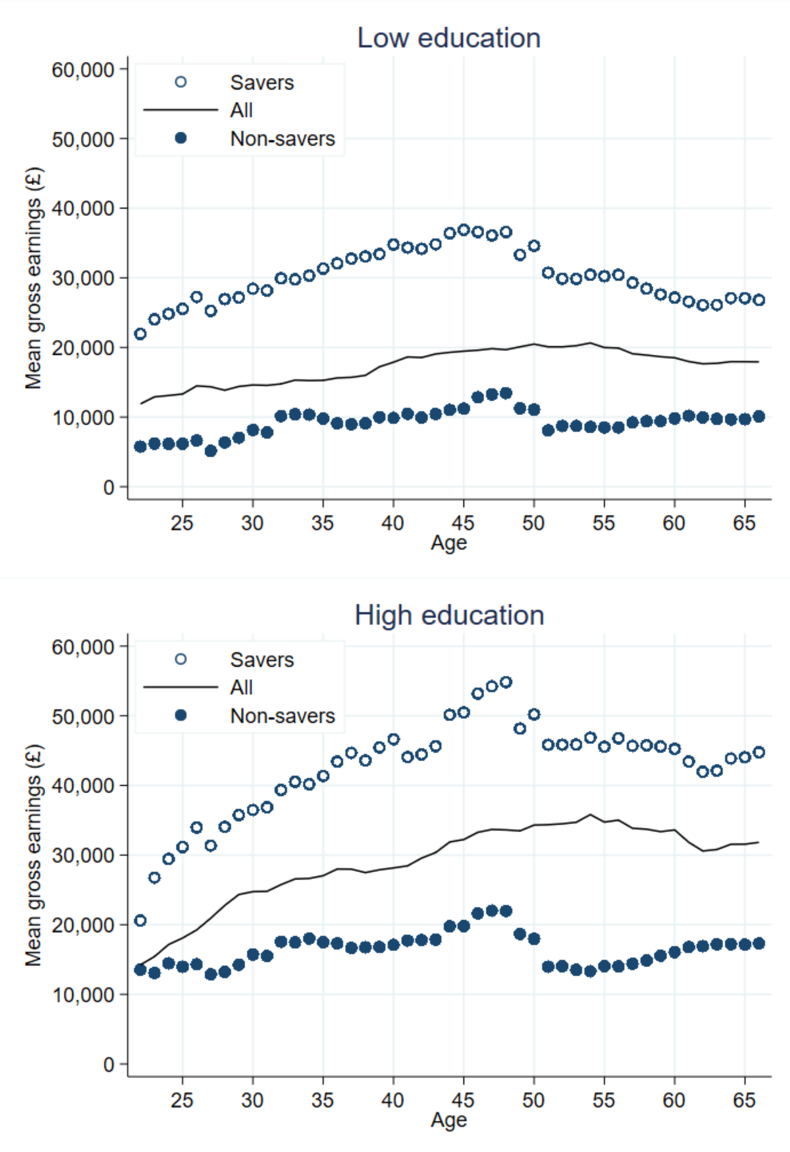

In many respects, however, the simulated savings profiles are similar. The proportion of people saving for retirement is not constant over working life, and increases at later ages due to the influences of earnings growth, children and student loan repayments. Furthermore, among savers, the savings rate is not constant over working life, but increases sharply around age 50, as it does in our simulations without uncertainty in earnings and employment.

Table 3.4 describes how the proportion of saving done in the last half of working life, and the proportion of retirement wealth arising from contributions made during that period, changes when uncertainty in earnings and employment is introduced. The simulated results when earnings growth, children and student loan repayments are factored into the model show that virtually all retirement saving would happen in the last half of working life. When earnings and employment uncertainty is introduced, as previously described, more saving happens earlier in working life. Half of mid-educated individuals do three-quarters or more of their retirement saving in the last half of life (contributing 62% or more of their retirement wealth). However, a quarter of individuals make less than 46% of their lifetime contributions (resulting in 30% or less of retirement wealth) in the last half of working life.

Table 3.4. Relative importance of saving in last half of working life – with and without earnings and employment uncertainty

Model: includes EG, C, SLR	Percentage of lifetime contributions made between ages 45 and 66			Percentage of retirement wealth arising from contributions made between ages 45 and 66		
Education:	Low	Mid	High	Low	Mid	High
No uncertainty	100	97	94	100	95	91
With earnings and employment uncertainty:						
25th percentile	31	46	48	18	30	32
Median	66	76	74	49	62	61
Mean	59	65	65	49	56	56
75th percentile	88	90	89	79	82	80

Note: Model incorporates earnings growth (EG), children (C) and student loan repayments (only applicable for high education, SLR). Rate of return is 3%. 'Retirement wealth arising from contributions' factors in the real return between when contributions are made and age 67.

Sensitivity to the availability of other assets

These simulated saving profiles are likely to be somewhat sensitive to the simplification in the model that there is only one saving asset available to individuals – the illiquid retirement asset. If individuals also had the capacity to save in a liquid asset, then they would likely want to save in that form in order to smooth their living standards against earnings and employment shocks during working life. In particular, they would likely save in that asset during years with a positive earnings shock, and save less in the pension asset during those years as result.

With contingent employer contributions

In Figure 3.13, we illustrate the simulated proportion of individuals saving at each age (top panel) and the median saving rate among those saving (bottom panel) when we also include contingent employer contributions in our model. In other words, this is analogous to the model run in Section 3.5, but with the addition of earnings and employment uncertainty. The proportion of individuals saving each year is much higher than in the absence of contingent employer contributions, but not 100% due to the possibility of unemployment (which is higher among the low-educated group). Among those saving, the profile of the median saving rate looks similar to that simulated in the absence of uncertainty; that is, the median saver contributes the minimum required to get the employer match until age 50, and then the average contribution rate increases sharply at age 50.[15]

[15] The mean saving rate among savers is higher than the median (and higher than the minimum required to get the employer match), as those individuals experiencing positive earnings shocks save at a higher rate. However, the mean saving rate still increases noticeably around age 50.

Figure 3.13. Proportion of individuals saving and the average saving rate in the presence of earnings uncertainty – with contingent employer contributions

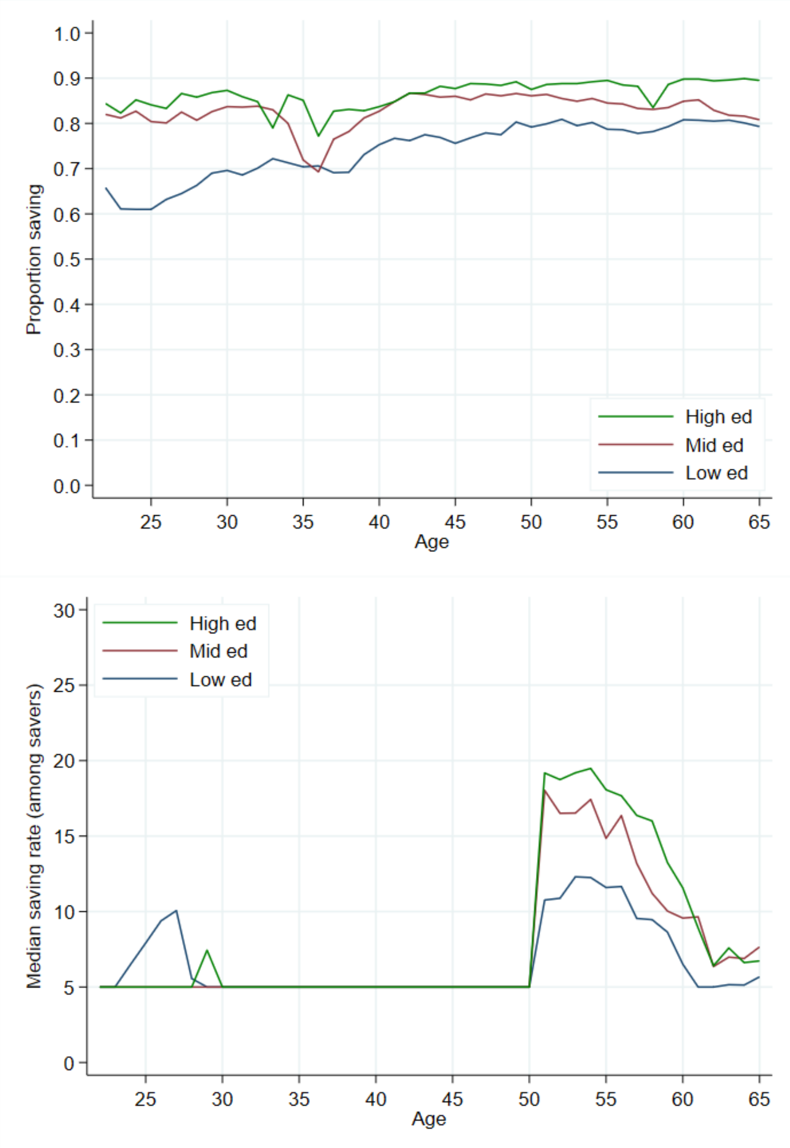

Note: Model includes earnings and employment uncertainty, children, student loan repayments and contingent employer contributions.

Table 3.5 describes the proportion of saving done in the last half of working life, and the proportion of retirement wealth arising from contributions made during that period, when both earnings and employment uncertainty and contingent employer contributions are included in the model. In the absence of uncertainty, the mid-educated individual did 77% of their retirement saving (contributing 64% of their retirement wealth) in the last half of working life. With uncertainty, the proportion of saving done in the last half of life is only slightly lower for the median individual, at 63% (46%). However, some individuals would do a greater proportion of saving than this, while others less. 64% of low-educated individuals, 71% of mid-educated individuals and 73% of high educated individuals would make more than half of their lifetime contributions between the ages of 45 and 66.

Table 3.5. Relative importance of saving in last half of working life – with and without earnings and employment uncertainty, and with contingent employer contributions

Model: includes EG, C, SLR, CEC	Percentage of lifetime contributions made between ages 45 and 66			Percentage of retirement wealth arising from contributions made between ages 45 and 66		
Education:	Low	Mid	High	Low	Mid	High
No uncertainty	74	77	83	61	64	72
With earnings and employment uncertainty:						
25th percentile	46	35	44	32	22	30
Median	63	68	71	46	53	56
Mean	58	62	64	47	51	52
75th percentile	82	84	84	71	72	73

Note: Model incorporates earnings growth (EG), children (C), student loan repayments (only applicable for high education, SLR) and contingent employer contributions (CEC). Rate of return is 3%. 'Retirement wealth arising from contributions' factors in the real return between when contributions are made and age 67.

4. Discussion

Summary of simulated results

Using a simple stylised economic model we have illustrated that there are good reasons to expect retirement saving rates to vary over the life cycle. In particular, earnings growth should lead people to save a higher proportion of their earnings in years in which their earnings are higher, and those with children would be expected to save substantially more either before their children arrive or after their children have left home. Taken together, these two factors suggest many individuals should postpone their saving for retirement until the last half of working life when children no longer impose costs on their parents – though the strength of this conclusion depends on the timing and costliness of children, the extent of earnings growth, and the rate of return on saving. When uncertainty over future employment and earnings is factored into the model, more saving would be expected earlier in working life, but on average saving would still be backloaded and saving rates would still be expected to increase substantially when children leave home.

Financial incentives to save in a pension – for example, in the form of contingent employer contributions – can heavily influence life-cycle saving profiles. With an arrangement where employer contributions may only be made if employees make a certain minimum contribution – similar to the automatic enrolment default arrangement – our simulated individuals would want to save at least that minimum amount throughout their working life. However, because that minimum would not be enough to maintain their working-life living standards into retirement, they would still substantially increase their contributions later in working life. This result holds for the different earnings profiles we model (with and without earnings uncertainty), and for a variety of assumed real rates of return on savings.

Other factors affecting the timing of retirement saving

The simple stylised model used in this work is not intended to be an accurate reflection of real life, but it is a useful tool for illustrating how and why different life patterns would be expected to affect saving rates over the life cycle. Here we briefly discuss some things that are not modelled, but which might have a bearing on the timing of retirement saving.

Housing

We discussed in Section 3 that a limitation of the model is that there is only one saving asset – the illiquid retirement saving vehicle – and that if the individual also had access to a liquid asset (or to borrowing) then the profiles of spending and saving in early working life might look quite different to those we model.

Another important asset that we do not model is housing. Nearly four out of five of individuals in their 50s and 60s are owner-occupiers, though the homeownership rate among generations born more recently is, at least for now, strikingly lower (Cribb, 2019). The purchase of owner-occupied housing, and the timing of this over the life cycle, could have important implications for the timing of retirement saving.[16] The effects will vary substantially across individuals though, depending, for example, on their preferences for housing versus spending on other goods and services, on the relative return on local housing versus pension saving, on their risk aversion and on their initial levels of wealth. Many individuals may be better off saving for a deposit and delaying starting saving in a pension until they have purchased a property, unless there are strong financial incentives – for example from employer contributions – to contribute to a pension. From that point on, how an individual would want to allocate their income between repaying a mortgage or saving for retirement will depend on the financial return on pension saving compared with the mortgage interest rate, on how risk averse they are and, again, on the financial incentives to save in pension. Some individuals may choose to pay off their mortgage as quickly as possible, and increase their pension saving at that

[16] The purchase of housing will also affect the overall *level* of saving (in non-housing assets) required for retirement, though exactly how will depend on whether or not an individual views it as an asset that they will liquidate (for example, through downsizing or equity release) to fund spending in retirement.

point (for example, if they are averse to holding debt, or if their mortgage interest rate is higher than the return on pension saving), while others may choose to hold a mortgage for as long as possible and save in a pension for a longer period.

Factoring in the existence of owner-occupied housing, and individuals' desire to purchase property as well as accumulate liquid financial resources for retirement, is unlikely to alter the main conclusions of this report – that retirement saving rates would be expected to vary over the life cycle, and to increase in later life as earnings grow and the expenses from children reduce. If anything, these qualitative findings may be strengthened if individuals want to divert saving towards the accumulation of housing wealth earlier in life.

Household pooling

Our model simulates the saving behaviour of an individual, albeit one who bears half of the cost of any children. In reality, many households have two adult earners, and decisions about saving for retirement may be taken jointly. In particular, the behaviour of an individual in a couple may deviate from what our model suggests for an individual if the two members of the couple have different returns to retirement saving – for example, due to different employer match arrangements (giving a benefit from the individual working for the employer with the better match arrangement doing more of the household pension saving), or different tax incentives to save (giving a benefit from the individual with the higher marginal tax rate doing more of the household pension saving, as they would benefit from greater tax relief). These factors could have an important effect of the level of pension saving done by each member of a couple, but this is unlikely to change the fundamental shape of the saving profile with respect to age that the household as a whole would want to choose.

Expectations of inheritances

Individuals in our model do not receive (or expect to receive) inheritances. The only sources of resources in our model are earned income, unemployment benefit (when earnings are uncertain) and the state pension. In practice, many individuals expect to receive an inheritance, and about one in four of those in their 30s expect to use a future inheritance to finance retirement.[17] While the expectation of future

[17] See Bourquin, Joyce and Sturrock (2020) and Crawford (2018).

inheritances, as well as the actual receipt of inheritances, has been shown to affect saving behaviour throughout working life,[18] this is again likely to affect the level of saving done overall, rather than the shape of the saving profile during working life.

Implications

The findings of this research use a simple economic model to demonstrate that rational well-informed individuals would not ideally want to have a constant retirement saving rate over their entire working life.

Government and industry policymakers should therefore carefully consider these life-cycle factors when debating and developing policies to encourage greater pension saving in the future. In particular, our results suggest that adjustments to automatic enrolment that result in non-constant saving rates over the life cycle – in particular, higher rates of saving at later ages – are likely to be better than across-the-board increases in default contribution rates. Examples of such policies that should be considered include default employee contribution rates that depend on age, increases in employee contribution rates that are triggered by earnings increases, and nudges to encourage individuals to increase their pension saving when their children leave home or when they finish debt repayments such as student loans or mortgages.

The findings also illustrate an important downside of defined benefit pension arrangements: that contributions cannot be easily varied over working life to suit an individual's circumstances or the timing of their capacity to save (in contrast to defined contribution arrangements that can come with a lot of flexibility). Accrual of benefits in career average schemes – which have replaced final salary schemes in the UK public sector – is relatively flat over working life. In the absence of age-varying accrual, there is little scope for employee contributions to vary over working life (as, for example, having flat accrual but employee contributions that rise with age would mean employer contributions being oddly more generous for younger employees than older employees – even if they were doing the same job on the same pay). Some employees therefore face a difficult choice of whether to join a defined benefit pension scheme with contributions at a higher level than they would ideally make at that stage in their lives, or not to join the pension and forgo

[18] See Bourquin, Joyce and Sturrock (2021).

the generous employer contribution. This lack of flexibility is an oft overlooked cost of public-sector pensions, which are often viewed enviously for their high generosity.

References

Anyaegbu, G. (2010), 'Using the OECD equivalence scale in taxes and benefits analysis', *Economic and Labour Market Review*, 4, 49–54.

Belfield, C., Britton, J. and van der Erve, L. (2017). 'Higher Education finance reform: raising the repayment threshold to £25,000 and freezing the fee cap at £9,250', Institute for Fiscal Studies (IFS) Briefing Note BN217, https://www.ifs.org.uk/publications/9964.

Bourquin, P., Joyce, R. and Sturrock, D. (2020), 'Inheritances and wealth inequality within generations', Institute for Fiscal Studies (IFS) Report R173, https://www.ifs.org.uk/publications/14949.

Bourquin, P., Joyce, R. and Sturrock, D. (2021), 'Inheritances and inequality over the life cycle: what will they mean for younger generations?', Institute for Fiscal Studies (IFS) Report R188, https://www.ifs.org.uk/publications/15407.

Coleman, L., Dali-Chaouch, M. and Harding, C. (2020), 'Childcare Survey 2020', Coram Family and Childcare, https://www.familyandchildcaretrust.org/childcare-survey-2020.

Costa Dias, M., Joyce, R. and Parodi, F. (2018), 'Wage progression and the gender wage gap: the causal impact of hours of work', Institute for Fiscal Studies (IFS) Briefing Note BN223, https://www.ifs.org.uk/publications/10358.

Crawford, R. (2018), 'The use of wealth in retirement', Institute for Fiscal Studies (IFS) Briefing Note BN237, https://www.ifs.org.uk/publications/12959.

Cribb, J. (2019), 'Intergenerational differences in income and wealth: evidence from Britain', *Fiscal Studies*, 40, 275–99.

Cribb, J. and Emmerson, C. (2020), 'What happens to workplace pension saving when employers are obliged to enrol employees automatically?', *International Tax and Public Finance*, 27, 664–93.

Deaton, A. and Paxson, C. (1994), 'Intertemporal choice and inequality', *Journal of Political Economy*, 102, 437–67.

Finch, D. and Gardiner, L. (2017), 'As Good As It Gets: The adequacy of retirement income for current and future generations of pensioners', Resolution Foundation Intergenerational Commission Report Wealth Series, https://www.resolutionfoundation.org/publications/as-good-as-it-gets-the-adequacy-of-retirement-income-for-current-and-future-generations-of-pensioners/.

Finch, D. and Pacitti, C. (2021), 'Building a Living Pension: Closing the pension savings gap for low-to-middle income families', Resolution Foundation Report, https://www.resolutionfoundation.org/publications/building-a-living-pension/.

Pensions Policy Institute (2019a), 'Increasing savings in automatic enrolment: analysis sponsored by Which?', PPI Research Report, https://www.pensionspolicyinstitute.org.uk/media/3180/20190603-increasing-savings-in-automatic-enrolment.pdf.

Pensions Policy Institute (2019b), 'Generation VeXed: solving the retirement puzzle', PPI Research Report, https://www.pensionspolicyinstitute.org.uk/media/3332/20191106-generation-vexed-report-final.pdf.

Scott, J., Shoven, J., Slavov, S. and Watson, J. (2021), 'Is automatic enrolment consistent with a life cycle model?', National Bureau of Economic Research (NBER) Working Paper No. 28396.

Appendix

Estimating earnings profiles and earnings uncertainty

In our model, earnings are one of the following: (a) constant across working life (our baseline model); (b) change with age but are certain (in Sections 3.1 and 3.3–3.5); (c) change with age but are uncertain (in Section 3.6).

Here, we describe the method used to construct the levels of earnings and the way in which they change when earnings are uncertain (i.e. case c). To obtain the level of earnings when earnings are certain but change with age (i.e. case b), we take the mean of all of the levels of earnings from the simulated uncertain earnings and smooth the relationship between earnings and age by fitting a quartic in age. To obtain the level of earnings when they are constant (i.e. case a), we take the level of earnings such that the discounted value of lifetime earnings is the same as the discounted value of earnings in case b.

When earnings are uncertain, individuals are either not in work and have zero earnings, or are in work and have positive earnings. If they are in work, an individual's level of earnings is the sum of an average earnings profile, which is certain and known in advance, and an uncertain component. We first describe how the certain profile is constructed, before turning to the uncertain component.

Estimating earnings profiles

The earnings profiles are estimated using data from the Family Expenditure Survey and its successor surveys. We use data from the years 1978–2018 (all years for which a measure of individual education is available) and all individuals born between 1930 and 1989 who are between the ages of 22 and 64 and in work. For each of our three education groups, the earnings profiles aim to capture the average level of earnings that an individual of a particular generation with that level of education, who was in work at each age, would receive. The earnings profiles we use in our model are for individuals born in the 1980s. As these individuals have not reached retirement age yet, we must infer their future earnings profile from those born earlier. We infer these earnings profiles using the following assumptions.

1. For a given education group, the *relative* levels of earnings between two ages are the same across cohorts. For example, if graduates in earlier cohorts earned, on average, twice as much at age 50 than they did at age 30, the same will be true for graduates born in the 1980s.
2. For a given education group, there can be differences in the absolute levels of earnings across cohorts but these will be the same across ages. For example, if graduates born in the 1980s earned, on average, 20% more at age 25 than graduates from earlier cohorts earned at age 25, then 1980s graduates will also earn 20% more than the graduates of earlier cohorts at age 50.
3. In addition to changes in earnings that occur as a result of being at different ages, the levels of earnings for a particular cohort–education group may vary due to economic conditions. Such changes are assumed to affect all cohort and education groups in the same way. For example, if earnings in 2009 are lower than they would otherwise be expected to be (because of, for example, the recession in that year), then this is assumed to have reduced the earnings of all cohort–education groups by the same percentage amount.

The first (or second) assumption is necessary in order to extrapolate the earnings profile into the future for the 1980s cohort. The third assumption is necessary to disentangle temporary economic shocks, and their effects on earnings, from changes in earnings due to the normal changes in earnings with age. Our method finds the earnings profiles that are consistent with the above assumptions and provide the closest fit to the earnings data.[19]

[19] Specifically, we run an ordinary least-squares regression of individual log earnings on indicator variables for cohort–education groups, the interaction of a full set of age dummies and education group, a linear time trend, and a series of indicator variables for each year that are constrained to sum to zero (the final restriction is suggested by Deaton and Paxson, 1994). When creating the simulated profiles, we set the time effects equal to those for 2018.

Estimating levels of earnings and earnings transitions when earnings are uncertain

To estimate the uncertainties in earnings, we use data that follow individuals over time from the UK Household Longitudinal Study. We first take the difference between individuals' earnings and the deterministic profile that was estimated in the step described above. This tells us, for each individual, how much their earnings are above or below the average for someone of their age, education group and generation. We call this the individual's earnings 'shock'. The distribution of these shocks gives us the distribution of shocks that can be experienced by an individual of that given education group and age. For example, if 5% of individuals in the data have earnings that are at least 150% higher than the deterministic earnings profile, the same will be true in our simulated earnings.

We then use a method to estimate, for those of each education level, age and level of earnings 'shock', what the probability is that they will have an earnings 'shock' of a different possible level the following year. We combine this with (i) information on the probability that those of different levels of education, age and earnings shock move from being in work to being out of work and (ii) information on the probability, for those of each education level and age who are out of work, of moving into work and of having different levels of earnings shock. Ultimately, this gives us, for each education level and age, probabilities of moving between different parts of the distribution of earnings and the probability of moving in and out of work, given where that individual previously was in the earnings distribution.

Figure A.1 shows these 'transition' probabilities for the case of a mid-educated individual at age 50. For example, this tells us that someone who is not in work at age 50 is simulated to have an 82% change of not being in work at age 51 and a 2.5% chance of moving into work and being in the bottom 5% of earners (1%–5% category). Someone in the top 5% of earners (96%–100% category) at age 50 has an 84% chance of being in the same category the following year and an 11% chance of seeing a fall in earnings taking them to between the 86th and 95th percentiles.

Figure A.1. Probabilities of individuals moving to selected parts of the earnings distribution at age 51, by position and work status at age 50, for mid-educated individuals

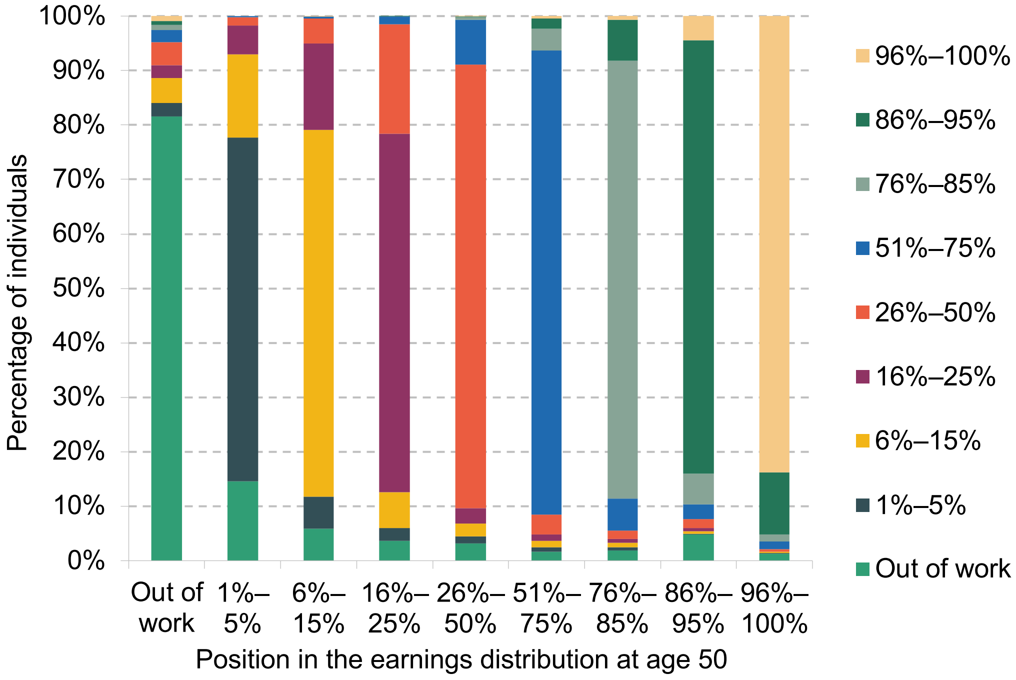

Source: Estimates using the UK Household Longitudinal Study.