

Sturrock, David

Working Paper

Wealth and welfare across generations

IFS Working Papers, No. 23/15

Provided in Cooperation with:

Institute for Fiscal Studies (IFS), London

Suggested Citation: Sturrock, David (2023) : Wealth and welfare across generations, IFS Working Papers, No. 23/15, Institute for Fiscal Studies (IFS), London, <https://doi.org/10.1920/wp.ifs.2023.1523>

This Version is available at:

<https://hdl.handle.net/10419/284241>

Standard-Nutzungsbedingungen:

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

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

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

Terms of use:

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

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

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

David Sturrock

23/15

Working paper

Wealth and welfare across generations

Wealth and Welfare across Generations*

David Sturrock[†]

April 10, 2023

Abstract

Although older generations have substantially more wealth than their recent predecessors did at the same age, younger generations do not. Bringing together UK data on those born between the 1930s and 1980s and a lifecycle model of saving, I quantify whether this is due to changes in preferences or changes in the circumstances each generation has faced. Changing circumstances can rationalise slowing generation-on-generation wealth growth. I find no evidence that later-born generations are less patient. I quantify the implications of changing circumstances for consumption and welfare. Later-born generations are predicted to have higher consumption, despite accumulating no greater wealth, than their predecessors because their earnings are more ‘backloaded’, they have fewer children, and face lower taxation.

Key words: Intergenerational inequality; wealth accumulation; lifecycle

JEL classification: D15, D31, E21, E24

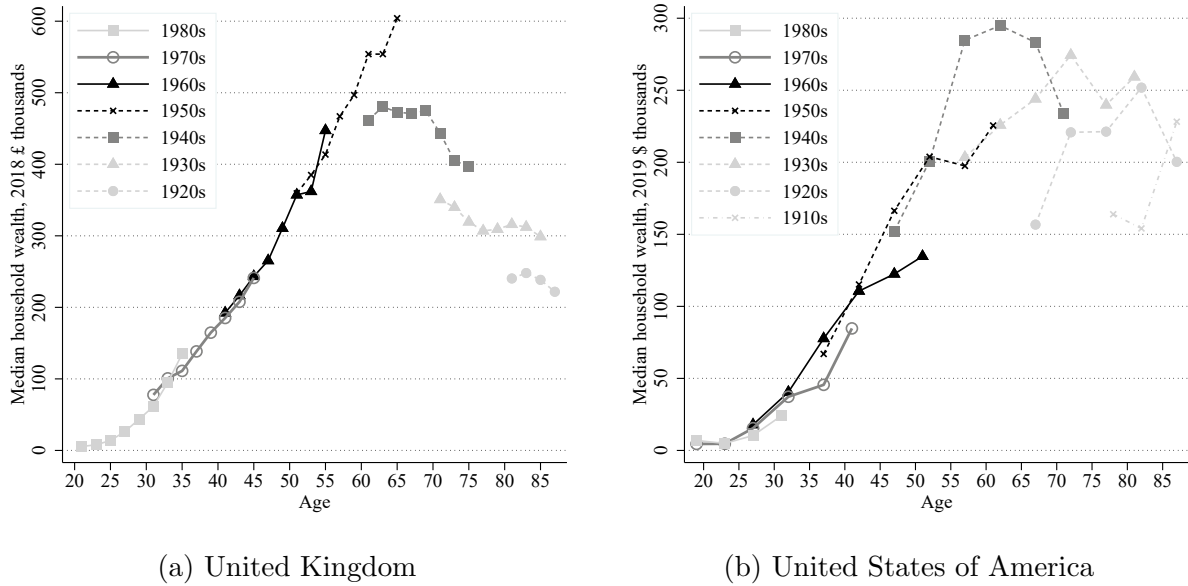
*Funding from the Economic and Social Research Council (through Research Grant number ES/V001248/1 and through the ESRC Centre for the Microeconomic Analysis of Public Policy (CPP) (ES/M010147/1)) is gratefully acknowledged. I thank participants at seminars at the Institute for Fiscal Studies, University College London and Princeton University, and at the International Institute of Public Finance Annual Congress (2020). I am also grateful to Richard Blundell, Rowena Crawford, Mariacristina De Nardi, Eric French, Rory McGee, Cormac O’Dea, Morten Ravn, Richard Rogerson and Gianluca Violante for helpful comments and suggestions. All errors are my own. Correspondence to david.sturrock.17@ucl.ac.uk.

[†]Institute for Fiscal Studies and University College London

1 Introduction

Across advanced economies, older generations have substantially more wealth than their recent predecessors, but younger generations do not.¹ In the UK and USA, generations born up until the middle of the 20th century have substantially higher average household wealth than those born 10 years earlier had at the same age. However, generational wealth progression for subsequent generations has stagnated and, in the USA, gone into reverse (see Figure 1).² While, all-else-equal, higher wealth expands future consumption possibilities, drawing conclusions from intergenerational wealth comparisons is difficult because generations are observed at different points in their lifecycle and face very different economic and demographic circumstances.

Figure 1: Median household net wealth by age and decade of birth



Note: UK data is from the Wealth and Assets Survey (2006 to 2018). The UK household net wealth measure is the sum of all housing, financial, private pension and physical wealth, less all debts. The USA data is the ‘net worth’ measure from the Survey of Consumer Finances (1989 to 2019) and, in contrast to the UK measure, excludes private defined benefit pension wealth.

There are three types of potential explanations for these generational wealth patterns, each with different implications for generational differences in consumption and for policy.

¹See Bartels and Morelli (2021) for analysis of Germany and Italy, Bauluz and Meyer (2021) for evidence from France and the USA, Bauluz and Meyer (2022), Jaeger and Schacht (2022) and Gale et al. (2020) for further analysis of the USA and Cribb (2019) for evidence from the UK.

²For a comprehensive discussion of these patterns and comparison over a longer time period and larger number of generations, see Bauluz and Meyer (2022) and Jaeger and Schacht (2022).

One possibility is that generational growth in lifetime economic resources and consumption possibilities has slowed. This could result from the increases in household earnings and asset prices, which drove wealth growth across older generations, failing to materialise for younger generations, or even making wealth accumulation more difficult (by lowering future returns and making homeownership less accessible, for example).³ This interpretation underlies claims that the young have received a less favourable economic settlement than the old and that redistributive policies should account for this.

A second possibility is that younger generations will receive a greater share of their lifetime income later in their lifecycle and so do more of their saving at older ages. Younger generations enter the workforce later and are set to retire later and receive larger inheritances than prior generations (Gale et al. (2020); Bourquin et al. (2021)). In this case, the lack of generation-on-generation growth in wealth at younger ages need not imply slowing growth in lifetime income and consumption.

A third possible explanation is that younger generations prioritise current consumption, at the expense of saving, more than older generations did. This interpretation underlies concerns about younger generations’ ability to achieve an adequate retirement income and calls to increase household saving rates by, for example, increasing default pension contribution rates.⁴ This possibility is made more consequential by the decline of employer-provided defined benefit pensions and growing longevity at older ages.

Understanding why generation-on-generation increases in wealth have stopped, and what this is likely to mean for the welfare of different generations, is therefore crucial to determining how policy should respond to these trends. Yet, to date, there has been no comprehensive empirical quantification of the circumstances that have changed across generations and their ability to explain observed wealth patterns.

In this paper, I bring together a wide range of micro-data sources to estimate how the following economic and demographic circumstances have changed across the generations born in the UK between the 1930s and 1980s: the earnings process, the tax and welfare system, the public pension system, occupational pensions, the rate of return on wealth, inheritances, longevity, and household size and composition. I then estimate a heterogeneous agents lifecycle model of wealth accumulation that incorporates these changing circumstances

³The idea of ‘asset price redistribution’ from the young to the old is examined in Fagereng et al. (2022).

⁴For example, the UK Government’s 2017 review of Automatic Enrolment into workplace pensions claims “People are saving, but they are not yet saving enough to ensure that they will have the retirement they may want.”. This interpretation has gained attention in popular discussions of saving choices. For example, in May 2022, the New York Times reported that “many adults under 35 are [...] saving less, spending more and pursuing passions.” Australian businessman Tim Gurner forged the image of the avocado-eating millennial in 2017, saying “When I was trying to buy my first home, I wasn’t buying smashed avocado for \$19 and four coffees at \$4 each. [...] We’re at a point now where the expectations of younger people are very, very high.”

across generations. This model features rich heterogeneity across households and stochastic processes for employment, earnings, returns to wealth, inheritances, longevity and end-of-life medical costs, that capture the different risks and constraints faced by each of the six decade-of-birth generations I examine. I use this model to quantify the extent to which the observed generational wealth patterns in the UK can be explained by changing circumstances, as opposed to changes in preferences for saving, across generations. Finally, I consider whether this slowing of generational wealth growth should be expected to translate into a slowing in the growth of living standards across generations.

I find that the changes in circumstances faced across generations can largely explain the empirical patterns in wealth. Changing circumstances alone can generate the large generational increases in wealth between the 1930s, 1940s and 1950s-born generations as well as the lack of generational wealth growth for the three generations born subsequently. While I estimate levels of patience that are lower for younger than for older generations, the change is quantitatively small. Consequently, declining patience can rationalise at most a 7% lower level of wealth at age 65. This compares to an estimated doubling of wealth at age 65 between the 1930s and 1950s generations. Furthermore, even the lowest estimated level of patience (0.992 for those born in the 1980s) is towards the higher end of estimates from the lifecycle saving literature (see Crawford and O’Dea (2020) for the UK, for example). I find no evidence that younger generations appear impatient, either relative to their predecessors or in absolute terms.

The decline of generation-on-generation increases in earnings is the main factor rationalising the fact that generational wealth growth has stopped. Average lifetime earnings were around 50% higher for the 1950s-born compared to the 1930s-born generation but the 1980s-born are expected to earn only around 20% more than those born in the 1950s, on average. This translates quite directly into a slowing of generation-on-generation growth in wealth. Second most important is the change in the shape of the age-profile of earnings. The expansion of the length of working-life among older generations can rationalise a 10% higher level of retirement wealth for the 1950s-born generation compared to the 1930s-born generation. For younger generations, higher levels of education - resulting in later-entry into the workforce and more steeply increasing earnings profiles - can rationalise the slowing of generational wealth growth at younger ages. I find that, all-else-equal, these changes would lead to a 27% *lower* level of wealth at age 33 for the 1980s-born generation compared to the 1950s-born. A number of other factors including the slowing of increases in life expectancy, the withdrawal of generous occupational defined benefit pension arrangements, and the increasing size of inheritances across generations also contribute. The generational shift whereby those born later are having fewer children, and having them later, increases

wealth accumulation in early adult life. Without this change, generational wealth growth at younger ages would have been slower still.

Turning to the implications of these trends for consumption and welfare, I find that welfare is expected to grow more rapidly across working-age generations than is wealth. While the changing age profile of earnings drives higher wealth at retirement and lower wealth at younger ages, it has less significant implications for consumption. It represents a retiming of lifetime income rather than a change in its level. Younger generations also have smaller households and face lower taxation.

Finally, given the importance of housing wealth in households' portfolios and the large increase in house prices compared to incomes between the mid-1990s and late-2000s, I produce a counterfactual simulation of wealth and consumption in which house prices grow in line with average earnings after 1995. Relative to this counterfactual, I estimate that the house price boom increased median consumption of the 1930s-born generation at age 70 (the end of the boom period) by 13%, whereas it is expected to *decrease* median consumption at that age for the 1980s-born by 13%.

The main contribution of this paper is to make the first comprehensive and quantitative assessment of the drivers of the slowing of generation-on-generation increases in wealth in an advanced economy. This is important given the different policy implications of the various explanations given for these trends. Younger generations' wealth levels are consistent with a degree of patience that is both high and similar to that for older generations. This means that addressing a perceived shortfall in wealth through policies to increase saving rates may not be welfare increasing.⁵ This suggests that if governments wish to see generational increases in wealth begin again then this will require either an increase in the growth of earned incomes or the returns to wealth, or an increase in the share of these that flow (directly, or indirectly through redistribution policies) to younger generations.

The second contribution of this paper is to quantify the implications of the changing economic and demographic circumstances for generational growth in consumption and welfare. This is important because when the timing of income, returns to wealth, timing of consumption needs, government policy, and length of lifetimes can vary, wealth at a given age is not a sufficient statistic for lifetime consumption and welfare. I find that, because some of the slowing of generational wealth growth is due to the retiming of earnings and corresponding 'steepening' of the wealth age profile, generational growth in living standards should be expected to slow less rapidly than wealth comparisons might first suggest.

⁵That is not to say that saving policies already in place, such as automatic enrolment into workplace pensions, are ill-advised, since current wealth levels reflect the past effects of these policies. However, my findings do not make a case for policies to further increase saving rates at younger ages.

Related literature

A recent literature has begun to document generational differences in wealth accumulation in advanced economies. Jaeger and Schacht (2022) shows that in the USA, median household wealth grew across generations until the 1940s-born. For those generations born since, median wealth has declined across generations. Bauluz and Meyer (2022) also documents this pattern and uses an accounting framework to decompose changes in the wealth to income ratio into saving, capital gains and inheritances. A slowing or reversal of generation-on-generation increases in household wealth have also been documented in the UK, France, Italy and Germany (see, for example, Bartels and Morelli (2021), Bauluz and Meyer (2021) and Cribb (2019)). I move beyond documenting these patterns to disentangle their drivers. By explicitly modelling the changing economic and demographic factors that different generations face, and households' responses to them, I can decompose the full impact of each of these factors, and quantify their implications for consumption and welfare.

A small number of papers have considered the role of changing economic conditions in driving generational differences in household outcomes. Within lifecycle frameworks, Borella et al. (2020) quantifies the implications of changes in wages, life expectancy and medical expenses for labour supply and wealth accumulation in the US, Paz-Pardo (2021) finds that the increase in earnings risk can explain half of the decline in homeownership across US generations, and Borella et al. (2022) estimates the effect of US tax reforms on saving behaviour. Glover et al. (2020) considers the impact of the great recession on different cohorts, through earnings and asset price channels. Studying the Netherlands, Kapteyn et al. (2005) find that differences in labour productivity and social security entitlements can explain cohort wealth differences. Gale et al. (2019) and Gale et al. (2020) discuss the economic conditions and prospects for wealth accumulation of 'millennials'. Malmendier and Nagel (2011) find that individuals who experienced lower returns to wealth during their life are less willing to take financial risks and to invest in the stock market. Relative to these studies, I consider a more extensive set of economic and demographic circumstances that could drive wealth differences, allowing me to make the first comprehensive quantitative decomposition of wealth differences over six generations.

There is a literature that has assessed the optimality and adequacy of saving among the currently retired. Engen et al. (1999) and Scholz et al. (2006) solved and simulated lifecycle models of wealth accumulation, assuming a level of patience. Comparing simulated and observed wealth levels, they conclude that the saving choices of those born in the 1930s and early 1940s are largely consistent with optimal behaviour. Using data on lifetime earnings and accumulated wealth of those born in the 1940s in England, Crawford and O'Dea (2020) find that households' saving choices can be rationalised by conventional levels of patience.

These studies have focused on older generations, for whom a full assessment of their working life saving decisions is possible. This paper turns the attention to today’s working age generations. This is important because the decline of generous occupational defined benefit pensions, combined with increases in longevity, mean that households’ decisions over private saving are increasingly consequential for their retirement living standards. It is therefore crucial to assess whether the lack of generational growth in wealth may represent undersaving on the part of younger generations.

There is increasing attention given to the extent and determinants of wealth inequality in advanced economies, including top-end wealth shares (Saez and Zucman (2016); Smith et al. (2022); Hubmer et al. (2021); Kuhn et al. (2020)), racial differences in wealth (Derenoncourt et al. (2022)), and wealth inequalities by parental background (Charles and Hurst (2003); Fagereng et al. (2021)). I contribute by turning the focus to *generational* inequalities in wealth. As part of the explanation of recent decades’ wealth trends, Auclert (2019), Greenwald et al. (2022) and Fagereng et al. (2022) have demonstrated how declining real interest rates and consequent asset price changes redistribute between households depending upon their portfolio composition and consumption plans. I also richly model how asset price shocks feed through to households’ wealth and consumption. I contribute by placing these effects alongside other changes in economic conditions and policy, allowing a more comprehensive assessment of different generations’ consumption and welfare.

This paper proceeds as follows. Section 2 gives an overview of the key economic and demographic trends across generations. Section 3 lays out the quantitative lifecycle model of wealth accumulation. Section 4 describes the estimation of the model’s exogenous processes, which capture the economic and demographic circumstances of each generation. Section 5 sets out the estimation of the model’s preference parameters. Section 6 sets out the results and Section 7 concludes.

2 Key economic and demographic trends

I set out the key economic and demographic trends using a number of data sources described in detail in Section 4. In general, the unit of analysis is the household (single individual plus their partner and any dependent children) with a couple’s age defined as their mean age.

Household earnings and employment: Figure 2 panel (a) shows mean household earnings at each age for each generation. It illustrates that earnings at comparable ages grew strongly across older generations, increasing by 34% between the 1930s and 1940s-born generations and increasing by 18% between the 1940s and 1950s-born generations, but sub-

sequent generation-on-generation growth in earnings has been much reduced. Also notable is the ‘steepening’ of the age profile of earnings across generations. While average earnings grew by 68% between age 20 and age 35 for the 1950s-born generation, for the 1970s-born generation they grew by 154%. Panel (b) shows that across the core earnings years of the lifecycle, the proportion of households where someone is in work has been very similar across generations, at around 80%, but that there was a trend of increased working at older ages across the 1950s-born and 1940s-born generations, driving higher earnings at older ages, compared to their predecessors, and less working at younger ages when looking across generations (contributing to the ‘steepening’ of the average earnings profile). There was also a substantial increase in earnings risk and earnings inequality which saw the variance of log household earnings roughly double over the period from around 1980 to the early-1990s (Blundell and Etheridge (2010)).

Figure 2: Average earnings and employment rate of households, by age and decade of birth



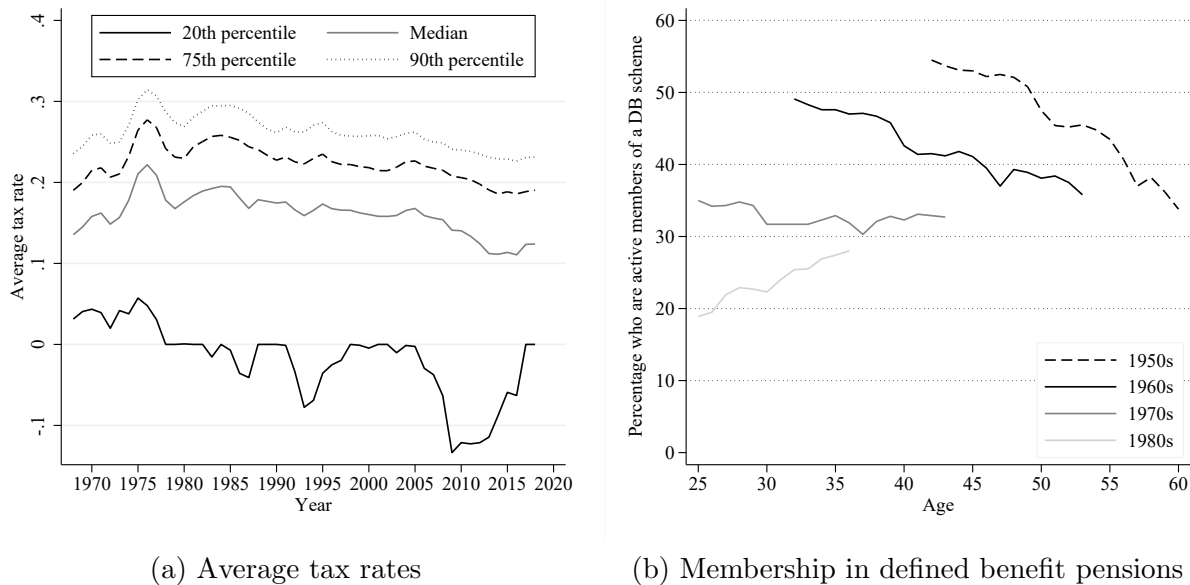
Note: A household is defined as an individual and any partner and dependent children. Age and decade of birth of couple households is defined based on the mean of their age and year of birth. A household is defined as employed if any of its members are in paid employment. Source: Family Expenditure Survey and its successors, 1968-2018.

Tax and welfare systems: As in the USA, tax rates on earned income have been on a decreasing path since the late 1970s. The basic rate of income tax was 30% in 1975-76 and now stands at 20%. A reduction in higher rates culminated in the consolidation, in 1988-89, of a range of higher rates between 40% and 60% into one 40% band. These sharp falls in income tax rates have been partly offset by increasing rates of National Insurance

contributions, a payroll tax. Tax thresholds have generally increased over time in real terms. Turning to welfare, the main rate of unemployment benefit has stayed relatively constant over time in real terms. However, there has been a significant expansion of the type and generosity of a range of other welfare payments, including the introduction of tax credits in the early 2000s (paid to those on low incomes) and the expansion over time, and particularly in the 1990s, of the number of individuals eligible for housing benefit and disability benefits. From 2013, there were widespread real-terms cuts to benefits.⁶

Figure 3 panel (a) shows selected points of the distribution of households' average tax rates in each year. We can see the gradual reduction in household tax rates across the top-half of the distribution from the mid-1970s onwards. At the bottom of the distribution, we can see the effects of the introduction of tax credits in the mid-2000s. Some fluctuations are attributable to changing economic conditions as well as policy. For example, the early 1990s and 2009, when tax rates fell at the bottom of the distribution, both saw significant increases in unemployment.

Figure 3: Selected points on the distribution of average tax rates and proportion of employees who are members of a defined benefit pension arrangement



Sources: (a): Family Expenditure Survey and its successors; (b): Cribb (2019) estimates using the Annual Survey of Hours and Earnings. Note: A household's average tax rate is defined as one minus their post-tax income as a percentage of their pre-tax earnings and includes the effects of both taxes and welfare transfers.

⁶For a comprehensive account of the changes in the UK tax system over time see Pope and Waters (2016) and for a comprehensive review and history of the benefits system see Hood and Oakley (2014) and Hood and Norris Keiller (2016).

Public pension system: The UK public pension began as a relatively flat-rate system where entitlement depended on years of work but not earnings. A significant earnings related component was added in 1978, but the generosity of accrual to this component was gradually eroded over time. From 2016 onwards, the state pension will be flat rate again, with no earnings-related component and broad activity requirements mean that entitlement is near-universal.⁷ Different generations have therefore faced accrual of different pension benefits depending on which systems and rules were in place during their working-life. The overall trend across the generations examined is of a declining entitlement for higher-earners and an increase in generosity for those less attached to the labour market and with lower earnings.

Occupational pensions: A key trend since the 1990s, and affecting cohorts from the 1950s onwards, is the decline of defined benefit (DB) pensions being offered by private sector employers. Due to increases in longevity, declines in returns, and changes to regulations, many private sector DB funds moved into deficit and employers have tended to close schemes to new entrants. New accrual of DB pension rights is uncommon outside of the roughly 20% of employees working in the public sector. Figure 3 panel (b) shows the differences in DB prevalence across employees in different generations from the 1950s onwards. Defined contribution (DC) pensions have become much more prevalent as DB has declined but employers tend to offer rates of contributions to DC pensions that are worth much less than the pension rights accrued in the DB pensions they have replaced. In the 1990s and early 2000s, the average value of pension rights accrued from an extra year in a DB scheme was 18.9% of earnings in the private sector and 25.5% of earnings in the public sector (Disney et al. (2009)). Even taking into account the employee contributions required as part of these DB arrangements, these rates represent a much more valuable employer pension provision than the DC contributions that are more common for younger generations.

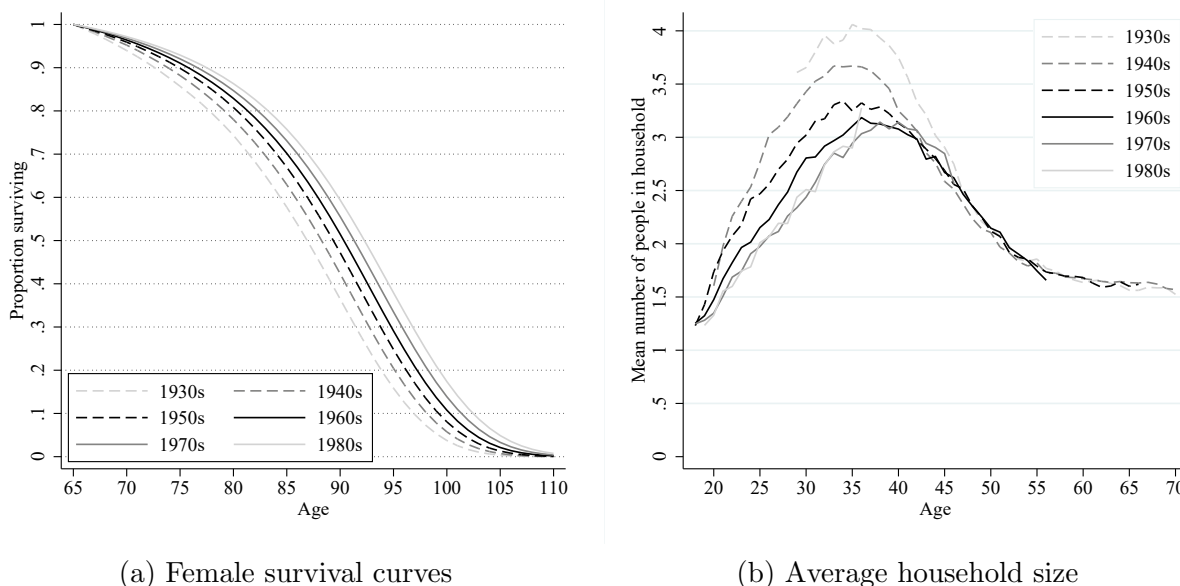
Longevity: In the UK as in many countries around the world, longevity at older ages has increased substantially for those born later, and this trend is expected to continue. Figure 4 shows cohort survival curves for 65-year-old women born in the middle year of each generation. While it is projected that 36% of women born in 1935 who survived to age 65 will go on to survive to age 90 or older, for those born in 1985 this figure is projected to be 59%. There is an even starker increase for men. While the expansion in longevity is expected to continue across generations, the rate of expansion is expected to slow. A man born in 1930 who reached the age of 65 was expected to live for a further

⁷For a full explanation of the history of the UK State Pension system, see Bozio et al. (2010) and Banks and Emmerson (2018).

16.4 years. This rose to 20.5 years for a man born twenty-five years later, in 1955, a 25% increase. Those born a further 25 years later in 1990, and who attain the age of 65, are expected to live on average a further 23.4 years, a increase in life expectancy at age 65 of 14%.

Household size and composition: Later-born generations are having fewer children and having them later in life. Figure 4 shows the mean number of people per household at each age, for each generation.⁸ I take the mean over all households, including singles and couples, with the household age for couples defined as their mean age. Those born in the 1930s had an average of four people per household at age 35. This ‘peak’ number of people is progressively lower and later amongst later-born generations, and was just over three for the 1970s generation, reached at age 40, with the 1980s-born following a similar pattern.

Figure 4: Trends in female survival probabilities and in household size, by decade of birth



Sources: (a): Office for National Statistics cohort survival curve for England and Wales for the middle birth year of each generation, conditional on survival to age 65; (b): FES.

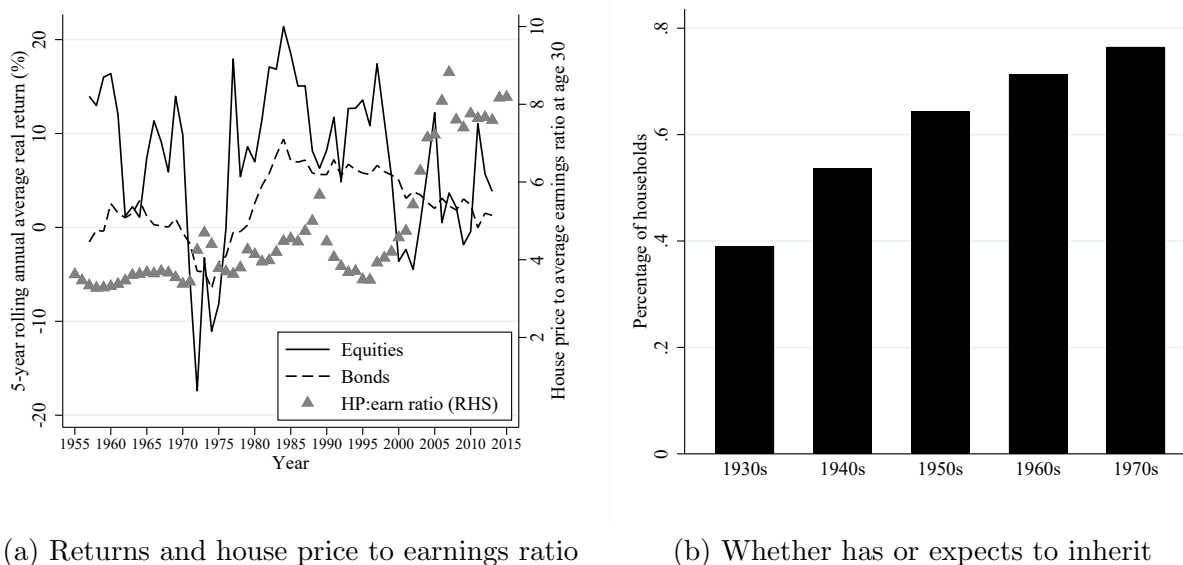
Returns to wealth: Figure 5 panel (a) shows a rolling 5-year average of the annual real total return to the equities and bonds over time. It also shows the house price to average household earnings ratio in each year. Equities had particularly bad years in the early-1970s and saw strong returns through most of the 1980s. Real equity returns have tended to be lower over the 2000s and 2010s than over the 1980s and early-1990s, reflecting the lack of a boom period comparable to the 1980s and the negative effects of

⁸A household is defined as an individual and their partner, if they have one, and any dependent children.

the financial crisis of 2008 and 2009. Returns to safe financial assets peaked in early-1980s and have declined steadily since, in line with the general decline in safe returns over time across advanced economies. Until the late-1980s, house prices oscillated around levels four times the value of average annual household earnings. They then rose dramatically, roughly doubling compared to earnings over the 1990s and early 2000s. This increase in house prices compared to earnings coincided with a decline in rates of homeownership for young people, with the homeownership rate at age 30 falling from around 60% for those born in the 1950s and 1960s to around 40% for those born in the 1980s.

Inheritances: As a result of the substantial growth in wealth across older generations and decline in the rate of child-bearing, inheritances have become more common across generations. A much larger proportion of younger households expect to inherit than did inherit in older generations. Figure 5 panel (d) shows that less than 40% of households of those born in the 1930s received an inheritance, but almost 80% of those born in the 1970s expect to inherit. Later-born individuals also expect to inherit larger sums. While fewer than one in twelve households born in the 1930s inherited over £100,000, over 40% of households born in the 1970s expect to inherit this amount or more (adjusting for inflation).

Figure 5: Trends in returns, asset prices and inheritances: 5-year rolling average real return to equities and bonds and house price to average earnings ratio by year, and proportion of households who have or expect to receive an inheritance by decade of birth



Note: Panel (a) shows the 5-year rolling geometric average of the total annual real return (capital gains and income) for both UK equities and government medium-length maturity bonds. Sources: (a): Jorda et al. (2019) and Nationwide House Price index; (b) Wealth and Assets Survey.

3 Model

In brief, households are one of several *ex ante* heterogeneous types who choose each period how much of their available resources to consume and how much to save. Households face uncertainty over employment, a stochastic component of earnings, returns to wealth, whether they have living parents and inheritance receipt at the time parents die, mortality and end-of-life costs. Government provides a tax-and-transfer system and public pension system and firms provide occupational pensions. Household utility comes from equivalised consumption and bequests, which depend on assets left over at death.

3.1 Household types, demographics, and preferences

Types: Households, indexed i , belong to one of a set of *ex-ante* heterogeneous types, defined by a combination of:

- Generation of birth, $g_i \in \{1930s, 1940s, 1950s, 1960s, 1970s, 1980s\}$
- Education, $ed_i \in \{low, mid, high\}$
- Earnings fixed effect, $\zeta_i \in \mathbb{R}$
- Whether the household has a defined benefit pension, $db_i \in \{0, 1\}$.

I denote the vector that defines household i 's type as $\omega_i = (g_i, ed_i, \zeta_i, db_i)$. Education levels *low*, *mid* and *high* corresponding, respectively, to: compulsory schooling only (i.e. high school dropout), further schooling but no higher education (i.e. high school graduate), and some higher education (i.e. some college education).⁹

Demographics: Households are modelled from the age of 24 until death. Each period represents one year. From the age of 65 onwards households may die and their probability of survival to the next year varies by household type and age and is denoted $s_{\omega,t+1}$. Death occurs at age 110 at the latest. The household's size changes in a deterministic way with age to account for the arrival of children and the death of the first member of a couple. Household size impacts decisions in the model through the equivalisation factor, $\theta_{\omega,t}$, which varies with age and household type.

⁹I do not separately model single and couple households and households are modelled as unitary agents. I define the generation of birth of a couple household as the decade corresponding to their mean year of birth. The education level of a household corresponds to the highest education level across the two members of the couple. Defined benefit pension membership takes the value one if either member of a couple has a defined benefit pension.

Preferences: Households get utility from consumption and bequests. Within-period utility is a constant relative risk aversion function (with parameter γ) of equivalised consumption, $c_{i,t}$:

$$u(c_{i,t}) = \theta_{\omega,t} \frac{(c_{i,t}/\theta_{\omega,t})^{1-\gamma} - 1}{1-\gamma} \quad (1)$$

There is a warm-glow bequest motive (following De Nardi (2004)). Bequests are denoted b_i and the parameters ϕ_1 and ϕ_2 govern the strength of the bequest motive and the extent to which bequests are a luxury good, respectively. Bequest utility is given by:

$$\phi(b_i) = \phi_1 \frac{(b_i + \phi_2)^{(1-\gamma)}}{1-\gamma}. \quad (2)$$

3.2 Sources of uncertainty

At age 24, households get an initial draw of employment status and the stochastic component of earnings and have a living parent household. At the start of each subsequent period, a living parent household may die (resulting in an inheritance draw), and employment status and the stochastic component of earnings evolve. The household's type and final earnings realisation determines their pension income. From age 65, the household's own survival is realised at the beginning of each period. Death results in the realisation of an end-of-life cost shock with remaining assets bequeathed. Each period, there is a stochastic return to assets. I now described these elements in detail.

Employment and earnings: Households receive earnings, $e_{i,t}$ from age 24 until some known latest retirement age K_ω , that varies by generation. I denote the household's binary employment status as $E_{i,t}$. When households are not employed, representing voluntary or involuntary unemployment or early retirement, earnings are zero. When employed, log household earnings are the sum of a deterministic component that varies by household type and age, $f_{\omega,t}$, the household fixed effect, ζ_i , and a persistent stochastic component, $\eta_{i,t}$:

$$\ln(e_{i,t}) = f_{\omega,t} + \zeta_i + \eta_{i,t} \quad E_{i,t} = 1 \quad (3)$$

$$e_{i,t} = 0 \quad E_{i,t} = 0 \quad (4)$$

Following Arellano et al. (2017), I assume that the stochastic earnings component, $\eta_{i,t}$ is drawn from a distribution that varies with generation, education, age, and its lagged value $\eta_{i,t-1}$ and is given by the series of conditional quantile functions $Q_{\omega,t}(\cdot|\eta_{i,t-1})$. This form flexibly allows for nonlinear persistence, non-normality, and age-dependence of earnings shocks.

These empirically important aspects of earnings processes have quantitatively important implications for wealth accumulation choices in lifecycle models (De Nardi et al. (2019)). Employment status is drawn from a distribution that varies by generation, education, age, earnings fixed effect and lagged employment status. The draws of the stochastic earnings component and employment status are independent across time and independent of each other and the draws of the stochastic earnings component are independent of the draw of the household fixed effect. These conditions can be written as

$$\eta_{i,t} = Q_{\omega,t}(u_{i,t}|\eta_{i,t-1}) \quad (5)$$

$$E_{i,t} = \mathbb{1}\{v_{i,t} > \bar{v}_{\omega,t}(E_{i,t-1})\} \quad (6)$$

$$u_{i,t} \stackrel{iid}{\sim} U(0, 1) \quad (7)$$

$$v_{i,t} \stackrel{iid}{\sim} U(0, 1) \quad (8)$$

$$u_{i,t} \perp\!\!\!\perp v_{i,s}, \quad \forall \quad t, s. \quad (9)$$

Public pensions and occupational pensions: From age K_ω onwards, households receive a public pension payment. The public pension in year K_ω is a type-specific function of the household's final earnings, $pub_\omega(e_{i,K_\omega-1})$. Households may be a member of an occupational defined benefit pension scheme, as denoted by the binary variable db_i . Households who are members must contribute fraction $q_{i,t}$ of their earnings as contributions to the DB scheme. The defined benefit pension income in year K_ω varies by type and is denoted, DB_ω .¹⁰ Those who are not members of the DB scheme (i.e. $db_i = 0$) receive no such income but instead receive an employer pension contribution, modelled as a negative value of $q_{i,t}$ i.e. an addition to gross earnings. From $K_\omega + 1$ onwards, the public and defined benefit pension income grows in proportion to the equalisation factor.¹¹ Formally:

$$\begin{aligned} p_{i,t} &= 0 & t < K_\omega \\ p_{i,t} &= pub_\omega(e_{i,K_\omega-1}) + DB_\omega \cdot db_i & t = K_\omega \\ p_{i,t} &= p_{i,K_\omega} \frac{\theta_{\omega,t}}{\theta_{\omega,K_\omega}} & t > K_\omega \end{aligned} \quad (10)$$

¹⁰In reality, individuals accrue defined benefit pension entitlements based on their number of years paying contributions into the scheme and some measure of average salary or final salary during those years of contributions. In a typical UK final-salary private sector DB scheme, for example, an individual would accrue entitlement to an additional one sixtieth of their final salary for each year of service (Cribb and Emmerson, Cribb and Emmerson). The model therefore abstracts from the dependence of DB income on idiosyncratic earnings realisations.

¹¹As will be set out in detail in Section 4, at ages 65 and older, the equalisation factor captures changes in household size due to the death of one member of a couple and this is assumed to reduce pension income.

Inheritances: Each household has a parent household. The parent household has probabilities of survival to each age given by $\{S_{\omega,t}^p\}_{t=24}^{t=110}$ and the variable $P_{i,t} \in \{0, 1\}$ denotes whether or not the parent household is alive at time t . When the parent household dies, the household receives an inheritance draw, $h_{i,t}$ from the type-specific distribution, F_{ω}^h .

Asset returns: Assets are allocated across several asset classes, namely gross housing wealth, risky assets, safe assets, physical wealth, cash, and mortgage debt. This allocation is exogenous and described below. At age t , the returns to each of these asset classes are denoted, respectively, by $r_{h,t}$, $r_{e,t}$, $r_{b,t}$, $r_{p,t}$, $r_{c,t}$ and $r_{mort,t}$. Returns are correlated across asset classes within period but are independently drawn over time.

End-of-life costs and bequests: When the household dies, they draw an end-of-life cost, κ_i from the distribution F_g^{κ} , which varies by generation. The bequest for a household that dies after the end of period t is the greater of their net wealth, less the end-of-life costs, and zero:

$$b(a_{t+1}) = \max\{0, a_{i,t+1} - \kappa_i\}. \quad (11)$$

3.3 Budget constraints

Gross income, $y_{i,t}$, is the sum of earnings (minus occupational pension contributions), pension income and inheritances:

$$y_{i,t} = e_{i,t}(1 - q_{i,t}) + p_{i,t} + h_{i,t}. \quad (12)$$

Each period, households decide how much of their gross income to save, denoted $z_{i,t}$, and consume their resulting after-tax income, according to the budget constraint:

$$c_{i,t} = \tau_t(y_{i,t} - z_{i,t}), \quad (13)$$

where $\tau_t(\cdot)$ is a function, varying with age, that captures the tax and welfare system. Start-of-period assets are denoted $a_{i,t}$. I define end-of-period assets as the sum of start-of-period assets and net saving:

$$\tilde{a}_{i,t} \equiv a_{i,t} + z_{i,t} \quad (14)$$

Assets at the start of next period, a_{t+1} are equal to end-of-current-period assets multiplied by a stochastic return, r_{t+1} :

$$a_{i,t+1} = \tilde{a}_{i,t}(1 + r_{t+1}(\tilde{a}_{i,t})) \quad (15)$$

The allocation of total assets to gross housing wealth, risky assets, safe assets, physical wealth and cash and less mortgage debt is given by a vector of portfolio shares, π_t^ω and a housing leverage ratio, lev_{t-1} . These portfolio shares are exogenous and vary by generation, age and the level of assets in a way that captures a typical portfolio for a household of that age, generation and wealth level. The realised rate of return for a household of type ω and who ended the previous period with assets $\tilde{a}_{i,t-1}$ is given by

$$r_t^\omega(\tilde{a}_{i,t-1}) = \pi_{h,t}^\omega(\tilde{a}_{i,t-1}) \frac{r_{h,t} - lev_{t-1} \cdot r_{mort,t}}{1 - lev_{t-1}} + \pi_{e,t}^\omega(\tilde{a}_{i,t-1}) \cdot r_{e,t} + \pi_{b,t}^\omega(\tilde{a}_{i,t-1}) \cdot r_{b,t} \\ + \pi_{p,t}^\omega(\tilde{a}_{i,t-1}) \cdot r_{p,t} + \pi_{c,t}^\omega(\tilde{a}_{i,t-1}) \cdot r_{c,t}. \quad (16)$$

The combination of the distribution of returns to each asset class and the exogenous portfolio allocation defines a distribution for the return to assets that varies by age, generation and level of end-of-period assets and is denoted $F_{\omega,t}^r(\tilde{a}_{i,t})$. This dependence of returns on the level of assets encapsulates the so called ‘scale dependence’ of returns.

Borrowing cannot exceed a type-specific limit, $\bar{a}_{\omega,t}$, equal to twice average type earnings up to age 74 and zero for ages 75 and older:

$$\bar{a}_{\omega,t} = 2\mathbb{E}_{24}[e_{i,t}|\omega] \quad \text{for } t < 75 \\ \bar{a}_{\omega,t} = 0 \quad \text{for } t \geq 75$$

Given that returns and income are uncertain, the borrowing limit for end-of-period assets in period t is defined recursively as:

$$BL_{\omega,T} = 0 \\ BL_{\omega,t} = \min \left\{ \frac{BL_{\omega,t+1} - y_{\omega,t}^{min}}{(1 + r_{t+1}^{max}(BL_{\omega,t}))}, \bar{a}_{\omega,t} \right\}$$

where y_t^{min} is the minimum possible income in period t and $r_{t+1}^{max}(a_t)$ is the maximum possible rate of return given end-of-period assets a_t as $r_{t+1}^{max}(a_t)$.¹² Households’ choices must therefore satisfy

$$\tilde{a}_{i,t} \geq -BL_{\omega_i,t}, \quad \forall t. \quad (17)$$

¹²As returns are a function of end-of-period assets, the limit for end-of-period borrowing depends on itself. When solving the model, the rate of return is the same for all negative assets levels meaning that the end-of-period borrowing limit is unique at each age and can be calculated recursively.

Definition of total private wealth: I define total wealth as the sum of net assets, $a_{i,t}$, and the expected discounted value of accrued occupational defined benefit pension income entitlements, $W_{i,t}^{db}$:

$$W_{i,t} = W_{i,t}^{db} + a_{i,t}. \quad (18)$$

The expected discounted value of accrued occupational defined benefit pension income entitlements is defined as

$$W_{i,t}^{db} = \sum_{\tau=t}^{110} \frac{\frac{S_{w,\tau}}{S_{w,t}} db_{\omega,t}(\zeta_i)}{(1 + r_\tau^d)} \cdot \frac{\sum_{\tau=1}^{t-1} \mathbb{E}_{24}[e_{i,t}|\omega_i]}{\sum_{\tau=1}^{K_\omega-1} \mathbb{E}_{24}[e_{i,t}|\omega_i]} \quad (19)$$

where r^d is a discount rate used to discount future flows of DB pension income.¹³ The way that defined benefit wealth is valued does not impact households' decisions but is relevant for taking the model to the data.

3.4 Household problem and value functions

I define 'cash-on-hand', denoted M_t , as the sum of start-of-period assets and income.¹⁴ The state variables of the model are the household's type, age, cash-on-hand, employment status, the persistent component of earnings and whether the parent household is alive. The household's lifecycle can be split into three sets of periods: working life, early retirement, and late retirement.

Late retirement periods: In periods after the household is retired and after the latest age of death of the parental household, there are no shocks to earnings or employment or parental survival. Suppressing i subscripts for simplicity, the household's problem in these periods is

¹³The first part of this formula is the expected discounted value of future DB pension income. The second part implies that pension entitlements are accrued based on the share of earnings lifetime that have, on average, been earned by a household of that type by that age. This is a way of capturing the fact that, in reality, defined benefit pension entitlement is commonly equal to either average career earnings or final pre-retirement earnings, multiplied by some accrual rate times the number of years of scheme membership.

¹⁴The use of 'cash-on-hand' to reduce the size of the state space is standard in lifecycle models. Pension income is a known function of household type and final periods earnings, which are state variables. Conditional on pension income, household decision rules do not depend on the division of the rest of cash-on-hand into inheritance and start-of-period assets. This means I do not need two separate state variables for start-of-period assets and inheritance received and can replace the sum of income and start-of-periods assets with the cash-on-hand variable.

given by

$$V_t(M_t; \omega, e_{K_\omega-1}) = \max_{c_t} \left\{ u(c_t; \theta_{\omega,t}) + \right. \\ \left. \beta(1 - s_{\omega,t+1}) \cdot \int \int \phi(b(\tilde{a}_t(1 + r_{t+1}))) dF_{\omega,t}^r(r_{t+1}|\tilde{a}_{i,t}) dF_g^\kappa(\kappa) + \right. \\ \left. \beta s_{\omega,t+1} \int V_{t+1}(\tilde{a}_t(1 + r_{t+1}) + p_{t+1}; \omega, e_{K_\omega-1}) dF_{\omega,t}^r(r_{t+1}|\tilde{a}_{i,t}) \right\}$$

subject to

$$M_t = a_t + y_{i,t}$$

and the intra and inter-temporal budget constraints, given by equations (12) to (17).

Early retirement periods: In the years from the latest retirement period onwards and while the parent household may be alive, the household problem for a household with a living parent household is given by

$$V_t(M_t, P_t = 1; \omega, e_{K_\omega-1}) = \max_{c_t} \left\{ u(c_t; \theta_{\omega,t}) + \right. \\ \left. \beta(1 - s_{\omega,t+1}) \cdot \int \int \phi(b(\tilde{a}_t(1 + r_{t+1}))) dF_{\omega,t}^r(r_{t+1}|\tilde{a}_{i,t}) dF_g^\kappa(\kappa) + \right. \\ \left. \beta s_{\omega,t+1} s_{\omega,t+1}^p \int V_{t+1}(\tilde{a}_t(1 + r_{t+1}) + p_{t+1}, P_{t+1} = 1; \omega, e_{K-1}) dF_{\omega,t}^r(r_{t+1}|\tilde{a}_{i,t}) + \right. \\ \left. \beta s_{\omega,t+1} (1 - s_{\omega,t+1}^p) \int \int V_{t+1}(\tilde{a}_t(1 + r_{t+1}) + p_{t+1} + h_{t+1}, P_{t+1} = 0; \omega, e_{K_\omega-1}) dF_{\omega,t}^r(r_{t+1}|\tilde{a}_{i,t}) dF_{\omega,t}^h(h_{t+1}) \right\}$$

subject to

$$M_t = a_t + y_{i,t}$$

and the intra and inter-temporal budget constraints, given by equations (12) to (17). The problem when the parent household no longer survives (i.e. $P_{i,t} = 0$) is an analogous but simplified version of this problem.

Working life periods: In the periods before period K_ω , the evolution of earnings and employment are uncertain.¹⁵ To ease notation, I define $X_t = \{M_t, E_t, \eta_t\}$. The household problem

¹⁵For most of working life, households survive to the next period with certainty. In these periods, the problem does not fundamentally change but $s_{\omega,t+1} = 1$ and the second term in the right-hand-side of the Bellman equation drops out.

for a household with a surviving parent household is given by

$$\begin{aligned}
V_t(X_t, P_t = 1; \omega) = & \max_{c_t} \{u(c_t; \theta_{\omega,t}) + \\
& \beta(1 - s_{\omega,t+1}) \cdot \int \int \phi(b(\tilde{a}_t(1 + r_{t+1}))) dF_{\omega,t}^r(r_{t+1}|\tilde{a}_{i,t}) dF_g^\kappa(\kappa) + \\
& \beta s_{\omega,t+1} s_{\omega,t+1}^p \int V_{t+1}(X_{t+1}, P_{t+1} = 1; \omega) dF(X_{t+1}, P_{t+1} = 1|X_t, P_t; \omega) + \\
& \beta s_{\omega,t+1} (1 - s_{\omega,t+1}^p) \int \int V_{t+1}(X_{t+1}, P_{t+1} = 0; \omega) dF(X_{t+1}, P_{t+1} = 0|X_t, P_t; \omega) \}
\end{aligned}$$

subject to

$$M_t = a_t + y_{i,t}$$

and the intra and inter-temporal budget constraints, given by equations (12) to (17), and laws of motion for employment and the stochastic component of earnings, given by equations 3 to 9. The problem for a household without a surviving parent is an analogous and simplified version of the above problem.

3.5 Model solution and simulation

There is no analytical solution to this maximisation problem. I solve the model recursively, using numerical methods to obtain the household decision rules. Full details of the solution method are given in Appendix A. When simulating the return to assets, I use observed realisations of returns experienced by each generation, as discussed further in Section 4.5. When calculating moments of the simulated distribution of wealth within each generation, I weight each household according to the probability of survival to that age for their household type.

4 Estimation of generations' economic and demographic circumstances

Each generation is characterised by a set of exogenous parameters or processes in the model, for each of the following 'circumstances': earnings process; tax and welfare system; public pension system; occupational pension provision; rate of return on wealth; survival; household size and composition; inheritances. Some of these circumstances also vary by some of the other elements of the household's type. For example, household survival varies by generation but also by education and earnings type within each generation.

I set out the data and estimation methods used for each circumstance in turn. In some cases, more detailed description is given in the Appendix B. I strive to define variables comparably across datasets: year of birth and age of couple households are always defined as their mean year of birth and mean age. The level of education of a couple is defined as the highest education level attained within the couple (individuals are either low-educated (up to compulsory schooling only), mid-educated (high school graduate), or high-educated (some higher education i.e. college education)).

The model embodies rational expectations on the part of households. From the perspective of unbiased estimation of the model’s preference parameters, I require that what I assume about the distribution of future circumstances is in line with households’ expectations.


4.1 Earnings process

Data: The ideal dataset for estimating generation-specific earnings processes would be a long-running household panel. In the UK, panel data on earnings is available from 1991 onwards (UK Household Longitudinal Study, or UKHLS) and cross-sectional data on earnings is available from 1968 onwards (the Family Expenditure Survey, or FES, and its successors). I use both surveys.

Estimation: I allow for heterogeneity across household types in as rich a way as is feasible given the data. I use the FES data to estimate the deterministic age-profile of earnings for each generation-and-education group. To address the collinearity of age, period and cohort, I restrict time effects to take the form of a linear trend plus annual deviations from trend that sum to zero (see Deaton and Paxson (1994)). I estimate the process for the household fixed effect and persistent stochastic component of earnings using the UKHLS panel data, following the flexible quantile regression-based estimation method of Arellano et al. (2017). I estimate transitions rates in and out of employment using the UKHLS. Employment rates and earnings risk have changed over time (Blundell et al. (2013); Blundell and Etheridge (2010)) before 1991. To account for this, I adjust the transition rates and the variance of the of the persistent earnings shocks to match the FES data for generations observed in years before 1991. I pursue the flexible quantile function approach because it has been shown that there is important age-dependency, nonlinearity and conditional skewness in earnings innovations and that this has quantitatively important implications for the level and distributions of wealth when incorporated into a lifecycle savings model (see De Nardi et al. (2019) and Arellano et al. (2017)). Adequately capturing these dynamics is therefore likely to be important for our estimated preference parameters and resulting conclusions.

Figure 6 shows the profiles for mean earnings for each generation obtained by simulating the estimated earnings processes. Mean earnings is simulated to grow, generation-on-generation, but at a diminishing rate. Whereas average lifetime earnings for those born in the 1940s are 26% higher than for those in the 1930s, and the 1950s are 17% higher than for the 1940s, the subsequent generation-on-generation growth rates are 5%, 7% and 6% respectively. Differences in the shape of the age profile of earnings (i.e. change in earnings with age) across generations arise due to differences in employment rates by age and differences in earnings growth conditional on employment. I capture the fact that employment rates at older ages are higher for later-born generations and that the shift towards higher levels of education means that the growth rate of earnings with age is higher for later-born generations. Between the age of 24 and 50, average earnings are simulated to grow by 66% for the 1930s-born generation, 116% for the 1950s-born and 130% for the 1980s-born.

Figure 6: Predicted mean gross annual household earnings, by age and generation



Note: Predicted mean earnings are the mean of 180,000 simulations of the earnings process for each generation.

4.2 Taxes and welfare system

The UK tax system operates primarily at the level of the individual so an exact tax and welfare calculator cannot be applied to household earnings. I therefore specify the tax and welfare system as the following function of pretax earnings:

$$\tau_t(y_{i,t} - z_{i,t}) = \psi_t^a + \psi_t^b(y_{i,t} - z_{i,t})^{\psi_t^c} \quad (20)$$

where the parameters ψ_t^a , ψ_t^b and ψ_t^c define the tax and benefits system in year t . This form follows Heathcote et al. (2014), with the addition of a constant term. Loosely, ψ_t^a pins down the consumption floor set by the welfare system, ψ_t^b pins down the level of taxation and ψ_t^c determines the progressivity of the tax system. These parameters are allowed to vary by year.

Data and estimation: Using the FES data, I define household gross annual earnings and household net annual earnings. For each year, I estimate the tax and welfare parameters using nonlinear least squares. The mean R-squared for the estimated tax functions is 0.89. To put these functions into the model, I select, for each generation and year of age, the tax function for the middle birth year for that generation at that age e.g. for the 1940s generation, I assume that at age 30 they experience the 1975 tax system.

4.3 Public pension system

Estimation: I model a simplified version of the state pension system whereby state pension income is an affine function of final earnings that varies by household type. To estimate the parameters of this simplified state pension system, I simulate the household earnings process for each generation a large number of times and then, for each simulation, use a full state pension calculator (taken from Banks and Emmerson (2018)) to calculate the state pension entitlements that would accrue to each household given their earnings and employment history. I then estimate an OLS regression of household state pension entitlement on generation-education-and-earnings-group-specific indicators and the interaction of generation-education-and-earnings-group-specific indicators with household earnings in the year before retirement.

Table 1 shows the estimated pension system parameters for households in the mid-earning group of each cohort- and education-group. I express the intercept term as a percentage of average working-life earnings for the generation, to abstract from changes in the real-terms generosity of the pension system across generations that reflect state pensions keeping pace with rising average earnings, and so that percentages are comparable across education groups in each generation.

Table 1: Public pension system parameters for mid-earning households, by generation- and education-group

Generation	Constant (% of mean earnings)			Earnings-related component (%)		
	Low-ed	Mid-ed	High-ed	Low-ed	Mid-ed	High-ed
1930s	32.2	40.8	48.3	16.0	7.4	4.9
1940s	31.6	41.2	46.7	11.2	4.2	4.5
1950s	28.3	37.1	42.0	8.7	1.8	1.1
1960s	26.1	31.6	33.1	5.4	0.4	0.5
1970s	26.8	29.7	29.2	7.3	0.9	0.9
1980s	28.4	30.8	31.9	7.0	1.4	0.0

Note: Table shows the two parameters that define state pension entitlements as a function of period $K - 1$ household earnings for the mid-earning group in each generation-and-education group. The constant component is expressed as a percentage of that generation’s mean earnings so that the figures are comparable across education groups within each generation.

The average generosity of the state pension system is highest for those born in the 1930s and 1940s, and declines over the three subsequent generations, who experienced less generous earnings-related pension systems, before rising slightly for 1980s generation due to an increase in the generosity of the flat rate component of the pension from 2016 onwards. The degree to which pension entitlements are related to earnings, as measured by the difference in α between education groups, as well as the size of β , becomes weaker moving across generations.

4.4 Occupational pension provision

Data and estimation: To estimate the DB pension function and the proportion of households with DB pension income, I estimate household DB pension income and DB membership using the English Longitudinal Study of Aging (ELSA), a household panel survey representative of the English population aged 50 and older. To estimate the prevalence of DB pensions for younger generations and DC pensions for all generations, I use estimates based on the Annual Survey of Hours and Earnings (ASHE), a representative sample of UK employees’ pay and pension provisions.

Table 2 gives the proportion of households in each generation and education group that I model as having DB pension entitlements, and the value of DB income as a percentage of mean earnings for that group. DB prevalence is steady over the 1930s to 1950s generations but declines thereafter. Estimated employer DC contributions for each generation (averaging over all of working life) rise to 2.5% of earnings for the 1980s-born generation. This increase

reflects the introduction of automatic enrolment into workplace pensions in 2012.

Table 2: Occupational pension system parameters by generation- and education-group

Generation	% with DB pension			DB income as % of mean earnings			Employer DC contribution %
	Low-ed	Mid-ed	High-ed	Low-ed	Mid-ed	High-ed	
1930s	30.8	48.6	68.3	33.7	36.9	47.6	0.3
1940s	28.1	48.4	68.1	34.8	35.4	43.1	0.4
1950s	22.1	48.6	67.4	31.9	36.5	40.4	0.9
1960s	20.0	37.3	56.1	31.9	36.5	40.4	1.6
1970s	20.0	35.0	42.8	31.9	36.5	40.4	2.1
1980s	20.0	35.0	40.0	31.9	36.5	40.4	2.5

Note: Table shows the proportion of households that are estimated to have at least one occupational DB pension and the value of the DB pension as a percentage of average working-life earnings for that group. Those not part of an occupational DB scheme receive employer DC contributions as an additional to their gross earnings, at the rate reported in the final column.

4.5 Rates of return on assets

Data: Information on the returns to asset classes by year are taken from Jorda et al. (2019). Mortgage rates are the effective mortgage rate series taken from the Bank of England’s “A Millenium of Economic Data.” Future rate of return assumptions are based on forecasts from Oxford Economics, the Financial Conduct Authority and the UK Office for Budget Responsibility. All returns are after inflation. Returns to housing include rental yield, capital gain and depreciation, returns to equities include capital gains and dividends, and returns to bonds include capital gains and interest income. The real return to physical wealth is assumed to be zero. The share of assets held in different asset classes is estimated using data from WAS.

Estimation: I estimate portfolio shares held in the following asset classes: net housing wealth, risky financial assets, safe financial assets, cash and physical assets. I use data from waves 1 to 5 of WAS on households’ portfolio shares in each asset class (assigning DC pension assets across these asset categories). I calculate the portfolio share held in each asset class in a number of net assets bins, for both households with and without a DB pension.¹⁶ To extrapolate portfolio shares to years outside of the WAS data period, I assume that the share of wealth held as housing depends not on the real level of wealth held but on the level of


¹⁶I tested whether, conditional on wealth level, there were differences in portfolio composition by education group or generation but found no significant differences.

wealth as a fraction of the average house price in that year.¹⁷ I use estimates for housing leverage as a function of age from Crawford and O’Dea (2020), estimated using WAS. The realised rate of return experienced by a household of type ω at age t is given by the product of the portfolio shares in each asset for a household of their type at that age (π_t^ω) and the realised return to each asset type in the relevant years, taking a geometric mean of the 10 returns at each age (corresponding to the 10 birth years) in each generation. I assume that households expect returns to be drawn from a distribution that reflects the long-term trends in returns and the historic distribution of returns around that trend (further details are given in Appendix B).

The overall effect of the trends in returns and the changes in portfolio composition across generations is illustrated in Figure 7. This figure shows the modelled levels of expected returns to assets as a function of the wealth, at age 30, for each generation. I express wealth as a multiple of that generation’s average earnings at age 30. This illustrates that younger generations face lower expected returns at given levels of wealth to earnings ratios, compared to older generations. For each generation, returns increase with wealth because households with more wealth tend to hold more of their wealth as housing, and, at higher levels of wealth, as risky financial assets. However, for younger generations, the increase in returns with wealth is less rapid at lower wealth levels due to the fact that they must accumulate more wealth (as compared to their earnings) before the housing share of their portfolio significantly increases.

¹⁷Average house prices by year are taken from Nationwide’s house price index.

Figure 7: Expected return to assets at age 30, by assets as multiple of generation's average earnings at age 30



Note: Each series shows the households' expected return to assets at age 30 as a function of assets (expressed as a ratio to that generation's average annual earnings at age 30).

4.6 Survival probabilities

Data: I use cohort life tables from the Office of National Statistics and ELSA, which is linked to administrative death records.¹⁸

Estimation: I estimate education-specific survival curves for each combination of education level, earnings type and sex, for each sex and year of birth. Separately for males and females, I take the ELSA data and estimate a Cox proportional hazard model where the explanatory variables are education level and within-generation-and-education-group household wealth tertile. I use wealth tertile as our measure of the individual's household earnings group as many older households are not in work and wealth rank is a good proxy for lifetime earnings rank. I apply the estimated hazard ratios to the official life tables to obtain male and female survival curves for each generation, education and earnings type. I use these to create household survival curves by household type by using FES data on the distribution of households across single male, single female and couple types (when aged 25 to 60) and assuming survival realisations are drawn independently across members of couples.

¹⁸I use year-of-birth-and-sex-specific life tables for England. Cohort-based survival curves for the whole UK are not available for all of the generations I examine, hence I use England and Wales only. However, England and Wales make up almost 90% of the UK population so these likely closely resemble survival curves for the whole UK.

4.7 Household size and composition

Data: I use data from the FES, which contains information on the number and age of individuals in each household.

Estimation: Using the FES I calculate the mean OECD-modified equivalisation factor by generation and education-group, for ages 24 to 65.¹⁹ For generations that are not observed through all of adult life, I extrapolate using the growth rate of the equivalisation factor with age from earlier or later generations. At ages 65 and older, I take the distribution of equivalisation factors at age 65, and assume that these change consistent with the rate of transition of couple households to single households implied by the estimated survival probabilities.

4.8 Inheritances

Data: I use WAS, which has data on lifetime inheritances received both before and during the survey period. The first wave asked individuals whether they expected to receive an inheritance and the amount that they expected to receive if so, selecting from 7 possible ranges. The survival curves for parental households draw on the same data used to estimate the household survival curves.

Estimation: I take the maximum inheritance reported as received or expected within each household and discretise the distribution of inheritances into eight groups, including zero. For each generation, education and earnings type, I calculate the distribution of responses across these eight categories. For the 1930s and 1940s generations, for which the majority did not receive inheritances, the distribution varies by education level but not earnings type. Appendix Figure B.3 shows the proportion of individuals who are expected to inherit more than given amounts, by generation.

4.9 End-of-life care costs

In England, retirees with assets above a low threshold are responsible for paying all costs of assistance with living at home and costs of living in a residential or nursing care home.²⁰

¹⁹Under the OECD modified scale, a single adult has an equivalisation factor of 1 and each additional adult or child aged 14 or older adds an additional 0.5 to this and a child 13 or under adds an additional 0.3 to this equivalisation factor.

²⁰Different systems have operated in Scotland and Wales. In 2021, the UK government announced a cap on lifetime contributions to social care costs but this was then delayed indefinitely. This is after our data period and so I treat households as having expected the prior system to continue.

Data on end-of-life costs in the UK is limited and does not allow us to estimate how end-of-life care costs have changed across generations. I use an estimated distribution of costs for current older people and assume that other generations experience the same distribution of costs relative to their annual earnings. This assumption may be thought reasonable given that labour costs are the main component of end-of-life care costs.

Data: I use the distribution of individuals' older age social care costs, estimated for the 2011 Commission on Funding of Care and Support. This is estimated for those aged 65 in 2010-11.²¹ While these estimates are for total costs over all of older age, the vast majority of care spending comes in the final months and years of life. Forder and Fernandez (2011) estimated that half of UK care home residents had a stay of 1 year and 7 months or less, while only a quarter had stays of 3 years and 7 months or more.

Estimation: I discretise this distribution of social care costs into seven bins. This distribution pertains to the 1960s generation. To obtain a distribution for other generations, I grow the cost amount in each bin by the ratio of the generation's average annual earnings to the average annual earnings of the 1960s generation.²²

5 Estimation of lifecycle model preference parameters

The lifecycle model has four preference parameters: the coefficient of relative risk aversion, the discount factor and the two bequest motive parameters. Here, I describe the method of estimation of the preference parameters when these are assumed to be common across the 1930s, 1940s and 1950s generations. Estimation under alternative assumptions is analogous, with differences described in the following section.

I fix the coefficient of relative risk aversion at 3, following Crawford and O'Dea (2020) and Scholz et al. (2006). I denote the vector of the remaining three preference parameters as $\Delta = (\beta, \phi_1, \phi_2)$ and estimate these parameters jointly using the method of simulated moments. In line with the literature (see, for example, French and Jones (2004) De Nardi et al. (2010) and De Nardi et al. (2021)), I take the vector of other estimated model parameters, denoted $\hat{\chi}$, as given in the estimation. This vector includes the estimated 'circumstances' set out in the previous section. I denote the true preference parameter vector Δ_0 and estimate it by finding

²¹Figure 1 of the January 2022 Social Care Charging Reform Impact Assessment presents this distribution in 2020-21 prices: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/10449-social-care-charging-reform-impact-assessment.pdf.

²²The distribution for the 1960s generation is as follows: 27% of households are expected to have zero social care costs, 23% to spend £11,000, 18% to spend £36,000, 10% to spend £67,500, 8% to spend £87,500, 8% to spend £125,000, 6% to spend £225,000.

the vector, $\hat{\Delta}$, that minimises the distance between certain model-generated moments and their data analogues, given $\hat{\chi}$.

The moments that I use for estimation are the 25th percentile, 50th percentile and 75th percentile of total net wealth at each age for the 1930s, 1940s and 1950s generations. All moments are informative about all elements of the parameter vector. However, intuitively, the discount factor is identified by the levels of median wealth, the strength of the bequest motive, ϕ_1 , is identified by the speed at which wealth is drawn down at older ages, and the extent to which bequests are a luxury, ϕ_2 , is identified by the differences in the degree of wealth inequality and how it evolves at older ages.

Let $W_{g,t}^q(\Delta, \chi)$ be the model predicted q^{th} quantile of total net wealth for generation g at age t and π_q the proportion of simulated households with wealth below this level i.e. for the median $\pi_q = 0.5$. Let $W_{i,t}$ denote household i 's total net wealth in time t as measured in the WAS data. Assuming that observed assets have a continuous density at the true parameter vector (Δ_0, χ_0) , we obtain moment conditions of the form

$$E[\mathbb{1}\{W_{i,t} \leq W_{g,t}^q(\Delta_0, \chi_0)\} - \pi_q | i \text{ is alive at } t \text{ and in generation } g] = 0 \quad (21)$$

for every age t , generation g and for $q \in \{25, 50, 75\}$. Let $\hat{G}(\Delta, \hat{\chi})$ be the vector formed by replacing $W_{g,t}^q(\Delta_0, \chi_0)$ with $W_{g,t}^q(\Delta, \hat{\chi})$ in Eq. (21) and stacking across all ages and generations and estimation quantiles. The estimated parameter vector is given by minimising the GMM criterion function as follows:

$$\hat{\Delta} = \arg \min_{\Delta} \hat{G}(\Delta, \hat{\chi})' \hat{\mathbf{W}} \hat{G}(\Delta, \hat{\chi}), \quad (22)$$

where $\hat{\mathbf{W}}$ is some square weighting matrix. I use a weighting matrix whose diagonal elements are equal to the the diagonal elements of the inverse of the variance-covariance matrix of the data used in estimation, with off-diagonal elements equal to zero. The asymptotic distribution of parameter estimates and model overidentification tests are discussed in Appendix C.

I obtain $\hat{\Delta}$ by finding the parameter vector that minimises the GMM criterion function. I start by solving and simulating the model a large number of times for a set of parameter vectors randomly drawn using a Sobel sequence. I construct the model-predicted quantiles of the wealth distribution and calculate the GMM criterion for each parameter vector. Starting from the parameter vector from this set with the smallest criterion value, I then search over the parameter space using the Nelder-Meade simplex algorithm.

Parameter estimates: Table 3 shows the estimated preference parameters. I show a set of estimates obtained by fitting to only the median level of wealth at each age for each generation and a set of estimates obtained by fitting to the 25th percentile, median, and 75th percentile of wealth at each age for each generation. The estimated level of patience is close to 1 in both cases. The bequest motive parameters are often described in terms of their implications for the share of wealth left as a bequest, rather than consumed, in the terminal period when death next period is certain. The sets of parameters imply similar marginal propensities to bequeath out of final-period wealth of 95% and 96%, respectively, for those making bequests. This is similar to other estimates in the literature (see De Nardi et al. (2021) for a discussion). Positive final-period bequests are made by those with at least £6,517 in wealth and at least £17,697 in wealth, under the two sets of estimates, respectively.²³ Appendix Figure D.1 shows that both sets of parameters give a close fit to the 25th, median and 75th percentiles of wealth at each age for each generation, and a good fit to the generation-on-generation growth in median consumption, which is untargted in estimation. In the following analysis, I use the estimates based on fitting median wealth only.

Table 3: Preference parameters estimated when fitting data for 1930s, 1940s and 1950s-born generations and assuming common preferences across generations

	Fitting 50th percentile (1)	Fitting 25th, 50th, and 75th percentile (2)
$\hat{\beta}$: patience	0.999 (0.0007)	1.008 (0.0004)
$\hat{\phi}_1$: bequest intensity	7,156 (271)	14,642 (373)
$\hat{\phi}_2$: bequest curvature	160,699 (11,638)	436,394 (10,833)

Note: Standard errors are in parentheses. The coefficient of relative risk aversion is set to 3 in both specifications. The bequest curvature parameter is expressed in pounds in 2018 prices.


²³For comparability with the literature, I calculate bequests as a function of cash-on-hand for an individual who faces death next period with certainty, faces no end-of-life care costs and will receive a rate of return on wealth of 2.3% with certainty.

6 Results

6.1 Changing circumstances or changing levels of patience?

I first show results from two ‘exercises’ that assess whether changing circumstances or changing preferences explain why generational wealth growth was strong but has stopped for younger generations. First, I simulate the model for all six generations, varying the generation-specific circumstances with each generation but holding preferences fixed at the levels estimated by fitting to the wealth data for the 1930s, 1940s and 1950s-born generations. Figure 8 shows that the model generates the large increases in median wealth across the older generations observed in the data. The model is also able to replicate the similar levels of median wealth, compared to the preceding generation, for the 1960s generation onwards. Therefore, changing circumstances are able to generate the qualitative pattern in the data, within a model in which preferences are unchanging across generations. Importantly, as shown in Appendix Figure F.1 (a), the model generates large differences in wealth between the three older generations during working life, meaning that the lack of generation-on-generation growth in wealth for younger generations is not a necessary consequence of the model structure.

Figure 8: Comparison of median wealth in WAS data and model simulations under fixed preferences



Note: Model simulations use the same estimated level of patience (0.999) for each generation, combined with the generation-specific model inputs capturing economic and demographic circumstances. Source: WAS waves 1 to 6 and model simulations.

While closely fitting the wealth patterns seen in the data, the model fit for younger generations in Figure 8 is not exact. In a second exercise I therefore re-estimate the model, allowing patience to vary across generations and assess whether these levels of patience, and the simulated levels of wealth that they imply, are substantively different from each other and from the estimate reported in Table 3. In this estimation, I fix the bequest motive parameters at the levels reported in Table 3. I do this for two reasons. First, the bequest motive parameters are identified primarily by the wealth patterns of the generations who are observed at older ages. Bequest motive parameters that vary across all generations cannot be plausibly separately identified from varying patience, given that younger generations have not yet been observed into older age. Second, I want to make the strongest test of the hypothesis that there are changing levels of patience across generations. If I do not find strong evidence of changing levels of patience even when loading all variation in wealth that is not explained by changing circumstances onto patience, then we can be confident that the data gives us no reason to think that patience has changed across generations.


Table 4 shows the estimated levels of patience for each generation, and their associated confidence intervals. The table also reports, for each generation, the simulated levels of wealth at age 65 (which is close to peak wealth for each generation).²⁴ When allowing patience to vary by generation, I estimate that patience is stable across the three older generations, at 0.999, and declining across the younger three generations. The 1980s-born generation has the lowest estimated level of patience of 0.992. One way of interpreting the magnitude of this decline in patience is to look at the implications for wealth. The table shows that the difference in simulated wealth at age 65 under the generation-specific level of patience, relative to wealth when assuming a common level of patience across older generations. Figure 9 shows the full age-profiles of median wealth under the model simulated using the generation-specific estimates and that using the common level of patience, for the 1960s, 1970s and 1980-born generations. For the 1980s-born, the difference between the two wealth profiles is largest: wealth at age 65 is 7.3% lower under the generation-specific rather than common level of patience. This is substantive. It is nevertheless much smaller than the increase in simulated wealth across generations from the 1930s-born to 1950s-born. Median wealth at age 65 is simulated to be 51% higher for the 1940s-born than the 1930s-born and 30% higher for the 1950s-born than the 1940s-born (implying a 97% rise from the 1930s-born to 1950s-born). The lower patience on the part of younger generations could therefore not be the primary explanation for generational increases in wealth of this size having stopped.

²⁴The full median wealth profiles for each generation are shown in Appendix Figure F.1.

Table 4: Comparison of discount factors and simulated wealth at age 65, under common and varying discount factor

Generation	$\hat{\beta}_g$, [95% CI]	Wealth at age 65 (£,000)		% difference: common vs varying β	% difference: vs prior generation	
		β common	β varies		β common	β varies
1930s	0.999 [0.997,1.001]	307	306	-0.3	-	-
1940s	0.999 [0.997,1.001]	461	461	0.0	50.2	50.8
1950s	0.999 [0.997,1.002]	602	602	-0.1	30.5	30.4
1960s	0.998 [0.995,1.000]	621	610	-1.8	3.1	1.3
1970s	0.994 [0.990,0.997]	683	648	-5.0	10.0	6.4
1980s	0.992 [0.985,0.998]	661	613	-7.2	-3.3	-5.4

Figure 9: Comparison of median wealth under common and generation-specific discount factors



6.2 Which changing circumstances matter?

I have shown that changing circumstances can largely explain the wealth patterns that we observe. This raises the question of what changes in circumstances generated large increases in wealth across generations from the 1930s to the 1950s-born and what has changed such that this has stopped for later generations. To answer this question, I use the model to isolate the contribution of changes in each economic and demographic circumstance to generational wealth differences. First, I show how the simulated level of median wealth changes when

I vary each circumstance in isolation, holding all others constant. In Figure 10 I show, for each circumstance, how the simulated median level of wealth varies when changing the inputs to those for the 1930s, 1950s and 1980-born generations, holding the inputs for all other circumstances constant at the level for the 1950s-born generation.

Second, I use these counterfactual simulations to decompose the percentage difference in median wealth at age 65 between the 1950s-born generation and to the 1930s-born generation into parts attributable to each economic and demographic circumstance, a part attributable to the change in estimated patience across generations, and a part attributable to the interaction between factors. I choose age 65 as this is the oldest age reached by the 1950s-born generation in the data and close to peak wealth. In the same way, I decompose the difference in wealth between the 1980s-born and 1950s-born generation at age 65. These two decompositions are shown in Figure 11. I also show the decomposition of the change between these two pairs of generations at age 33, the oldest age reached by the 1980s-born generation in the data. These two decompositions are shown in Figure 12. I compare only three generations to keep the comparisons clear and because most circumstances have changed smoothly across generations, meaning that the profiles for intervening generations lie between those shown. In Appendix E, I set out precisely how these counterfactual simulations are constructed.

The key findings of this analysis are as follows. The circumstances where changes have been most consequential are the level and profile of earnings and the returns to assets. Changes to taxes and household composition are also important. In panel (a) of Figure 10, I show the results of counterfactuals where I increase and decrease the distribution of earnings in proportion to the difference in average annual household gross earnings across generations.²⁵ Average annual household working-life earnings were around 50% higher for the 1950s-born than the 1930s-born generation and this feeds through to simulated wealth levels. The growth rate of average annual household earnings between the 1950s and 1980s generation is estimated to be lower, at around 20%, and consequently the wealth profiles for those generations are more similar.

Panel (b) of Figure 10 shows the effects of the varying age profile of earnings (i.e. change in earnings with age) across these three generations. Differences in the age profile of earnings arise due to differences in employment rates by age and differences in earnings growth conditional on employment. Younger generations both enter and exit the labour force later in life than their predecessors. Their higher levels of education also mean they see faster earnings growth with age, conditional on employment. Consequently, wealth accumulation

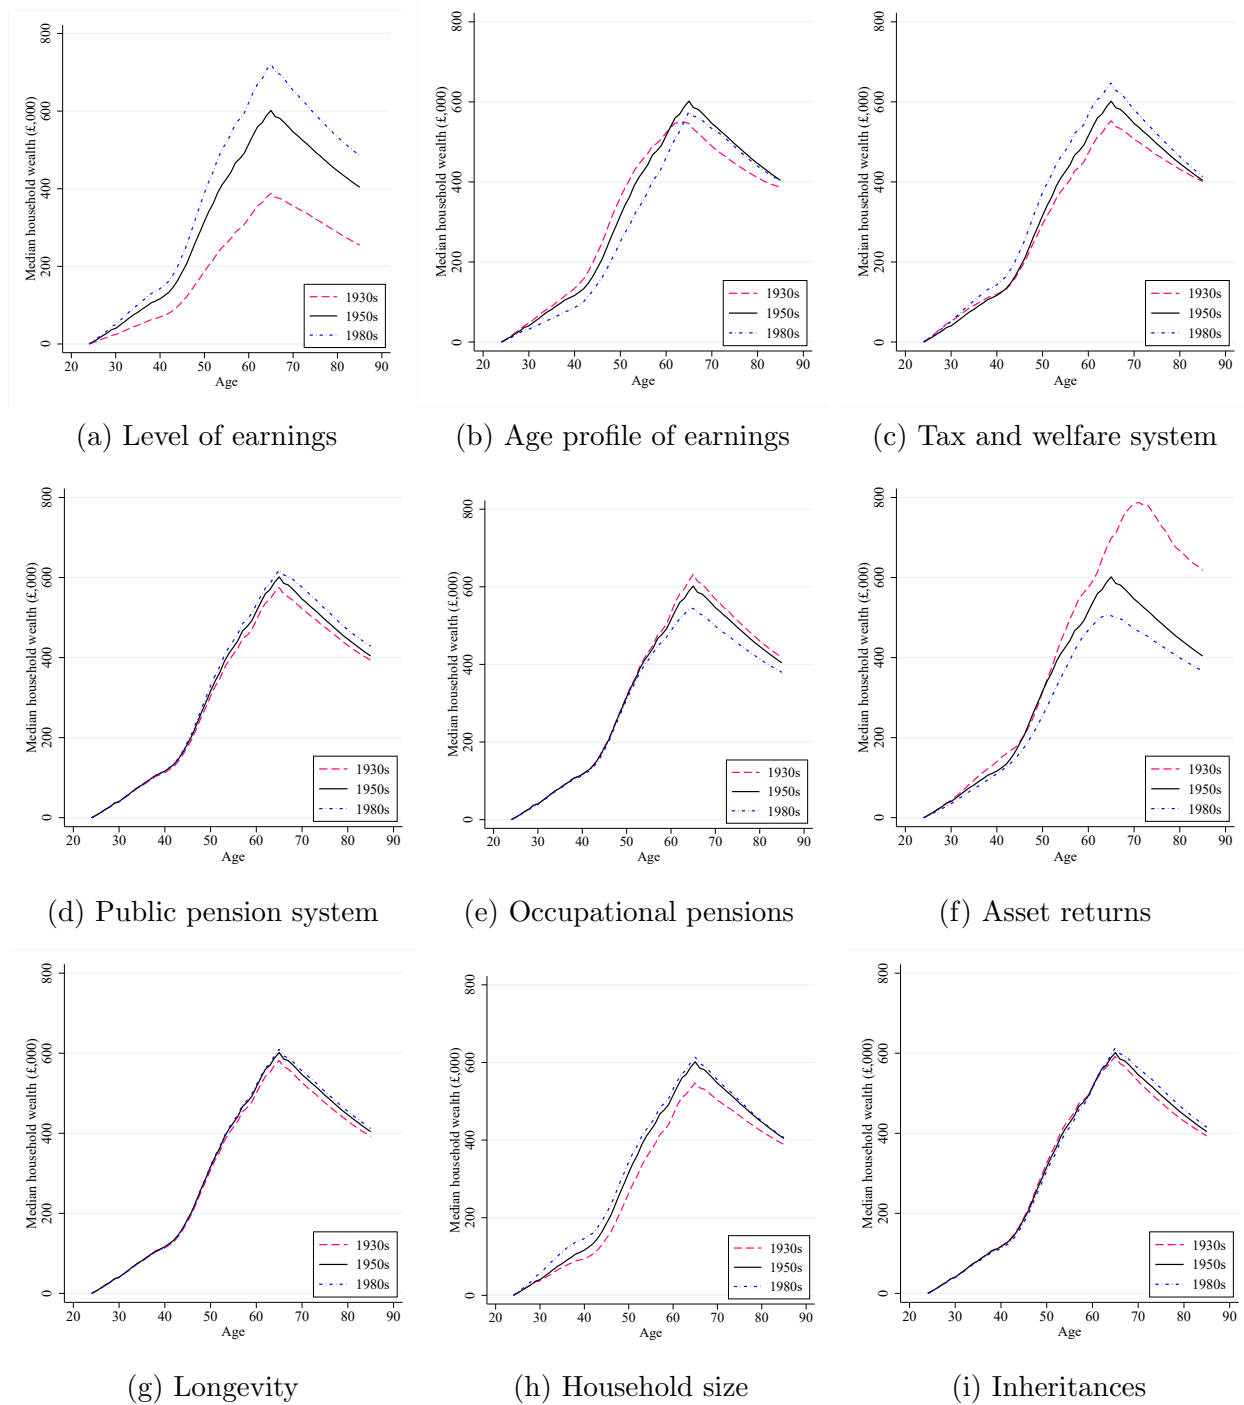
²⁵For the each generation, I calculate average annual earnings over ages 24 to 64. If generation X has average annual earnings that are on average Y% higher than the 1950s-born generation, I increase each point on the distribution of possible earnings levels for the 1950s generation by Y% to yield the counterfactual for generation X.

is pushed back into later life. Simulated wealth for the 1950s-born is 11% lower at age 33 but 10% higher at age 65, as compared to the 1930s-born. Comparing the 1980s-born to the 1950s-born, we see substantially slower wealth accumulation in early working life for the 1980s-born (27% less by age 33). However, the gap is closed by the end of working life. Changes in earnings risk across generations have negligible effects on wealth at retirement but do lead increased precautionary saving for younger generations in early life, with wealth 4% higher at age 33 due to the increase in risk between 1950s and 1980s generations.²⁶

Differences in experienced rates of return affect both the level and the profile of wealth accumulation. The biggest impact of the different rates of return experienced by different generations comes after age 50, when comparing the 1930s-born and 1950s-born generations' returns. While in their 50s and 60s, the 1930s-born generation approached the peak of their wealth and also experienced the returns to housing in the 1990s and 2000s and the high returns to equities in the 1980s. By contrast, the 1950s generation experienced these 'good' years at an earlier point in their lifecycle and experienced the period of lower returns following the financial crisis when in their 50s. The expectation that returns to most assets will be lower going forward than they were before the financial crisis means that median simulated wealth based on the 1980s generation's expected returns is 16% lower at age 65 than under the 1950s generation's returns.

²⁶For brevity, earnings risk is not shown in Figure 10 but is included in Figure 11 and Figure 12.

Figure 10: Effect on simulated median household wealth of varying each circumstance in isolation



Note: Each figure shows the effect of varying individual circumstances to their levels for the 1930s, 1950s and 1980s-born generations while all other circumstances are set to those for the 1950s-born generation.

Changes to the timing of partnering and having children, as well as changes in the number of children in each household, have significant implications for the age-profile of wealth accumulation. Those born in the 1930s had significantly larger households than those born in the 1950s, until they were in their late 40s. The estimated model predicts that this would lead to a greater amount of wealth accumulation for the later generation and, in particular, a greater proportion of wealth accumulation happening before the age of 50, with 14% more wealth held at age 33 and 10% more at retirement. Those born in the 1980s had much smaller households in their 20s and early 30s than did those born in the 1950s. Consequently, simulated median wealth is 37% higher at age 33. This difference reduces to 2% by age 65, indicating that much of the effect on wealth during working life is a retiming of saving.

A number of other factors have substantial, but smaller, impacts on generational wealth differences. Firstly, four factors contribute to slightly higher levels of wealth at age 65 for the later-born generations. In decreasing order of quantitative importance: The decline in rates of personal taxation has the effect of increasing lifetime disposable income and allows later-born generations to accumulate more wealth throughout working life. The simulated median level of wealth at age 65 is 9% and 8% higher for the 1950s-born compared to the 1930s-born and for the 1980s-born compared to the 1950s-born, respectively, as a result of the changes in taxation faced across generations.

Greater longevity leads to a slightly larger amount of wealth being held into retirement and it being drawn down more slowly through later age. The expansion of survival at older ages is simulated to have a slightly larger effect on wealth when comparing the 1950s-born to the 1930s-born generation (around 3% increase in wealth at age 65) than when comparing the 1980s-born to the 1950s-born (1% more wealth at age 65), in line with the slowing of generational increases in life expectancy.

The decline in the generosity of the state pension means that those born in the 1950s would be expected to save 5% slightly more for their retirement than those born in the 1930s. The change in the state pension system faced between the 1950s-born and 1980s-born generations is expected to lead to 3% more retirement wealth.

Inherited wealth grows in significance across generations. Those born later save slightly less during working life in anticipation of a larger future inheritances, but the uncertainty over the inheritance amount means this effect is small (a 3% reduction in wealth age age 33 when comparing to the 1980s-born to the 1950s-born). At age 65, the larger inheritances received by the 1950s generation, compared to the 1930s generation, leads to 2% more wealth being held, with an equivalent increase between the 1950s 1980s generations. While each of these components may be viewed as modest in size, together they would be expected to lead

to an 18% increase in wealth between the 1930s-born and 1950s-born at age 65 and a 14% increase in wealth between the 1950s-born and 1980s-born generations.

In addition to the negative effect of changes to rates of return, the other factor with a significant negative impact on generational wealth growth is changes to occupational pension provision. The decline in generous occupational DB pensions and their replacement with employer DC contributions is effectively a reduction in lifetime remuneration for those born later and so leads to lower wealth. This effect is substantive, being estimated to lead to a 5% reduction in wealth at age 65 when comparing the 1950s-born to the 1930s-born and a 9% reduction when comparing the 1980s-born and to the 1950s-born.

I show the impact of changing the assumed level of patience between its levels estimated for the 1930s, 1950s and 1980-born generations to illustrate its modest quantitative importance alongside the change in economic and demographic circumstances.

6.3 Implications of changing circumstances for consumption and welfare


There are several reasons why changing circumstances can have a different effect on equivalised consumption and welfare than on wealth. First, differences in household size or longevity mean that a given level of wealth can purchase different levels of annual equivalised consumption. Second, differences in returns to wealth and employer pension contributions imply differences in the level of saving out of earned income (and therefore forgone consumption) that is needed to achieve a given wealth level. Third, differences in net government transfers can drive differences in the relationship between private wealth and after-tax financial resources. Finally, changes in the timing of the receipt of income – holding consumption plans constant – or changes in the chosen timing of consumption – holding income streams constant – can mean that the timing of wealth accumulation changes without lifetime resources and lifetime consumption changing.

I use the model to show the change in equivalised consumption that results from each change in circumstances (shown in Appendix Figure F.4).²⁷ I find that while median wealth at age 65 is simulated to be 97% higher for the 1950s-born than the 1930s-born, simulated median equivalised consumption increases by 40% on average across these generations.²⁸ For the 1980s-born, median wealth is expected to be 2% higher at age 65 than for the 1950s-born


²⁷I verify that the model generates similar growth rates of consumption across generations as are observed in the data. Appendix D Figure D.2 shows the growth rates of median consumption between adjacent generations as simulated by the estimated model and as observed in the data, at the ages where data is available for each pair of generations.

²⁸The growth rate of consumption is based on taking the mean of annual equivalised consumption from the beginning of life until female life expectancy and comparing this value across generations.

Figure 11: Decomposition of generational differences in wealth



(a) Comparison of 1930s and 1950s-born at age 65




(b) Comparison of 1950s and 1980s-born at age 65


but median equivalised consumption is 14% higher over working life. This implies that we should expect generational increases in living standards to slow, but not to the extent seen for wealth. One major reason for this is that while the steepening of the earnings profile is a driver of generational wealth growth across the 1930s to 1950s-born and leads to lower wealth accumulation in early life for the 1980s-born, its effects on median equivalised consumption are more muted.

To assess the welfare impact of changes in circumstances, I calculate the compensating variation for each change in circumstances across pairs of generations. Define $\tilde{V}_{i,0}(a_0, \chi)$ as the value of household i 's lifetime utility – given *realised* values of shocks – as a function of initial assets, a_0 and circumstances χ . I define the compensating variation for a change

Figure 12: Decomposition of generational differences in wealth



(a) Comparison of 1930s and 1950s-born at age 33



(b) Comparison of 1950s and 1980s-born at age 33

in circumstances for generation g as the level of additional start-of-life assets, a_0^* , which, if held each member of generation g , would leave *ex-post* mean lifetime discounted utility unchanged when experiencing the change in circumstances i.e. for a change in circumstances from χ to χ' , a_0^* solves

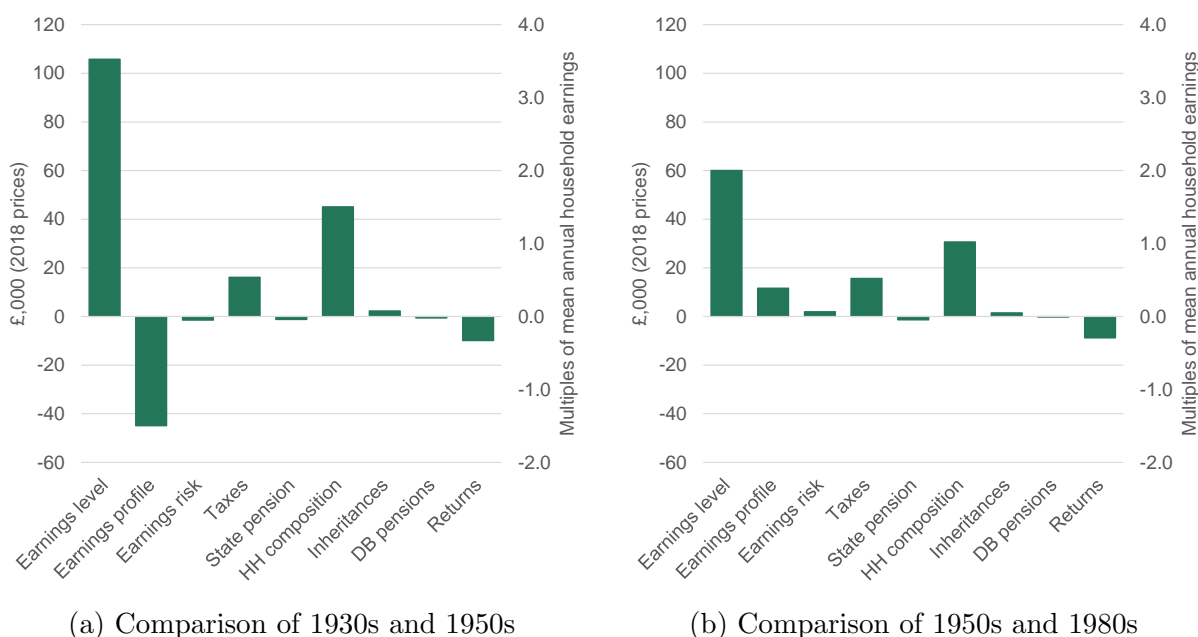
$$\int_{i \in g} \tilde{V}_{i,0}(a_0, \chi) di = \int_{i \in g} \tilde{V}_{i,0}(a_0^*, \chi') di. \quad (23)$$

Figure 13 shows the compensating variation for changes relative to the 1950s baseline.²⁹ Panel (a) shows the compensating variation for a change from one of the circumstances

²⁹I do not make the calculation for the change in survival probability as this depends on the assumed value of life.

being from the 1930s-born to all being from the 1950s-born and (b) shows the compensating variation for a change from all circumstances being those of the 1950s-born to one of the circumstances being from the 1980s-born. Someone would have to be given £106,000 or 3.5 times annual average 1950s earnings to accept the 1930s-born level of earnings rather than the 1950s-born level. Someone would have to be given £60,000 or 2.0 times annual average 1950s earnings to accept the 1950s-born level of earnings rather than the 1980s-born level. The other changes in circumstances that lead to substantial increases in welfare are the decreasing sizes of households and the decreasing burden of taxation. The fact that earnings come later in life for the 1950s-born compared to the 1930s-born is welfare-reducing (due to significant foregone returns). The decline in returns across generations has an compensating variation of around -£10,000 looking across both pairs of generations. Other changes have smaller effects.

Figure 13: Estimates of compensating variation of changing circumstances



Note: Figure shows the compensating variation associated with changing each circumstance from its 1930s level to its 1950s level (panel a) or from its 1950s level to its 1980s level (panel b).

6.4 What if there hadn't been a house price boom?

One major driver of the change in returns to wealth over time is the dynamics of house prices, which roughly doubled compared to average earnings between the mid-1990s and late 2000s. Given the importance of housing wealth in households' portfolios, I quantify the effect

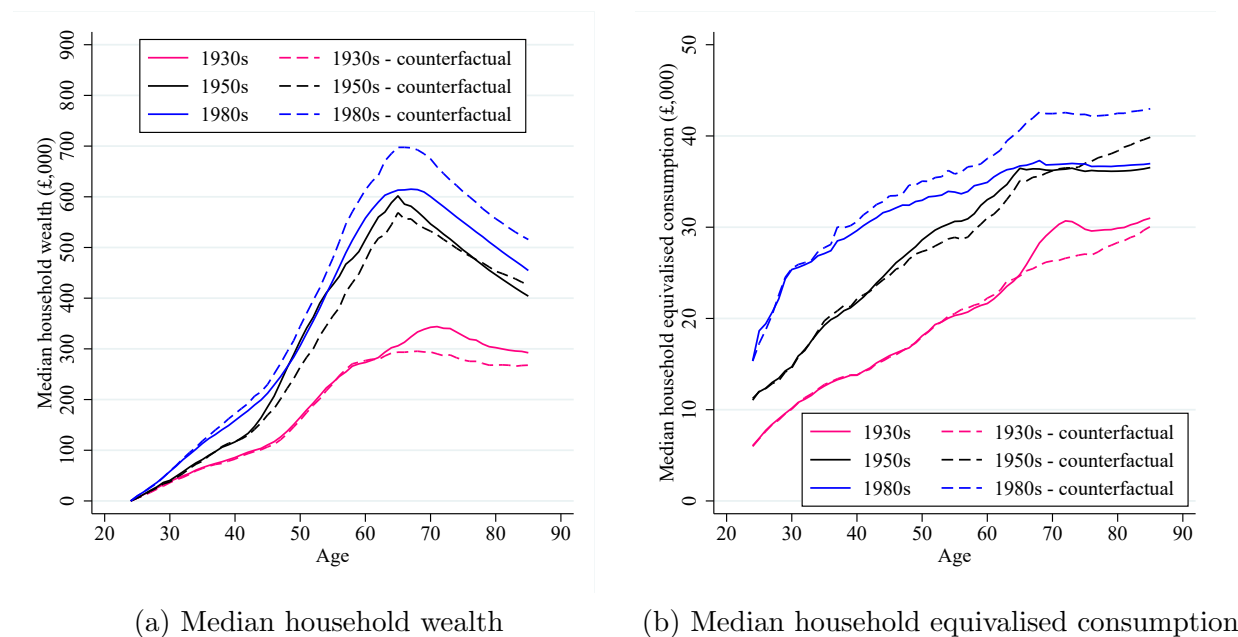
of this increase in house prices on the wealth and consumption of different generations. I compare the model simulations to a model counterfactual where the large increase in house prices compared to earnings did not take place. Specifically, in this counterfactual, house prices remain at their 1995 level, as a ratio of average household earnings, over the period 1995 to 2015, before reverting to their assumed long-run growth rate.³⁰ The flow return from housing, composed of the imputed rental yield less taxes and depreciation, is assumed to be unchanged in real terms, meaning that it is higher as a percentage of the house value in the counterfactual after 1995. Households' expectations are the same as in the baseline simulation until 2015, but households expect a higher return thereafter, in line with the lower prices and therefore higher percentage-terms rental yield.

This counterfactual implies much reduced capital gains on housing wealth during the 1990s and 2000s, and increased total returns going forward (as lower prices imply a higher flow of rents net of costs). The counterfactual also implies a change in households' portfolio composition because portfolio composition in the model is a function of the level assets as a fraction of average house prices. At moderate levels of wealth, a lower house prices imply a larger share of wealth held as housing. This captures the empirically important fact that increased house prices coincided with a decline in homeownership for those at younger ages.

Figure 14 compares the counterfactual levels of median wealth and equivalised consumption to their baseline levels, for the 1930s-, 1950s- and 1980s-born generations. Wealth is decreased significantly for those born in the 1930s and 1950s, at ages that correspond to the time when the house price boom arrived. Median consumption after the boom years is reduced. This is particularly the case for the 1930s-born generation, because they retired just as the house price boom happened, and can crystallise their capital gains as they spend down assets in retirement. Their consumption is 13% higher at age 70 as a result of the house price boom. For the 1980s-born generation, wealth and consumption are higher in the counterfactual because they hold more of their portfolio as housing wealth and it has a higher return over most of their life. Their consumption is 13% *lower* at age 70 as a result of the house price boom. This is a quantification of the 'asset price redistribution' of Fagereng et al. (2022). Without the house price boom, generational wealth and consumption growth is still strong across the older generations but growth in wealth between the 1950s and 1980s-born is stronger. Peak wealth growth between the 1950s and 1980s generations is by 23% in the counterfactual, compared to 2% in the baseline.

³⁰Mean house prices were 3.5 times mean annual household earnings in 1995. Average annual earnings is computed as mean household earnings for the population aged 16-64.

Figure 14: The effect of the increase in house prices compared to earnings on the 1930-born, 1950s-born and 1980s-born generations' median wealth and equivalised consumption



7 Conclusion

The reasons why generational wealth growth was strong between the 1930s and 1950s-born but has stopped thereafter are therefore as follows. The biggest single contribution comes from the slowing of generational earnings growth. The change in the 'shape' of the earnings profile also makes an important contribution. Over older generations, expansions in the length of working lives contributed to generational wealth growth. For younger generations, a steeper profile of earnings, stemming from higher levels of education and later entry into the workforce, implies lower wealth accumulation in early working-life but higher wealth accumulation thereafter. A number of other factors also play a smaller part. The slowing of the reduction in household size, the withdrawal of occupational DB pensions, and the slowing of the expansion of life expectancy are next most important when looking at wealth at age 65. Changes to the tax system and to the timing of partnering and children act to drive wealth growth at younger ages for the 1980s-born (compared to the 1950s-born) more strongly than for the 1950s-born (compared to the 1930s-born). Without these changes, generational wealth across younger generations would have been slower still. Importantly, generational wealth differences do not map directly to generational differences in consumption and welfare. Although generational progress in consumption has slowed, it is expected to be higher than that observed for current wealth, because younger generations are expected to see greater earnings progression as they age, and they have smaller households and face lower taxes.

References

- Altonji, J. G. and L. M. Segal (1996). Small-Sample Bias in GMM Estimation of Covariance Structures. Journal of Business & Economic Statistics 14(3), 353–366. Publisher: [American Statistical Association, Taylor & Francis, Ltd.].
- Arellano, M., R. Blundell, and S. Bonhomme (2017). Earnings and Consumption Dynamics: A Nonlinear Panel Data Framework. Econometrica 85(3), 693–734.
- Auclert, A. (2019). Monetary Policy and the Redistribution Channel. The American Economic Review 109(6), 2333–2367.
- Banks, J. and C. Emmerson (2018). A Lifetime of Changes: State Pensions and Work Incentives at Older Ages in the UK, 1948-2018. NBER Working Paper Series, 37.
- Bartels, C. and S. Morelli (2021, February). A tale of two countries: The long shadow of the crisis on income and wealth in Germany and Italy. Journal of Modern European History 19(1), 33–39.
- Bauluz, L. and T. Meyer (2021). The Great Divergence: Intergenerational Wealth Inequality in the U.S. and France. SSRN Electronic Journal.
- Bauluz, L. and T. Meyer (2022). The Wealth of Generations. SSRN Working Paper.
- Blundell, R. and B. Etheridge (2010). Consumption, income and earnings inequality in Britain. Review of Economic Dynamics.
- Blundell, R., H. Low, and I. Preston (2013). Decomposing changes in income risk using consumption data. Quantitative Economics 4(1), 1–37.
- Borella, M., M. De Nardi, M. Pak, N. Russo, and F. Yang (2022). The Importance of Modeling Income Taxes over Time: U.S. Reforms and Outcomes. NBER Working Paper Series 30725.
- Borella, M., M. De Nardi, and F. Yang (2020, January). The Lost Ones: The Opportunities and Outcomes of White, Non-College-Educated Americans Born in the 1960s. NBER Macroeconomics Annual 34, 67–115.
- Bourquin, P., R. Joyce, and D. Sturrock (2021). Inheritances and inequality over the life cycle: what will they mean for younger generations? Technical Report R188, The IFS. ISBN: 9781801030359.

- Bozio, A., R. Crawford, and G. Tetlow (2010). The history of state pensions in the UK: 1948 to 2010. IFS Briefing Note (BN105). ISBN: 9781903274774.
- Charles, K. K. and E. Hurst (2003). The correlation of wealth across generations. Journal of Political Economy 111(6), 1155–1182.
- Crawford, R. and C. O’Dea (2020). Household portfolios and financial preparedness for retirement. Quantitative Economics 11(2), 637–670.
- Cribb, J. (2019, September). Intergenerational Differences in Income and Wealth: Evidence from Britain. Fiscal Studies 40(3), 275–299.
- Cribb, J. and C. Emmerson. Workplace Pensions and Remuneration in the Public and Private Sectors in the UK. pp. 8.
- Cribb, J., A. Hood, and R. Joyce (2016). The economic circumstances of difference generations: the latest picture. IFS Briefing Note (BN187). ISBN: 978-1-911102-23-6.
- De Nardi, M. (2004). Wealth inequality and intergenerational links. Review of Economic Studies 71(3), 743–768.
- De Nardi, M., G. Fella, and G. Paz-Pardo (2019). Nonlinear Household Earnings Dynamics, Self-Insurance, and Welfare. Journal of the European Economic Association.
- De Nardi, M., E. French, and J. B. Jones (2010). Why do the elderly save? The role of medical expenses. Journal of Political Economy 118(1), 39–75.
- De Nardi, M., E. French, J. B. Jones, and R. McGee (2021). Why Do Couples and Singles Save During Retirement? Federal Reserve Bank of Richmond Working Papers 21(09), 1–65.
- Deaton, A. S. and C. H. Paxson (1994). Saving, Growth, and Aging in Taiwan. In Studies in the Economics of Aging, pp. 331–362. University of Chicago Press.
- Derenoncourt, E., C. H. Kim, M. Kuhn, and M. Schularick (2022). Wealth of Two Nations, The U.S. Racial Wealth Gap , 1860-2020.
- Disney, R., C. Emmerson, and G. Tetlow (2009, November). What is a Public Sector Pension Worth? The Economic Journal 119(541), F517–F535.
- Duffie, D. and K. J. Singleton (1993). Simulated Moments Estimation of Markov Models of Asset Prices. Econometrica 61(4), 929–952.

- Engen, E. M., W. G. Gale, C. E. Uccello, C. D. Carroll, and D. I. Laibson (1999). The Adequacy of Household Saving. Brookings Papers on Economic Activity 1999(2), 65–187.
- Fagereng, A., M. Gomez, M. Holm, B. Moll, and G. Natvik (2022). Asset-Price Redistribution. Working Paper.
- Fagereng, A., M. Mogstad, and M. Rønning (2021). Why do wealthy parents have wealthy children? Journal of Political Economy 129(3), 703–756.
- Forder, J. and J.-L. Fernandez (2011). Length of stay in care homes. PSSRU Discussion Paper.
- French, E. and J. B. Jones (2004). On the distribution and dynamics of health care costs. Journal of Applied Econometrics 19(6), 705–721.
- French, E. and J. B. Jones (2011). The Effects of Health Insurance and Self-Insurance on Retirement Behavior. Econometrica 79(3), 693–732.
- Gale, W. G., H. Gelfond, and J. J. Fichtner (2019). How Will Retirement Saving Change By 2050? Prospects for the Millennial Generation. Technical report. Publication Title: Peter G.Peterson Foundation US 2050 Project Issue: March.
- Gale, W. G., H. Gelfond, J. J. Fichtner, and B. Harris, H. (2020). The Wealth of Generations, with Special Attention to the Millenials.
- Glover, A., J. Heathcote, D. Krueger, and J.-V. Ríos-Rull (2020). Intergenerational Redistribution in the Great Recession. Journal of Political Economy 128(10).
- Greenwald, D., M. Leombroni, H. N. Lustig, and S. Van Nieuwerburgh (2022). Financial and Total Wealth Inequality with Declining Interest Rates. NBER Working Paper (28613).
- Heathcote, J., K. Storesletten, and G. L. Violante (2014). Consumption and labor supply with partial insurance: An analytical framework. American Economic Review 104(7), 2075–2126.
- Hood, A. and A. Norris Keiller (2016, November). A survey of the UK benefit system. Technical report. ISBN: 9781903274996.
- Hood, A. and L. Oakley (2014, November). The social security system: long-term trends and recent changes. Technical report, IFS. ISBN: 9781909463653.
- Hubmer, J., P. Krusell, and A. A. Smith. (2021, May). Sources of US Wealth Inequality: Past, Present, and Future. NBER Macroeconomics Annual 35, 391–455.

- Jaeger, P. and P. Schacht (2022). The Rise and Fall of Median Wealth in the U.S.: A Birth-cohort Story.
- Jorda, O., K. Knoll, D. Kuvshinov, M. Schularick, and A. M. Taylor (2019). The Rate of Return on Everything, 1870–2015. Quarterly Journal of Economics 134(3), 1225–1298.
- Kapteyn, A., R. Alessie, and A. Lusardi (2005). Explaining the wealth holdings of different cohorts: Productivity growth and Social Security. European Economic Review 49(5), 1361–1391.
- Kuhn, M., M. Schularick, and U. I. Steins (2020). Income and Wealth Inequality in America, 1949–2016. Journal of Political Economy 128(9).
- Malmendier, U. and S. Nagel (2011). Depression Babies: Do Macroeconomic Experiences Affect Risk Taking? The Quarterly Journal of Economics 126(1), 373–416.
- Meghir, C. and L. Pistaferri (2011). Earnings, Consumption and Lifecycle Choices. In Handbook of Labor Economics.
- Newey, W. K. (1985, September). Generalized method of moments specification testing. Journal of Econometrics 29(3), 229–256.
- Pakes, A. and D. Pollard (1989, September). Simulation and the Asymptotics of Optimization Estimators. Econometrica 57(5), 1027.
- Paz-Pardo, G. (2021). Homeownership and Portfolio Choice Over the Generations. SSRN Electronic Journal.
- Pope, T. and T. Waters (2016). A survey of the UK tax system. IFS Briefing Note (BN09). ISBN: 978-1-909463-68-4.
- Saez, E. and G. Zucman (2016, May). Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data. The Quarterly Journal of Economics 131(2), 519–578.
- Scholz, J. K., A. Seshadri, and S. Khitatrakun (2006). Are Americans Saving “Optimally” for Retirement? Journal of Political Economy 114(4), 607–643.
- Smith, M., O. Zidar, and E. Zwick (2022, December). Top Wealth in America: New Estimates Under Heterogeneous Returns. The Quarterly Journal of Economics 138(1), 515–573.

A Online Appendix: Numerical method for solving the lifecycle model

As there is no analytic solution to the household problem, it must be solved numerically. The model has 9 state variables: generation, household education level, household earnings fixed effect, whether the household has a defined benefit pension, the household persistent earnings component, employment status, age, whether the parental household is alive, and start-of-period cash-on-hand. In some periods, these state variables can take on only one value. For example, after the latest age at which a parent can die, there is only one value for whether the parental household is alive. I solve the model at a number of discrete points on 9-dimensional grid for the state variables. There are 6 generations, 3 education groups, 3 values for the household fixed effect and 2 values for whether or not the household has a defined benefit pension. This yields 108 household types (combination of values of these first 4 state variables). For each household type I create a beginning-of-period cash-on-hand grid which has 20 values at each age. These values span from the minimum possible beginning-of-period cash-on-hand (the sum of pretax income and the age-specific borrowing limit) to the maximum possible beginning-of-period cash-on-hand (the sum of the maximum possible pretax income and the maximum possible beginning-of-period assets). Maximum beginning-of-period assets is obtained by receiving maximum possible income in all prior periods, saving all income in all periods, and receiving the maximum possible return to assets in all periods, and is defined as:

$$a_{i,t}^{max} = \sum_{\tau=25}^{\tau=t-1} [(y_{\tau}^{max}) \Pi_{s=\tau}^{s=t-1} (1 + r_{s+1}^{max})] \quad (24)$$

where y_t^{max} is the maximum pretax income in period t that can be received by the relevant household type. Grid points are spaced so as to be more dense at lower asset levels. There are 87 periods (from age 24 to a maximum age of 110).

The choice variable for the household is consumption, or, equivalently, end-of-period assets. I calculate the household's optimal choice at each grid point as follows. For each household type, I solve the model recursively, beginning in the final period. At the end of the final period, the household will die with certainty and decides how much of start-of-period cash-on-hand to consume (which requires the payment of taxes) and how much to leave as end-of-period assets. The sources of uncertainty are the return on assets and the end-of-life cost. The assets remaining after the realisation of these shocks is the household's bequest.

The optimal choice of end-of-period assets, $a_{i,T}^*$ solves

$$V_T(M_T; \omega, e_{K-1}) = \max_{a_T^*} \left\{ u(\tau(M_T - a_T^*)) + \beta \int \int \phi(b(a_T^*(1 + r_{T+1}))) dF_{\omega, T+1}^r(r_{T+1} | a_{i,T}^*) dF_g^\kappa(\kappa) \right\}, \quad (25)$$

subject to the borrowing limit. For each point on the grid of state variables in period T , I find the level of end-of-period assets in period T that solves this value function using the Matlab function ‘fminbnd’, which uses golden section search and parabolic interpolation methods to find the maximum of this expression within the feasible set of end-of-period assets. Numerical integration over the distributions of returns and end-of-life costs uses the method of Tauchen. With the optimal choice of consumption and end-of-period assets, I obtain the maximised value function at that grid point.

I then move backwards through the late retirement periods of the household’s lifecycle, solving for the optimal choices of end-of-period assets at each grid point. Rather than storing the value of the realised level of cash-on-hand at each grid point conditional on survival (i.e. storing $V_t(M_t; \omega, e_{K-1})$), I integrate over the distribution of rate-of-return and end-of-life cost shocks, conditional the household type, final working-life earnings, and a grid of values for end-of-period assets, in the prior period. This grid for end-of-period assets is defined analogously to the grid for start-of-period cash-on-hand. This yields a set of values for the expected value of end-of-period assets, conditional on survival to the next period, for each combination of state variables, which I store. I denote this expected continuation value as $\hat{V}_t(a_{t-1}^*; \omega, e_{K-1})$. In the late retirement periods up to period $T - 1$, I therefore obtain the optimal choice of end-of-period assets, a_t^* as the solution to

$$V_t(M_t; \omega, e_{K-1}) = \max_{a_t^*} \left\{ u(\tau(M_t - a_t^*)) + \beta s_{t+1} \hat{V}_{t+1}(a_t^*; \omega, e_{K-1}) + \beta(1 - s_{t+1}) \int \int \phi(b(a_t^*(1 + r_{t+1}))) dF_{\omega, t+1}^r(r_{t+1} | a_{i,t}^*) dF_g^\kappa(\kappa) \right\}, \quad (26)$$

subject to the borrowing constraint. I find the solution to the value function using ‘fminbnd’, linearly interpolating the expected continuation value function, $\hat{V}_{t+1}(a_t^*; \omega, e_{K-1})$, between the end-of-period asset grid points at which values are stored.

In the early retirement periods, the optimal choice and value function vary by whether or not the parental household is alive. In working-life, choices and value functions additionally depend on the persistent earnings shock and employment status. The approach is analogous: in each period, I solve for the optimal choice of end-of-period assets at each grid point. I

then integrate out the realisation of shocks and the realisations of the evolution of the state variables, conditional on end-of-period assets and each combination of the state variables (other than cash-on-hand) in the prior period. I store this expected continuation value function and linearly interpolate it when solving for the optimal choice of end-of-period assets in the prior period. I work backwards from the oldest to the youngest age, solving for optimal choices at each grid point until reaching the initial age.

Each state variable other than cash-on-hand takes on a discrete set of values. When simulating the model, I linearly interpolate the policy function for end-of-period assets when the simulated level of cash-on-hand is not on one of the grid points for which a solution is calculated and stored.

B Online Appendix: Further details on estimation of economic and demographic circumstances

B.1 Earnings process

Data: I use two datasets. The first is the Family Expenditure Survey (later called the Expenditure and Food Survey and then the Living Costs and Food Survey, with this collection of surveys referred to here as “FES”), an annual household survey running since 1968. This measures household earnings and household members’ age and (since 1978) level of education. The second dataset is the UK Household Longitudinal Study and its predecessor the British Household Panel Study (I refer to the combined data as “UKHLS”). This annual household panel has run since 1991 and measures household earnings, and household members’ age and education. The two data sources are complementary: the FES runs for a longer time period, allowing us to observe older generations when they were at younger ages, while the UKHLS is a panel, allowing us to observe earnings and employment transitions. I use data for all years from 1968 to 2018 in the FES and 1991 to 2018 in the UKHLS and keep households with year of birth between 1930 and 1989. In both datasets, I construct a measure of employment and total gross annual earnings at the household level. Households consist of either a single individual or a couple, plus any dependent children the individual or couple has. This corresponds to the UK concept of a ‘benefit unit’ and defines adults who are co-resident but not in a couple as separate households. A household is employed if any member records positive earnings and is not employed otherwise. Earnings are the sum of all earnings from employment and self-employment recorded in the data. When selecting a balanced panel of benefit units in the UKHLS, I define a benefit unit as the same benefit unit in another period if it has the same members (excluding dependent children). I enforce

that a panel cannot run across the ‘seam’ between the British Household Panel Survey and Understanding Society years of the UKHLS surveys (due to a change in periodicity). When selecting balanced panels I also drop observations where consecutive interviews are less than 9 months or more than 15 months apart (this is 3% of interviews).

Estimation and identification: I employ a multi-step procedure common in the earnings process literature (see, for example, Meghir and Pistaferri (2011)). In the first step, I estimate the deterministic component of log earnings by estimating age-profiles for log household annual gross earnings using the FES (using years from 1978 onwards when education is measured). Let e_{it} , g_i , ed_i and age_{it} denote the earnings, generation, education level and age of household i in time t . I estimate the following specification using OLS:

$$\ln(e_{it}) = \sum_{G=1930s}^{1980s} \sum_{E=1}^3 \beta_{G,E} \mathbb{1}\{g_i = G \wedge ed_i = E\} + \sum_{A=16}^{64} \sum_{E=1}^3 \gamma_{A,E} \mathbb{1}\{age_{i,t} = A \wedge ed_i = E\} + \delta_t + \sum_{k=1}^{45} DP_{it}^k + \epsilon_{i,t} \quad (27)$$

where δ_t is a linear time trend and $\sum_{k=1}^{45} DP_{i,t}^k$ are a series of time effects constrained to sum to zero (see Deaton and Paxson (1994)). I control for the interaction of education and generation dummies as well as the interaction of a full set of year-of-age dummies with education dummies. This embodies the assumption that while the average level of earnings can vary across cohort-education groups, average log earnings are assumed to vary with age in a way which is depends on the individual’s education level, but does not depend on their generation. The time trend and Deaton-Paxson dummies can be interpreted as cyclical deviations around a linear time trend, embodying the macroeconomic trend and cycle. This form allows me to control flexibly for the effects of age, while allowing the returns to different levels of education to vary across cohorts.³¹

In the second step, I estimate the parameters governing the distribution of the household fixed effects and the distribution and evolution of the persistent stochastic component of earnings. The estimation of these components follows Arellano et al. (2017) and uses the UKHLS data. I estimate the equivalent of Eq. (27) using the UKHLS data and obtain the regression residuals for each household. I then use the estimation algorithm of Arellano et al. (2017) to estimate the distribution of the permanent, persistent stochastic, and transitory components of earnings and the quantile functions governing the transitions of the stochastic

³¹Separation of time effects into a linear time trend and a cyclical fluctuations around this means that I can extrapolate to years outside the data period under the assumption of a constant time trend and no future cyclical fluctuations.

component. This estimation procedure is conducted separately for each education group, allowing for the possibility that household heterogeneity and earnings dynamics will differ by education level. I use the same specification of the order of the hermite polynomials for each component of the process as used by Arellano et al. (2017) (tensor products of polynomials of degree 3 (in the lagged persistent component) and 2 (in age) for the transitions of the persistent component, and polynomials of degree 2 in age for the initial distribution of the persistent component and for the transitory component). Again following Arellano et al. (2017), I use quantile regressions of earnings on their lagged values and age to set initial parameter values. I use the same variances for the random walk proposals in the Metropolis-Hastings sampler as Arellano et al. (2017), which yields an acceptance rate of around 0.20-0.25.

In the UK, earnings risk increased markedly between the late 1970s and early 1990s (Blundell et al. (2013); Blundell and Etheridge (2010)). To account for this, I assume that this change took the form of an increase in the age-and-education-specific dispersion of the persistent earnings shocks over time and apply a year-and-education-specific loading to the persistent stochastic component of earnings in years prior to 1991. I estimate the percentage increase in the age-and-education-specific variance of log earnings over the period from 1980 to 1991 using the FES. For generations observed before 1991, this yields an age-and-education-specific loading for the persistent stochastic component that fits the increase in the variance of log earnings observed in the cross-sectional data. I discretise the estimated distributions of the fixed effects and distributions and transitions of the stochastic component using the approach of De Nardi et al. (2019). I estimate employment transitions non-parametrically, conditional on household type and age.

I pursue the flexible quantile function approach to modelling household earnings dynamics because it has been shown that there is important age-dependency, nonlinearity and conditional skewness in earnings innovations and that this has quantitatively important implications for the level and distributions of wealth when incorporated into a lifecycle savings model (see De Nardi et al. (2019) and Arellano et al. (2017)). Incorporating the richness of household earnings risk into our model is therefore likely to be important for our estimated preference parameters and resulting conclusions about how these vary across generations. Appendix B describes the estimation in more detail and shows that the estimated earnings process fits key cross-sectional and dynamic properties of the household earnings distribution.

Identification of the parameters of the deterministic components of the earnings process given cross-sectional data is standard. Given the assumption of common age profile across cohorts for each education group, we can separately identify the cohort and age effects for

each education group so long as multiple cohorts are observed at each age. The time trends are separately identified from the age and cohort effects given the Deaton-Paxson restriction.

Because the UKHLS data begins in 1991 and the youngest generations have not completed working life by 2018, I cannot conduct this estimation procedure separately by generation. I account for the changes in earnings risk that happened before 1991 by allowing the age-and-education-specific quantiles of the persistent stochastic component to vary by generation. For each age-and-education group the age-and-education specific quantiles are assumed to be some multiple of that estimated in the UKHLS data. Letting $Q_{ed,t}(\cdot)$ be the conditional quantile function estimated using the UKHLS data. The conditional quantile function for household type ω is given by

$$Q_{g,ed,t}(\cdot) = L_{g,ed,t}Q_{ed,t}(\cdot). \quad (28)$$

The generation-education-age-specific loadings are informed by the cross-sectional FES data. I estimate the average annual percentage change in the education-and-age-specific variance of log earnings over the period from 1980 to 1991. Each generation observed before 1991 therefore has a loading (less than 1) that varies by education and age, depending on the years in which that generation was at each age.

I discard the transitory component when putting the earnings process into the lifecycle model, on the basis that this partly consists of measurement error. I discretise the estimated distributions and transitions using the approach of De Nardi et al. (2019). I discretise the household fixed effect distribution into 3 equal sized bins. I discretise the distribution of the persistent component at each age into 3 bins covering the bottom 25%, middle 50% and top 25% of the distribution, respectively. Given I estimate the earnings process separately for each education group, this yields an earnings process with 27 possible positive values of earnings at each age for each generation.

The probabilities of being in employment are a function of generation, education, age and the household earnings fixed effect. I estimate these probabilities non-parametrically, by calculating the proportion of households that transition in and out of employment for each generation, education, age and tertile of the earnings distribution (as there are three levels of the household fixed effect). For generations not observed over the whole of their working lives, I base their employment transitions on those of the nearest generation that is observed at that age. To do this, I take the employment transitions for each household type for the adjacent generation and scale the probability of transition proportionally to fit the observed employment rate for each household type in the FES (which covers a longer time period).

As demonstrated in Arellano et al. (2017), the marginal distributions the stochastic com-

ponent and their transitions are identified given 4 periods of panel data and some conditions on the process, including a completeness condition. Intuitively, the completeness condition requires that there is some dependence of the persistence component such that they can be distinguished from the transitory component. As we have data on multiple cohorts for whom the initial and terminal ages are different, then by the assumption of the invariance of these distributions of over time, and with the parametric restrictions in the empirical specification of the model, we can recover these distributions at all ages.

Model fit: Figure B.1 and Figure B.2 show the comparison between the simulations of household earnings and employment rates by generation and education group and their equivalents in the FES data.


B.2 Public pension system

Estimation: I estimate the state pension system parameters as follows. I simulate the household earnings process for each generation for 90,000 working lifetimes. For each of these earnings histories, I use a full state pension calculator (taken from Banks and Emmerson (2018)) to calculate the state pension entitlements that would accrue to each household given their earnings and employment history. I obtain the state pension function parameters by estimating the following specification using OLS:


$$pub_{i,K} = \sum_{g=1930s}^{1980s} \sum_{ed=1}^3 \sum_{\zeta=1}^3 \alpha_{g,ed,\zeta} gen_i \times ed_i \times \zeta_i + \sum_{g=1930s}^{1980s} \sum_{ed=1}^3 \sum_{\zeta=1}^3 \beta_{g,ed,\zeta} gen_i \times ed_i \times \zeta_i \times e_{i,K-1} + \epsilon_{i,t} \quad (29)$$

In order to make the calculations of state pension entitlements, some information must be imputed to the simulated earnings histories. First, household earnings must be divided into individual earnings. To do this, I use the FES to estimate the share of earning households that are two-earner couples for each generation, education level and age and randomly assign that proportion of earnings histories to be two-earner households. For two-earner couples, the split of earnings between the primary and secondary earner is also estimated as a function of age and household type in the FES and imputed into the simulated earnings histories accordingly. State pension entitlement depends significantly on childcare undertaken and I assume that women aged between 25 and 34 who are in a couple but are not in work are caring for a child. With this information, state pension entitlements are calculated for the individual members of a household and then summed to give the household's state pension

Figure B.1: Comparison of mean earnings in earnings model simulations and FES data, by generation and education group




(a) 1930s-born




(b) 1940s-born




(c) 1950s-born



(d) 1960s-born




(e) 1970s-born




(f) 1980s-born


Figure B.2: Comparison of employment rate in earnings model simulations and FES data, by generation and education group




(a) 1930s-born




(b) 1940s-born




(c) 1950s-born



(d) 1960s-born



(e) 1970s-born



(f) 1980s-born

entitlement.

B.3 Occupational pension provision

Estimation: Using ELSA, I calculate the proportion of households with any DB pension income and the average value of annual household defined benefit pension income amongst those with DB income, for each generation-and-education group for those born in the 1930s, 1940s and 1950s. For generations born from the 1960s onwards, the DB pension function is assumed to be equal to those for the 1950s generation, with the prevalence of DB allowed to change. To estimate prevalence of DB pensions for the 1960s, 1970s and 1980s generations, I use the estimates of DB prevalence by generation from Cribb et al. (2016) and assume that differences in DB prevalence between generations hold within each education group. I assume that households who have a DB pension must contribute 4% of their earnings to the pension with the employer assumed to provide the remaining funding required to pay the pension income.³²

I estimate employer contributions for those without a DB pension in each year as follows. Of those individuals in ASHE that report not having a DB pension, I take the percentage of individuals who report having a DC pension and multiply this by an assumed employer contribution rate of 4%.³³

To calculate the expected present discounted value of accrued DB entitlements, I use survival probabilities and discount rates together with the formula described in section 3. The survival probabilities are those described below. For the period from retirement onwards, I discount the pension income stream based on the 15-year government gilt rate. For ages before 65 I discount using government discount rates. These two choices reflect the way that defined benefit pension income is valued in the Wealth and Assets Survey data that I use to estimate the model.

B.4 Rates of return to assets

I estimate the expected distribution of returns for each asset class as follows. For each asset class k , I take the annual real returns from all years from 1955 to 2015 and remove a linear trend by estimating the following equation using OLS:

$$r_{k,t} = \alpha_k + \beta_k \cdot t + \epsilon_{k,t}. \quad (30)$$

³²Disney et al. (2009) reports average employee DB contributions of 4.6% in the private sector and 3.9% in the public sector in the year 2000.


³³Rates of membership of pensions by type are reported in <https://www.ons.gov.uk/employmentandlabourmarket/peoplein>

Let the regression residual for asset type k in year t be $\hat{\epsilon}_{k,t} = r_{k,t} - \hat{\alpha}_k + \hat{\beta}_k \cdot t$. Define $\hat{\alpha}$, $\hat{\beta}$ and $\hat{\epsilon}_t$ as the vectors of estimates and residuals stacked over asset classes. Households of type ω and age $t - 1$ expect that r_t is a draw from the set $\{\hat{\alpha}'\pi_t^\omega + 1955 \cdot \hat{\beta}'\pi_t^\omega + \epsilon'_{1995}\pi_t^\omega, \dots, \hat{\alpha}'\pi_t^\omega + 2015 \cdot \hat{\beta}'\pi_t^\omega + \epsilon'_{2015}\pi_t^\omega\}$, i.e. the return is expected to be the long-run trend plus a draw from the historic distribution around that trend, given the household's portfolio. This method draws the stochastic components of returns from a year that is common across asset classes in order to capture the co-movement of returns across asset classes. For years after 2015, the deterministic component of returns is assumed to remain at its 2015 level. To use this distribution in the model, I discretise it into 7 bins for each generation, age and asset grid point (and its implied portfolio composition), taking the mean return within each bin. The seven bins cover 5%, 10%, 20%, 20%, 20%, 10% and 5% of the distribution, respectively.

B.5 Inheritances

Estimation: For the 1950s to 1980s generations, earnings type is measured by the within-age-group tertile of household average earnings (averaging earnings over all observations of the household). The responses to the inheritance expectations were given in 2006-08 when the 1980s generation were aged between 16 and 28. The distribution of expected inheritance amounts for the 1980s-born interviewed at younger ages is much lower than for older individuals and seems implausible given the wealth holdings of the generations of their parents. Consequently, I use the percentage of 1980s-born who report that they expect to inherit in each education and earnings group and then assume that the distribution of inheritance amounts, conditional on inheriting, is equal to that for the equivalent group in the 1970s generation.

Figure B.3: Distribution of inheritances, by generation



Note: Figure shows the estimated proportion of households that have inherited or will inherit more than certain amounts over their lifetime, by generation. Source: WAS, waves 1 to 3

To obtain the parental household survival probabilities, I use the FES to estimate the distribution of age gaps between parents and children and the implied distribution of birth years for the parents of each generation. I then take the implied average survival curve for a parental household, assuming the parental household is a couple whose mortality realisations are independently drawn and defining parental household death as the death of the second person to die.

C Online Appendix: Asymptotic distribution of parameter estimates and model overidentification tests

Here I describe the asymptotic properties of the estimated parameter vector and overidentification tests of the model's specification. Under the regularity conditions stated in Pakes and Pollard (1989) and Duffie and Singleton (1993), the estimator $\hat{\Delta}$ is both consistent and asymptotically normally distributed:

$$\sqrt{n}(\hat{\Delta} - \Delta_0) \xrightarrow{p} N(0, \mathbf{V}), \quad (31)$$

where the variance-covariance matrix is given by

$$\mathbf{V} = (1 + \tau)(\mathbf{D}'\mathbf{W}\mathbf{D})^{-1}\mathbf{D}'\mathbf{W}\mathbf{S}\mathbf{W}\mathbf{D}(\mathbf{D}'\mathbf{W}\mathbf{D})^{-1} \quad (32)$$

where \mathbf{S} is the variance-covariance matrix of the data and

$$\mathbf{D} = \frac{\partial G(\Delta, \chi_0)}{\partial \Delta'} \Big|_{\Delta=\Delta_0} \quad (33)$$

is the derivative of the population moment vector with respect to the parameter vector, Δ , at its true value and $W = \text{plim}_{n \rightarrow \infty} \{\hat{W}_n\}$.

Newey (1985) shows that if the model is properly specified then,

$$\frac{n}{1+\tau} \hat{G}' \mathbf{R}^{-1} \hat{G} \xrightarrow{p} \chi^2_{J-M} \quad (34)$$

where J and M are the length of G and Δ , respectively i.e. number of moments and number of estimated parameters and \mathbf{R}^{-1} is the generalised Penrose inverse of

$$\begin{aligned} \mathbf{R} &= \mathbf{P} \mathbf{S} \mathbf{P}' \\ \mathbf{P} &= \mathbf{I} - \mathbf{D}(\mathbf{D}' \mathbf{W} \mathbf{D})^{-1} \mathbf{D}' \mathbf{W} \end{aligned}$$

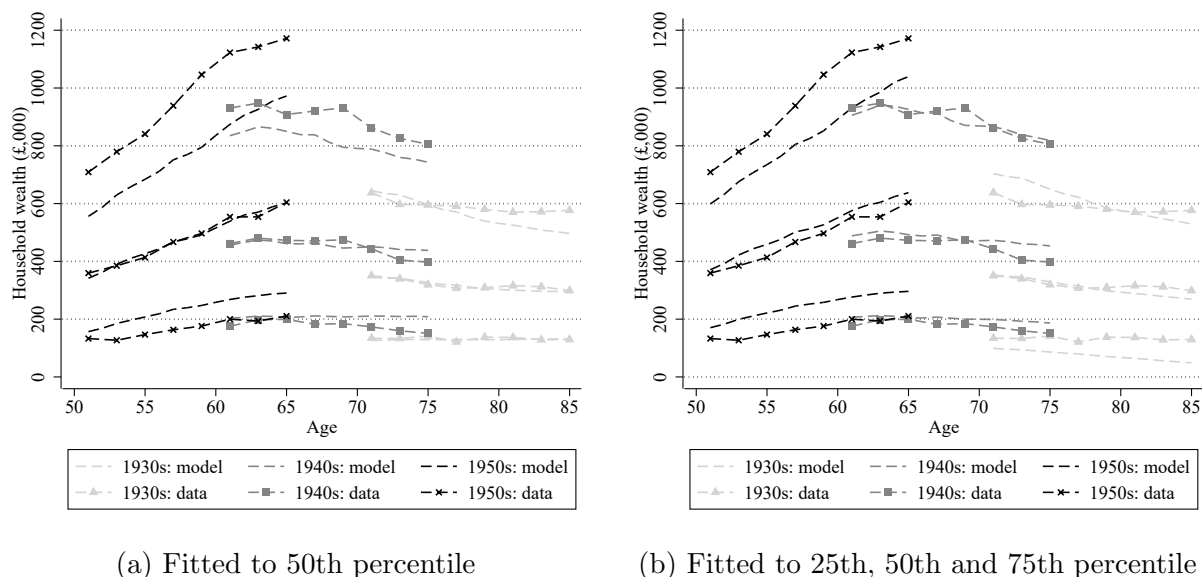
As noted in the main text, I use a weighting matrix whose diagonal elements are equal to those on the inverse of the variance-covariance matrix of the data used in estimation and whose non-diagonal elements are zero. This is because, as shown by (Altonji and Segal, 1996), while the asymptotically efficient weighting matrix arises when \mathbf{W} converges to \mathbf{S}^{-1} , such a weighting matrix can lead to substantial bias in finite samples. I estimate \mathbf{D} , \mathbf{S} and \mathbf{W} with their data analogues. Following French and Jones (2011), when constructing \hat{W} , I replace $W_{g,t}^q(\Delta, \hat{\chi})$ with the equivalent sample asset quantile before calculating the data variance-covariance matrix. The row of the derivative matrix, \mathbf{D} corresponding to the quantile q , generation g and age t and can be rewritten as:

$$\mathbf{D} = \frac{\partial W_{g,t}^q(\Delta_0, \chi_0)}{\partial \Delta'} \cdot f(W_{g,t}^q(\Delta_0, \chi_0) | i \text{ is alive at } t \text{ and in generation } g) \quad (35)$$

that is the derivative of the relevant wealth quantile for the relevant generation and age, with respect to the parameter vector, multiplied by the density of wealth at that quantile for that generation when observed at that age. I estimate the derivative vector by simulating the estimated model, perturbing each element of the parameter vector (in practice I increase the value of each parameter by 0.5%, 1% and 2% and then average any final statistics, such as standard errors, over their corresponding three values). I estimate the density at each conditional wealth moment using a kernel density estimator.


D Online Appendix: Model fit

Figure D.1: Comparison of 25th, 50th and 75th percentile of wealth in WAS data and model simulations given fixed preferences across generations



Note: The 3 pairs of series for each generation show the 25th percentile (lowest pairs of lines), median (middle pairs of lines) and 75th percentile (highest pairs of lines). Source: WAS and model simulations.

Figure D.2: Simulated and actual growth in median consumption relative to prior generation



Note: Figure shows the growth rate in median consumption in the data and model simulations, for ages for which consumption data is available for the each pair of generations. Source: FES and model simulations.

E Online Appendix: Definition of counterfactual model parameters

Here, I set out how counterfactual simulations in which one generation is faced with the circumstance(s) of another are defined. Each circumstance corresponds to a set of model parameters. For some circumstances (including the tax and welfare system and asset returns), the corresponding model parameters do not vary across household types within each generation. In these cases, I can simulate one generation facing the circumstances of an alternative generation by using the relevant model parameters from the alternative generation. In some other cases (most notably earnings, household size and composition, longevity and inheritances) the model parameters vary by household type within generation. Given that there are different shares of each education group within each generation, a counterfactual simulation that simply used the model parameters from each corresponding household type in an alternative generation would not accurately reflect the change in circumstances at the generation level. My approach for each circumstance is as follows:

Level of earnings

Let $F_{g,t}^\omega(e_t)$ denote the unconditional distribution of earnings at age t for a member of generation g of household type ω . Then the counterfactual distribution of earnings when facing the circumstances of some other generation g' is given by:

$$F_{g',t}^\omega(\alpha \cdot e_t) = F_{g,t}^\omega(e_t) \quad (36)$$

where

$$\alpha = \frac{\mathbb{E}_{24}[e_{i,t}|g_i = g', t < K]}{\mathbb{E}_{24}[e_{i,t}|g_i = g, t < K]}. \quad (37)$$

This rescales the distribution of earnings at each age for each household type such that the lifetime average earnings in the counterfactual simulation are the same as lifetime average earnings for the alternative generation, g' .

Survival probabilities

For each education-and-earnings-group within a generation, g , the counterfactual one-year-ahead survival probabilities when faced with the circumstances of generation g' are denoted as $s_{w',t}$ and defined as:

$$(1 - s_{w',t}) = \frac{\mathbb{E}[1 - s_{i,t}|g_i = g']}{\mathbb{E}[1 - s_{i,t}|g_i = g]}(1 - s_{w,t}). \quad (38)$$

This rescales the one-period-ahead hazard rate in each education-and-earnings group such that the aggregate hazard rate of the generation is equal to that for the alternative generation. Because the growth rate of the household equivalisation factor from age 65 onwards captures the survival probability of the first member of a couple, this is also changed analogously.

Household size and composition

For each education-group within a generation, g , the counterfactual equivalisation factor when faced with the circumstances of generation g' is denoted as $\theta_{w',t}$ and defined as:

$$\theta_{w',t} = \frac{\mathbb{E}[\theta_t|g_i = g']}{\mathbb{E}[\theta_t|g_i = g]}\theta_{w,t}. \quad (39)$$

This rescales the one-period-ahead hazard rate in each education-and-earnings group such that the aggregate hazard rate of the generation is equal to that for the alternative generation.

Inheritances

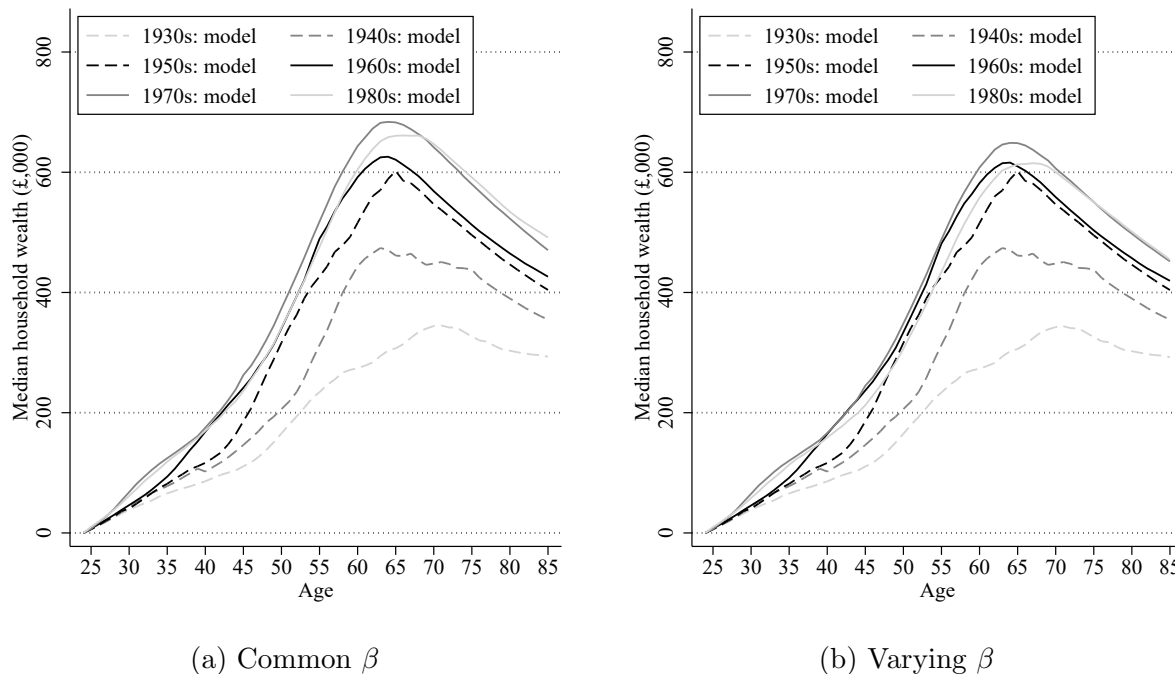
The CDF from which inheritances are drawn for a household of type ω is denoted F_ω^h . Let the unconditional distribution of inheritances received by generation g be given by the CDF F_g^h . Then for each education-and-earnings group within generation g , counterfactual CDF from which inheritances are drawn when faced with the circumstances of generation g' is denoted as $F_{\omega'}^h$ and defined as:

$$F_{\omega'}^h\left(\frac{F_{g'}^h(h)}{F_g^h(h)} \cdot h\right) = F_\omega^h(h). \quad (40)$$

This rescales the probability of receiving an inheritance of size h in each education-and-earnings group such that the aggregate probability of receiving an inheritance of this size is equal to that for the alternative generation.

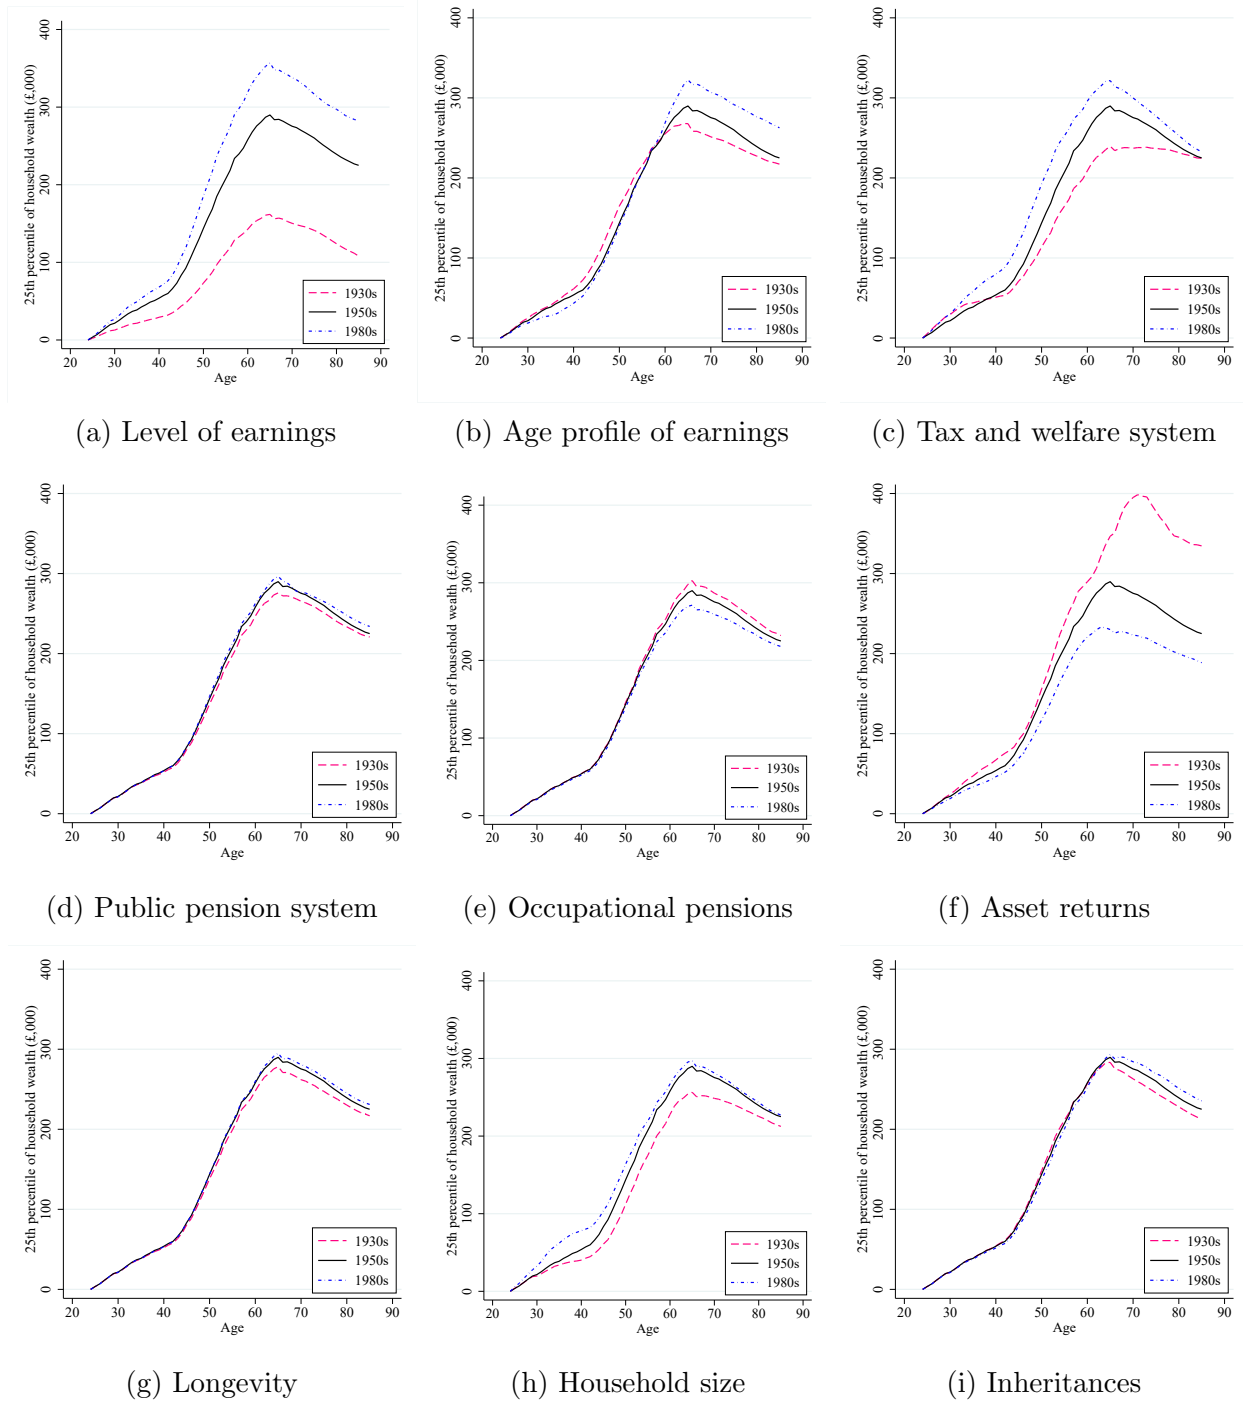
F Online Appendix: Supplementary results

Figure F.1: Simulations of median wealth under common and generation-specific patience



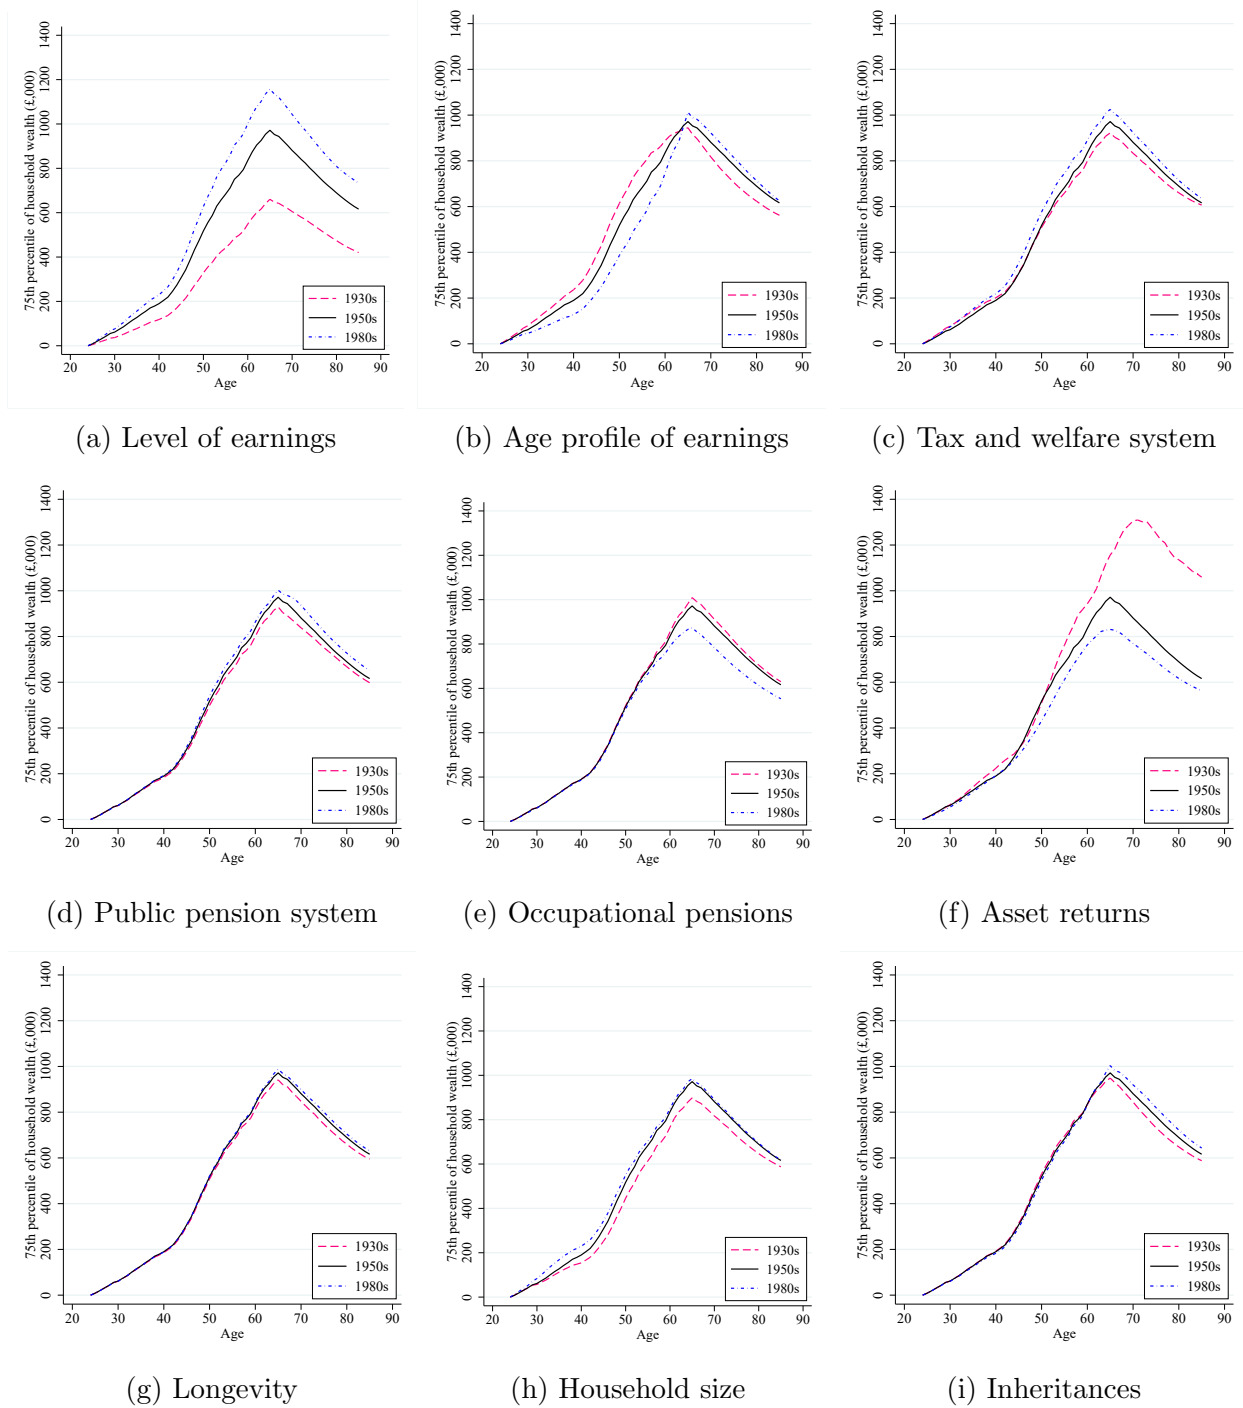
Note: Model simulations with patience (a) common and (b) varying across generations, combined with the generation-specific model inputs capturing economic and demographic circumstances.
Source: Model simulations.

Figure F.2: Impact of varying individual circumstances on the 25th percentile of household wealth



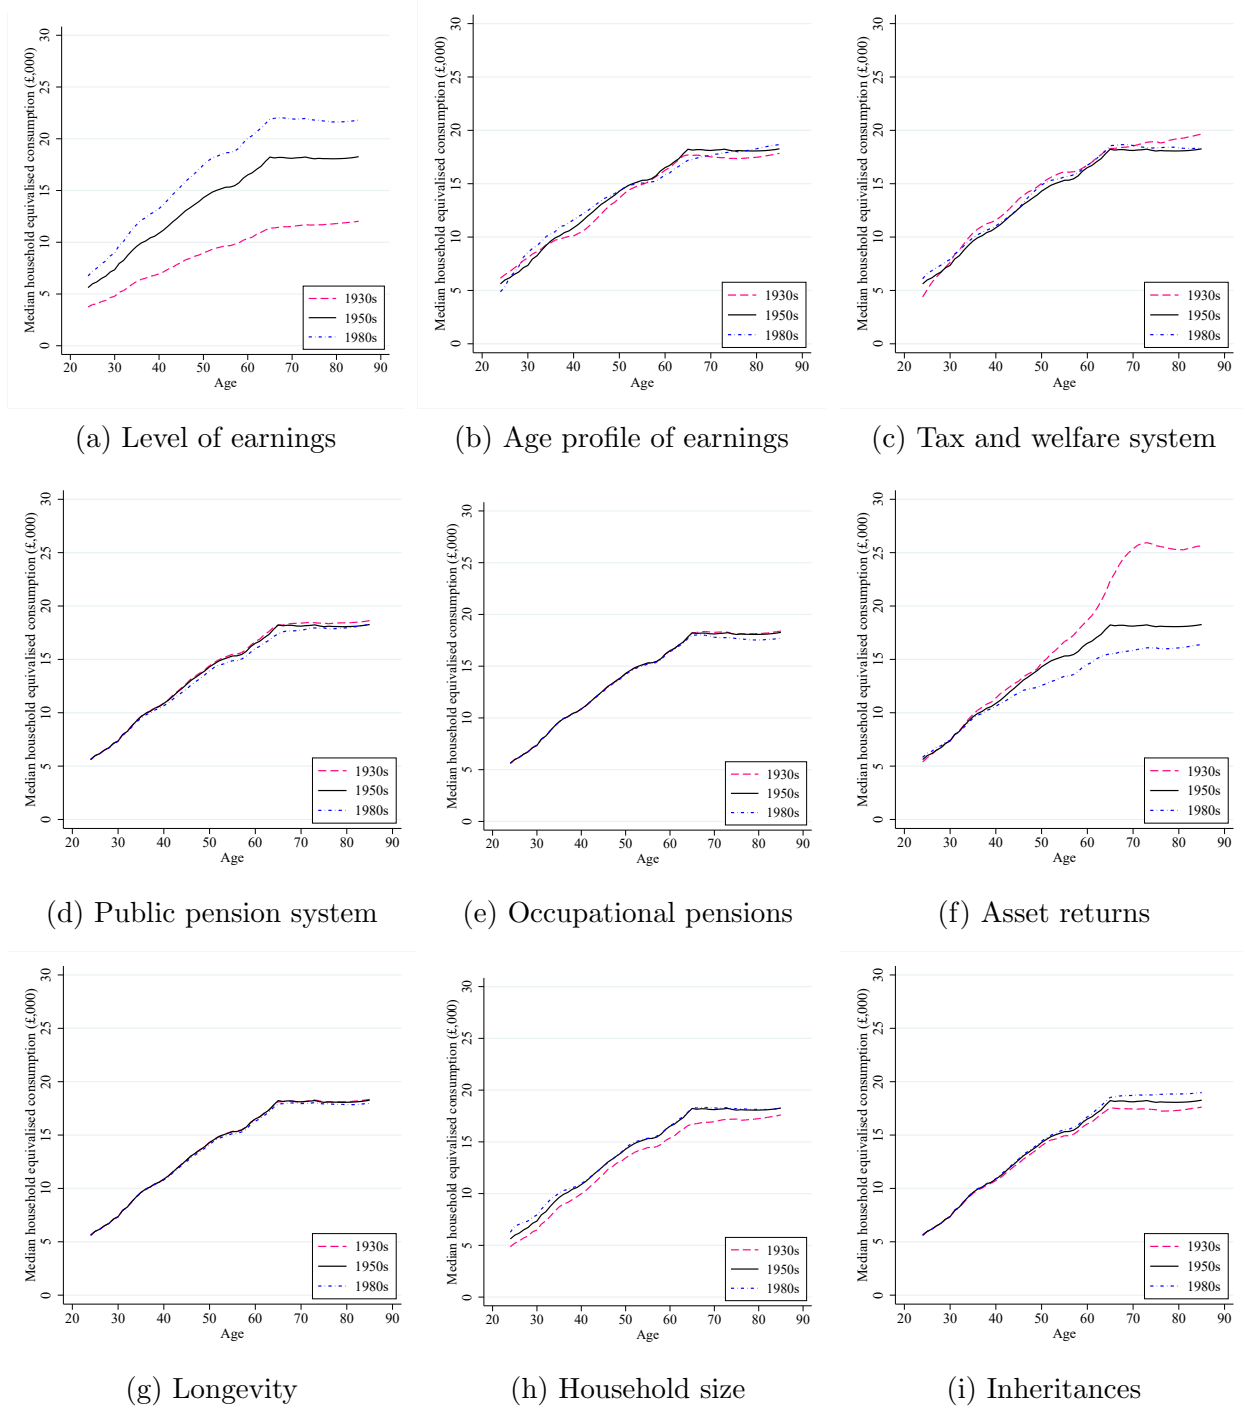
Note: Each figure shows the effect of varying individual circumstances to their levels for the 1930s, 1950s and 1980s-born generations while all other circumstances are set to those for the 1950s-born generation.

Figure F.3: Impact of varying individual circumstances on the 75th percentile of household wealth



Note: Each figure shows the effect of varying individual circumstances to their levels for the 1930s, 1950s and 1980s-born generations while all other circumstances are set to those for the 1950s-born generation.

Figure F.4: Impact of varying individual circumstances on median household equivalised consumption



Note: Each figure shows the effect of varying individual circumstances to their levels for the 1930s, 1950s and 1980s-born generations while all other circumstances are set to those for the 1950s-born generation.