BANK OF ENGLAND

Staff Working Paper No. 701
Demographic trends and the real interest rate

Noëmie Lisack, Rana Sajedi and Gregory Thwaites

December 2017

BANK OF ENGLAND

Staff Working Paper No. 701

Demographic trends and the real interest rate

Noëmie Lisack,[1] Rana Sajedi[2] and Gregory Thwaites[3]

Abstract

We quantify the impact of past and future global demographic change on real interest rates, house prices and household debt in an overlapping generations model. Falling birth and death rates can explain a large part of the fall in world real interest rates and the rise in house prices and household debt since the 1980s. These trends will persist as the share of the population in the high-wealth 50+ age bracket continues to rise. As the United States ages relatively slowly, its net foreign liability position will grow. The availability of housing and debt as alternative stores of value attenuates these trends. The increasing monopolisation of the economy has ambiguous effects.

Key words: Demographics, population ageing, neutral interest rates.

JEL classification: E13, E21, E43, J11.

(1) Bank of England. Email: noemie.lisack@bankofengland.co.uk
(2) Bank of England. Email: rana.sajedi@bankofengland.co.uk
(3) Bank of England. Email: gregory.thwaites@bankofengland.co.uk

The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England or its committees. We would like to thank Pavel Brendler, Juan Francisco Jimenez, Pierre-Olivier Gourinchas, Richard Harrison, Benjamin Johansen, Jonna Olsson, Jacob Robbins, Alan Taylor, Pablo Winant, participants at the EUI Alumni conference, DEGIT 2017 conference, and ESM, OECD, RBA and HMT research seminars, for their useful comments and suggestions. All errors remain ours.

The Bank's working paper series can be found at www.bankofengland.co.uk/news/publications

Publications and Design Team, Bank of England, Threadneedle Street, London, EC2R 8AH
Telephone +44 (0)20 7601 4030 email publications@bankofengland.co.uk

ISSN 1749-9135 (on-line)

1 Introduction

In the 2010s, advanced countries' long-term real interest rates fell well below zero, to levels unprecedented in a period of peacetime and stable inflation. While the global financial crisis has played a part, this was the continuation of a downward trend that began at least two decades previously. At the same time, the population of advanced countries has continued to age, with life expectancy and the old-age dependency ratio reaching new highs. This paper quantifies the link between these two important trends.

The advanced world is in the midst of a rapid and unprecedented ageing of its population, driven by a fall in birth rates and, more importantly, a rise in life expectancy. When old-age pensions were first institutionalised in the early 20th century, the chance of reaching pensionable age, and residual life expectancy at that point, were relatively low. In contrast, the overwhelming majority of the population can now expect to retire for several decades. Households need to accumulate increased resources through their working lives to fund at least part of this. Furthermore, as household wealth tends to fall only slowly over retirement, more of the population will be at relatively high-wealth stages of life. This rise in the population's effective propensity to hold wealth will in turn have profound effects on the financial system, its key relative price – the real interest rate – and the prices of other assets. To the extent that these trends are stronger or weaker in different countries, they will also give rise to international payments imbalances.

We build a calibrated neoclassical overlapping generations (OLG) model to quantify the impact of these factors. We find that the ageing of the aggregate advanced-country population can explain around 75% of the roughly 210bp fall in the Holston et al. (2017) measure of natural real interest rates since 1980. Demographic pressures are forecast to reduce real interest rates a further 37bp by 2050. Furthermore, past falls in interest rates, along with the life-cycle pattern of housing demand, means that demographics can explain around three-quarters of the 50% rise in real house prices between 1970 and 2009. Given that the purchase of housing is predominantly carried out using household credit, these developments also explain the doubling of the household debt-to-GDP ratio between 1970 and the start of financial crisis.

We show that concerns about future rises in real interest rates as baby-boomers retire (see for example Goodhart and Pradhan, 2017) are largely misplaced, for two main reasons. First, it is the stock of wealth, rather than the flow of saving, that determines the interest rate in neoclassical models, and this stock falls only slowly and partially over the course of retirement. The rise in the share of population in high-wealth stages of life will therefore tend to raise the capital-output ratio even in the absence of any behavioural reaction to higher life expectancy. In contrast to Carvalho et al. (2016), therefore, we argue that while the retired may save less, they hold more wealth, and this stock pushes down on the interest rate. Second, the ongoing rise in life expectancy will, other things equal, tend to raise average wealth at any point in life as households anticipate needing to provide for a longer retirement. The ageing of the baby boom generation - a transient rise in the birth rate - merely changes the timing of these long-run life expectancy effects.

While the size and timing of the effects we find are sensitive to model calibration and specification, the mechanism underlying our model is quite general. The rise in life expectancy will tend to raise the capital-labour ratio: households save more for longer retirements and spend longer in high-wealth phases of their lives, and this extra wealth finds its way to firms and finances their capital investment. The rise in the capital-labour ratio will tend to push down on the marginal product of capital to an extent that depends on how easy it is to use extra capital in production - i.e. how rapidly diminishing the marginal product of capital is or, equivalently, how steep the investment line is in a simple saving-investment diagram. In turn, the marginal product of capital is a key determinant of the real interest rate - the only one in a simple model such as ours, and one of several in a model with risk premia and/or product markups (Eggertsson et al., 2017; Caballero et al., 2017). The fall in interest rates may encourage or dissuade further saving, thereby strengthening or weakening the effect of the original demographic shock. This depends on whether households react to more expensive retirement consumption by putting away more or less money for it when working - i.e. whether the savings function slopes upward or downward.

We use our model to study the role that housing plays in mediating these effects. In our model, households save in order to smooth consumption over the life cycle and accumulate resources for bequests. If productive capital is the only store of value, all of the burden is placed on capital to meet any increase in desired wealth holdings. So the presence of alternative stores of value can affect the impact of demographics on the interest rate. In particular, an asset, such as land, that yields a positive dividend and does not depreciate, will prevent the interest rate from going negative in a frictionless model such as ours. We show that, in practice, the presence of housing attenuates the fall in interest rates induced by demographic change, although the effects appear to be quantitatively small in our baseline calibration.

We also look at the effect of introducing imperfect competition into the model, following several recent papers that have proposed falling competition as a potential driver of the real interest rate (Philippon and Jones, 2016; Eggertsson et al., 2017). The introduction of a rising markup helps to rationalise a rising profit share alongside a fall in real interest rates and the marginal product of capital (Barkai, 2016; Caballero et al., 2017). Unlike the existing literature, we show that the effect of supernormal profits depends crucially on whether claims to these profits are traded. A tradable claim on supernormal profits is another competing store of value, so once again its presence attenuates the effect of demographics on the real interest rate.

Using the heterogeneity by age and birth year, we look at the implications of demographic changes for the level and distribution of welfare across time. We find an increase in welfare across cohorts over time, which seems to be predominantly driven by increased longevity, although this effect is conceptually difficult to quantify. Other than longevity, changes in life-time consumption are the main drivers of realised utility, and seem to be detrimental to the baby-boom generation. Finally, lower interest rates help to boost welfare by help agents smooth consumption over time.

Finally, we delve further into the open-economy dimension of our model by considering the demographic transition from the perspective of each country, taking the world interest rate as given. Specifically, while the world interest responds to the aggregate demographic trend, differences in the size and speed of demographic change across countries will lead to capital flows between countries. We show that demographics explains around 30% of the dispersion in advanced-country net foreign asset positions, with observed international imbalances generally smaller in magnitude than the values predicted by our model, suggesting the presence of frictions in international capital flows. Given the difficulty of taking such a low-frequency model to time-series data, the explanatory power of the model in the cross-section gives us greater confidence in its predictions about the average level of interest rates in the global economy.

We do not claim that demographic change is the only influence on interest rates over the long run. In common with the other papers in this literature, our model is very stylised, and in particular does not include an account of the risk premia and financial frictions that may have caused the return on capital to diverge from government bond yields (see for example Marx et al., 2016). The aim of this study is rather to isolate the effect of demographic change on savings behaviour and the real return on capital. A unified approach that quantifies these effects alongside financial frictions and risk premia is beyond the scope of this study.

Our paper is one of several addressing the impact of demographic change on the real interest rate, house prices or external payments, including Gertler (1999); Domeij and Flodén (2006); Krueger and Ludwig (2007); Waldron and Zampolli (2010); Kiyotaki et al. (2011); Backus et al. (2014); Carvalho et al. (2016); Eggertsson et al. (2017); Gagnon et al. (2016); Marx et al. (2016). Relative to these papers, our main contributions are twofold. First, we focus on global demographic change and use a unified framework to determine the effect this will have on the world real interest rate and country-specific net asset positions. Second, we examine the effect that these changes have had on house prices and credit, and conversely the role that housing plays in mediating the effect of demographics on interest rates.

The remainder of this paper is structured as follows. Section 2 sets out the key demographic trends over the past few decades. Section 3 describes the model and its calibration. Section 4 shows the results of model simulations in which we incorporate the demographic trends, and considers some robustness exercises. Section 5 present two extensions: one looking at the US as a closed economy and another looking at the role of monopoly power. Section 6 concludes.

2 Demographic Trends

This section documents the key demographic trends that motivate this paper. We use data from the UN Population Statistics, which runs from 1950 to 2015, and includes projections up to 2100. The focus is on an aggregate of advanced economies.[1]

The shares of different age groups in the total population over time presents two clear patterns (see Panel (a) of Figure 1). Firstly, we see a clear rise in the share of older generations, for example with the over-80s going from 1% of the population in 1950 to around 5% in 2015, and reaching a projected 14% by the end of the century. Secondly, the effect of the baby-boom shows as a 'bulge' moving through the population, entering the 20-44 age group from the 1970s and slowly disappearing by around 2040. Based on the same data, Panel (b) of Figure 1 confirms the increasing importance of older age groups. The average age of the population shows a clear upward trend, going from below 35 in 1950 to over 40 in 2015 and almost 50 by 2100. Here, the effect of the baby-boom is to increase the slope of this line in the 1970s-2040s, but the average age remains high even once the baby-boom generation have died out of the population.

The old-age dependency ratio (henceforth OADR), defined as the ratio of over-65s to 20-64 year olds, summarizes this evolution (Figure 2). Again the clear upward trend shows the rise in the share of older generations relative to the working population. The dashed lines, corresponding to the alternative fertility scenarios in the UN projections, give an indication of the degree of certainty around the projections: even in the high-fertility scenario, the OADR increases substantially from around 15% in 1950 to over 40% in 2100. In the medium-fertility scenario the final number is closer to 55%.

Having documented these trends, we now examine their causes. Panel (a) of Figure 3 shows the growth rate of consecutive 20-24 year old cohorts over time. We can see that the period 1970-1985 saw elevated growth rates, as high as 3%, corresponding

[1] In particular, we use Western Europe (Austria, Belgium, Denmark, France, Germany, Ireland, Italy, Netherlands, Portugal, Spain, Switzerland and the UK), North America (Canada and the US), Australia, Japan, and New Zealand.

Figure 1: Ageing Population

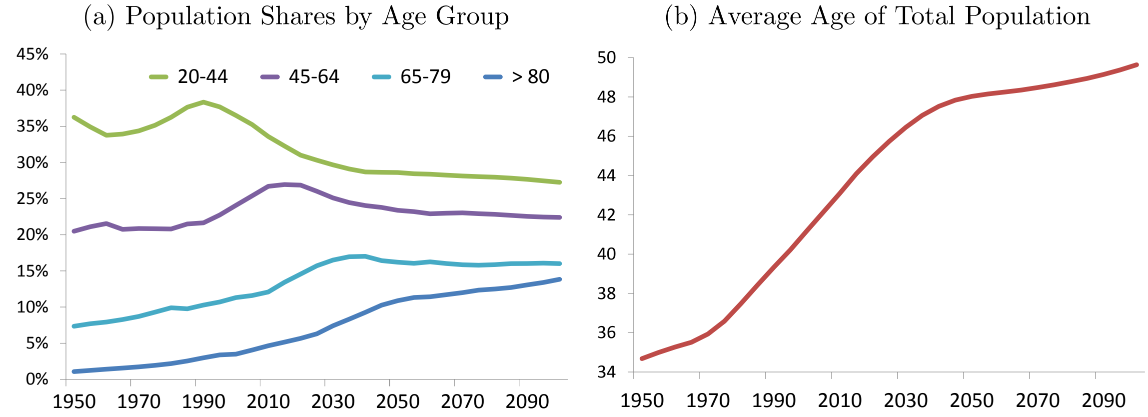

Source: UN Population Statistics (projections based on medium-fertility scenario)

to the baby-boom generations born between 1945 and 1960. Growth rates have since fallen, and even been significantly negative for several periods. Both of these affect the age structure of the population. In particular, as the large baby-boom cohorts grow old, the age distribution becomes skewed towards the older age groups. This is amplified by the smaller size of the new younger generations entering the population.

To further illustrate the baby-boom effect on the aggregate demographic trends, panel (b) of Figure 3 shows the counter-factual OADR when we assume that cohort growth in 1970-1985 was zero, hence removing the effect of the baby-boom. We can see that there is a non-negligible effect from these cohorts. When they are young, and on the denominator of the OADR, they lower this ratio relative to the counter-factual. As they get older and begin to move to the numerator of the ratio, they account for a steeper rise in the OADR. Nonetheless, once these cohorts have faded out of the population, the counter-factual OADR reaches the same high levels as the baseline projections. Hence the baby-boom does not account for the long run trend in the OADR.

The key determinant of the rise in the OADR is increasing longevity. Figure 4 shows

Figure 2: Old-Age Dependency Ratio (≥65/20-64)

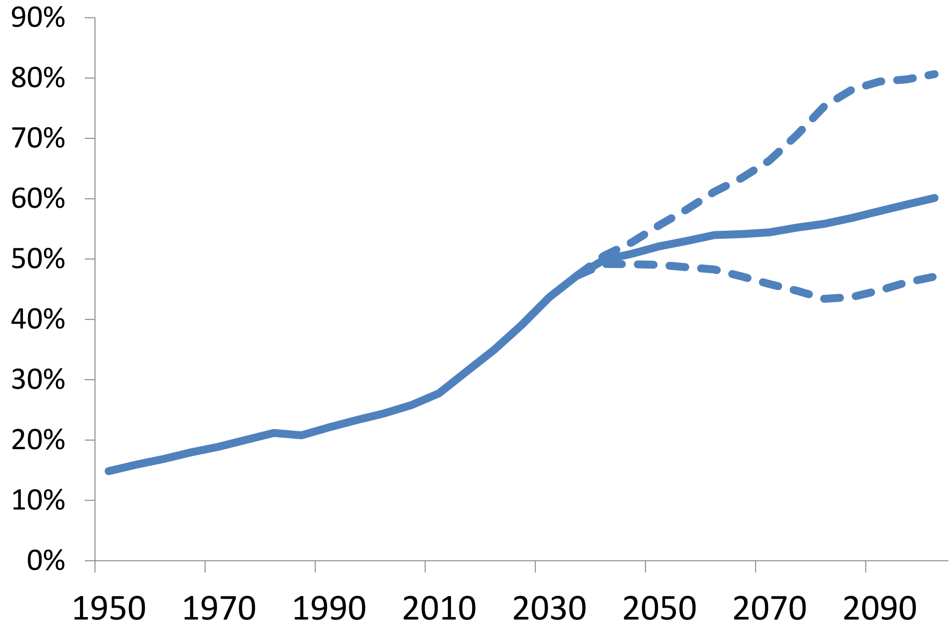

Source: UN Population Statistics (dashed lines show high- and low-fertility scenarios)

life expectancy conditional on living to age 60.[2] While a sixty year old in 1950 would not expect to live past the age of 77, by 2015 a sixty year old can expect to live until close to 85. By the end of the century this number rises past 90. As people face lower mortality rates later in life, and their life expectancy rises, older age groups account for a growing proportion of the total population.

As this data makes clear, ageing population in advanced economies has led to an unprecedented shift in the age structure of the population, and these effects will almost certainly persist for decades to come. The rest of this paper will employ an OLG model to uncover the macro-economic effects of these important trends.

[2]This measure is taken directly from the UN Population Statistics, and is defined as: "The average... years of life expected by a hypothetical cohort of individuals alive at age 60 who would be subject during the remaining of their lives to the mortality rates of a given period."

Figure 3: The Baby Boom

(a) Annual Growth Rate of Cohort Aged 20-24 (b) OADR Counterfactual

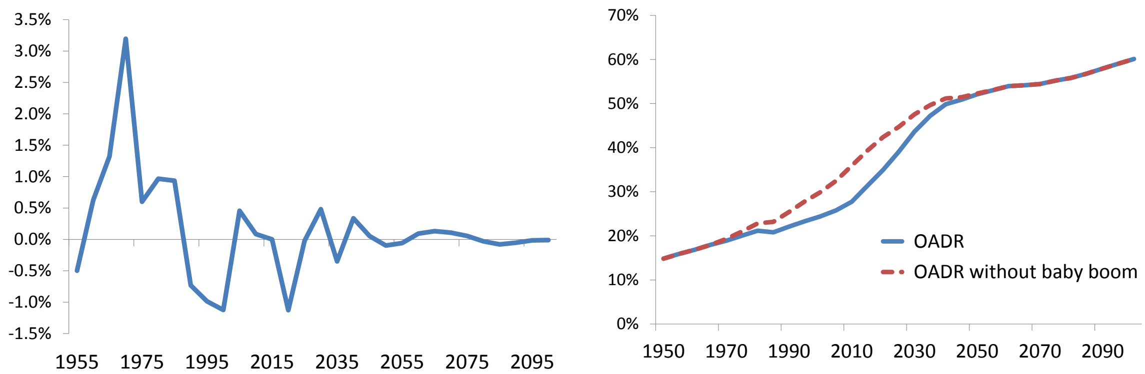

Source: UN Population Statistics (projections based on medium-fertility scenario)

3 Quantitative Model

We now present the quantitative model and its calibration. Results of the dynamic simulations based on this model will be presented in Section 4.

3.1 The Model

Our general equilibrium set-up includes overlapping generation households and a representative firm producing, in the first instance, in a perfectly competitive environment. We describe the two groups of agents in turn, before turning to issues of aggregation and market clearing.

Household

Agents are born at age 1 and can live up to T periods. They work from their first period of life until they reach retirement, at a fixed age T^r. They face a probability of death at (after) each age τ, denoted $(1-\psi_{\tau,t})>0$, and die with certainty at the

Figure 4: Life Expectancy at 60

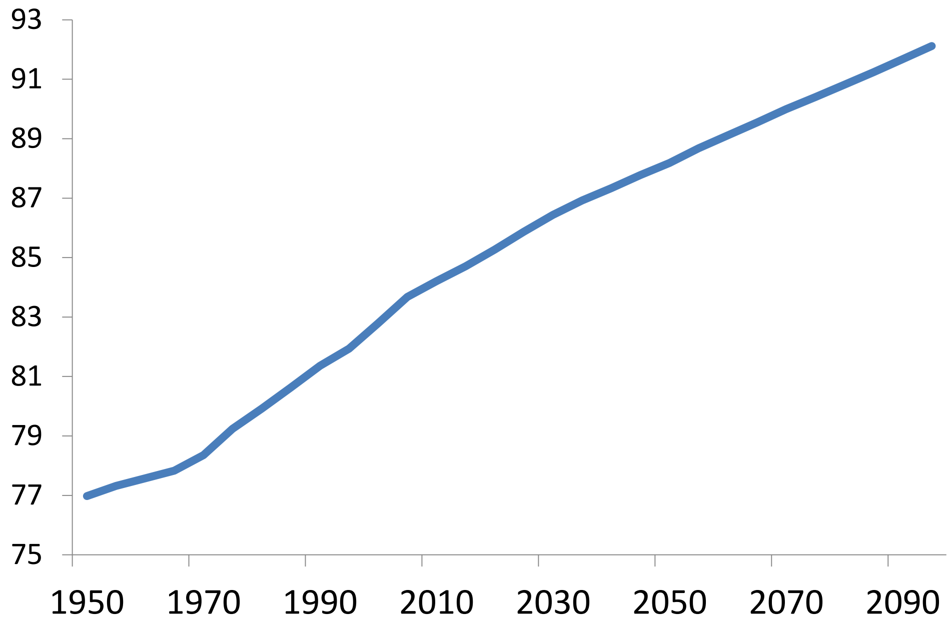

Source: UN Population Statistics (projections based on medium-fertility scenario)

maximum age, T, hence $\psi_{T,t} = 0\,\forall t$. This can be translated into the probability of surviving until each age, $\tilde{\psi}_{\tau,t} = \Pi_{j=1}^{\tau-1}\psi_{j,t}$, with $\tilde{\psi}_{1,t} = 1$.

Throughout their life, agents supply labour, l, inelastically, and gain utility from a consumption good, c, and a housing good, h, which is bought and sold at relative price p^h.

We denote by $x_{\tau,t}$ the value of a variable x, for a household born in period t, when they are aged τ. Hence the T-period optimisation problem faced by a representative household born in period t can be written as

$$\max_{\{c_{\tau,t},\, a_{\tau,t},\, h_{\tau,t}\}_{\tau=1}^{T}} \quad \sum_{\tau=1}^{T} \beta_\tau \tilde{\psi}_{\tau,t} \left(\ln c_{\tau,t} + \theta_\tau \ln h_{\tau,t}\right) + \beta_T \tilde{\psi}_{T,t} \phi \ln a_{T,t}$$

subject to

$$c_{\tau,t}+a_{\tau,t}+p^h_{t+\tau-1}(h_{\tau,t}-h_{\tau-1,t}) \leq w_{t+\tau-1}\epsilon_\tau l_{\tau,t}+(1+r_{t+\tau-1})a_{\tau-1,t}+\pi_{\tau,t} \qquad \text{for } \tau = 1, ..., T$$

where l is the inelastic labour supply, ϵ is the age-specific productivity level, w is the wage per efficiency units of labour, and a is a safe asset with return r.[3] We assume that $l_{\tau,t} = \epsilon_\tau = 0$ for $\tau \geq T^r$.

Agents are born without any assets, that is $a_{0,t} = 0$, but we allow the possibility of bequests, setting $\phi > 0$ so that $a_{T,t} > 0$. These bequests are distributed among subsequent generations as part of $\pi_{\tau,t}$, which captures all non-labour income, taken as exogenous by the households.

There are a fixed number of periods when the household is able to "move house", i.e. re-optimise their housing wealth, and hence outside of these "move dates" the household has an additional constraint $h_{\tau,t} = h_{\tau-1,t}$. We assume that agents are born without any housing wealth, and do not leave any housing wealth when they die, hence $h_{0,t} = h_{T,t} = 0$, which necessitates $\theta_T = 0$.

Denoting by $\lambda_{\tau,t}$ the Lagrange multiplier on the budget constraint at age τ, this problem gives rise to the following first order conditions:

$$\lambda_{\tau,t} = \beta_\tau \tilde{\psi}_{\tau,t} c^{-1}_{\tau,t} \qquad \forall \tau = 1, \ldots, T \qquad (1)$$

$$\lambda_{\tau,t} = (1 + r_{t+\tau})\lambda_{\tau+1,t} \qquad \forall \tau = 1, \ldots, T-1 \qquad (2)$$

$$\lambda_{T,t} = \beta_T \tilde{\psi}_{T,t} \phi a^{-1}_{T,t} \qquad (3)$$

$$\sum_{j=\tau}^{\tau'-1} \beta_j \tilde{\psi}_{j,t} \theta_j h^{-1}_{\tau,t} = p^h_{t+\tau-1}\lambda_{\tau,t} - p^h_{t+\tau'-1}\lambda_{\tau',t} \qquad \forall \tau \in \text{"move dates"} \qquad (4)$$

where τ' in (4) denotes the next move date after τ.

[3] Note that $t + \tau - 1$ is the period in which the generation born at time t is aged τ.

Firm

The firm's problem is to choose the aggregate factors of production, K_t and L_t, to maximise profit, taking as given the rental rate of capital, r_t^k, the wage per efficiency units of labour, w_t, and the production function, $Y = F(K, L)$. Note that L_t denotes the aggregate efficiency units of labour supplied by households. This problem can be written as

$$\max_{L_t, K_t} \quad F(K_t, L_t) - w_t L_t - r_t^k K_t$$

Taking the CES production function $F(K, L) = Z\left[(1-\alpha)L^{\frac{\sigma-1}{\sigma}} + \alpha K^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}$, we have the following first order conditions

$$\begin{aligned} w_t &= (1-\alpha)Z^{\frac{\sigma-1}{\sigma}}\left(\frac{Y_t}{L_t}\right)^{\frac{1}{\sigma}} \\ r_t^k &= \alpha Z^{\frac{\sigma-1}{\sigma}}\left(\frac{Y_t}{K_t}\right)^{\frac{1}{\sigma}} \end{aligned}$$

Capital is financed from the households savings and depreciates at rate δ every period. Before paying for the capital rental rate, the firm is left with $(1 - \delta + r_t^k)K_t$ at the end of each period t and the households receive an interest rate r_t on their savings, hence the zero profit condition of the firm implies

$$r_t^k = r_t + \delta$$

Aggregation

We denote the gross growth rate of the generation born at time t relative to the generation born at time $(t-1)$ with g_t. Normalising the size of the generation born at time 0 to 1, this means the size of the generation born at time t can be written as

$$s_t = g_t s_{t-1} = \Pi_{i=1}^t g_i$$

At each age, the size of the cohort reduces, with survival probability $\tilde{\psi}_{\tau,t} \leq 1$. Hence the total population in period t is given by

$$S_t = \sum_{\tau=1}^{T} \tilde{\psi}_{\tau,t-\tau+1} s_{t-\tau+1} = \sum_{\tau=1}^{T} \tilde{\psi}_{\tau,t-\tau+1} \Pi_{i=1}^{t-\tau+1} g_i$$

Let $\mathbf{x}_t$ denote the (Tx1) vector of a variable x, for one representative household of each generation alive at time t, in other words $\mathbf{x}_t = \{x_{\tau,t-\tau+1}\}_{\tau=1}^{T}$. Let ρ_t denote the (Tx1) vector of population sizes at time t, that is $\rho_t = \{\tilde{\psi}_{\tau,t-\tau+1} s_{t-\tau+1}\}_{\tau=1}^{T}$. The aggregate value of variable x at time t is denoted by $X_t = \rho_t' \mathbf{x}_t$. We denote by $\tilde{X}_t$ the value of X_t per aggregate capita, that is X_t/S_t. We can write this as $\tilde{X}_t = \tilde{\rho}_t' \mathbf{x}_t$ where $\tilde{\rho}_t = \rho_t/S_t$ denotes the vector of *relative* population sizes.

Market Clearing

Capital/Savings Market The value of the capital stock must equal the aggregate savings of the previous period

$$A_{t-1} = K_t$$

As introduced above, we denote per capita capital stock as $\tilde{K}_t = K_t/S_t$. For consistency, this implies that per capita savings are defined relative to next period's population, that is $\tilde{A}_t = A_t/S_{t+1} = \frac{S_t}{S_{t+1}} \tilde{\rho}_t' \mathbf{a}_t$.

Labour Market Aggregate labour supply must equal labour demand. Let $\epsilon \mathbf{l}_t = \{\epsilon_\tau \, l_{\tau,t-\tau+1}\}_{\tau=1}^{T}$ denote the vector of efficiency units of labour supplied by each generation at time t. Then

$$\rho_t' \epsilon \mathbf{l}_t = L_t \quad \Rightarrow \quad \tilde{\rho}_t' \epsilon \mathbf{l}_t = \tilde{L}_t$$

Housing Market As with the household savings, for consistency we define per capita housing relative to next period's population, that is $\tilde{H}_t = H_t/S_{t+1} = \frac{S_t}{S_{t+1}} \tilde{\rho}_t' \mathbf{h}_t$. Housing is effectively residential land, in that its supply is inelastic, hence we assume that the housing stock per capita is fixed at some level, $\tilde{H}$.[4] Market clearing then simply requires

$$\tilde{H}_t = \tilde{H} \quad \forall t$$

This implies that the aggregate housing stock, H_t, grows with the population, meaning that the economy is endowed with an additional $(\frac{S_{t+1}}{S_t} - 1)\tilde{H}$ units of housing each period.[5] This endowment is distributed across households through non-labour income, along with the bequests, as detailed below.

Bequests and non-labour income At each time t, the non-housing assets of the generations that died in the previous period must be distributed, along with the accrued interest, to living households through bequests, B_t, given by

$$B_t = (1 + r_t) \sum_{\tau=1}^{T} (1 - \psi_{\tau,t-\tau}) \tilde{\psi}_{\tau,t-\tau} s_{t-\tau} a_{\tau,t-\tau}$$

$$\tilde{B}_t = (1 + r_t) \frac{\sum_{\tau=1}^{T} (1 - \psi_{\tau,t-\tau}) \tilde{\psi}_{\tau,t-\tau} s_{t-\tau} a_{\tau,t-\tau}}{S_t}$$

Similarly, the housing wealth of the agents that died in the previous period must be distributed among remaining agents. This is aggregated analogously to savings above

$$\tilde{B}^h_t = \frac{\sum_{\tau=1}^{T} (1 - \psi_{\tau,t-\tau}) \tilde{\psi}_{\tau,t-\tau} s_{t-\tau} h_{\tau,t-\tau}}{S_t}$$

The additional housing endowment, added in each period to maintain a stable level of housing per capita, is added to the aggregate asset and housing bequests to form aggregate non-labour income, Π_t. In other words

$$\tilde{\Pi}_t = \tilde{B}_t + p^h_t \tilde{B}^h_t + p^h_t \left(\frac{S_{t+1}}{S_t} - 1 \right) \tilde{H}$$

[4] This is in line with Knoll et al. (2014), who find that the bulk of the increase in house prices is attributable to the increase in the value of residential land.

[5] The alternative would be to allow this additional housing to be produced, with a technology which transforms the consumption good into housing. While this assumption does not materially affect our results, it is more in line with the interpretation of housing as being in fixed supply.

This non-labour income is evenly distributed among households above a given age, T^b, while younger households are not entitled to any non-labour income. This assumption is aimed to reflect the fact that older households are more likely to see their family members die and to inherit their assets and housing wealth. Furthermore, a flat bequest distribution across households above this age ensures that bequests do not create strong distortions on the household consumption and saving choices.

Goods Market Aggregating the budget constraints of all households alive at a given time t, and substituting the equilibrium conditions described above, gives us the familiar resource constraint

$$\tilde{Y}_t = \tilde{C}_t + \tilde{I}_t \tag{5}$$

where $\tilde{I}_t$ is the net increase in aggregate savings, given by

$$I_t = A_t - (1-\delta)A_{t-1} \quad \Rightarrow \quad \tilde{I}_t = \frac{S_{t+1}}{S_t}\tilde{A}_t - (1-\delta)\tilde{A}_{t-1}$$

Hence the resource constraint (5) simply implies that all goods produced at time t are either consumed or saved as capital.

3.2 Calibration

Each period in the model represents 5 years. We assume that working life begins at age 20 and no agents live beyond age 90, setting $T = 14$.

The focus of our calibration will be to match life-cycle profiles of labour productivity, housing wealth and net worth, as well as aggregate housing wealth-to-GDP, debt-to-GDP and real interest rates. All of these moments will vary over time in the dynamic transition path due to the demographic trends. Given the data availability, we target average life-cycle patterns for the years 1990-2010. For the aggregate moments, we target their average values over the 1970s, in order to allow the model to determine the transition over the past few decades.

3.2.1 Data

Life-cycle Profiles

Given limited cross-country data availability, we will assume that US households are representative of all advanced-economy households in terms of the life-cycle profiles of labour productivity, housing wealth and net worth. Hence, we can use the Survey of Consumer Finance (SCF) to match life-cycle profiles for productivity, ϵ, net worth, a, and housing wealth, h.

Specifically, we calibrate productivity to match *"Wage Income"* data from the SCF, which corresponds to total labour income, irrespective of hours worked. Hence, since hours worked are inelastic in the model, we are effectively subsuming all life-cycle hours and wage decisions into the productivity profile. The estimated labour income profile falls close to zero from around age 65, and in fact *median* wage income is exactly zero from the 65-70 age group. Hence we assume retirement begins at age 65, that is $T^r = 10$. To calibrate housing wealth over the life-cycle, we take the sum of *"Primary Residence"* and *"Other Residential Real Estate"* in the SCF. The SCF includes a measure of *"Net Worth"* that aggregates all financial and non-financial assets and liabilities: to ensure that the profile of total net worth in the model corresponds to this observed net worth, we calibrate non-housing assets, a, to match the SCF *"Net Worth"* minus housing wealth as defined above. Note that, in this way, housing wealth measures only housing *assets*, and any *debt* related to housing, such as mortgages, are included in other assets, a.

To create the life-cycle profiles for each of these variables, we put the survey respondents into 5-year age buckets corresponding to the life-cycle of households in the model, and calculate the average level of each variable for each age group, using the sampling weights provided in the SCF. This gives us an estimated life-cycle profile for each survey year from 1989-2013. We then take the average over the survey years, weighting by the number of observations in each age group in each survey year.[6] This procedure gives us an estimated life-cycle profile for each of the three variables,

[6] Using the coefficients on the age group in fixed effects panel regressions yields similar results.

corresponding to the average cross-sectional age profile between 1989-2013.

Aggregate Variables

We take three aggregate variables as targets: the real interest rate, housing wealth-to-GDP and debt-to-GDP. In order to allow the model to determine the evolution of these variables over the last few decades, we target their average value in the 1970s.

For the real interest rate we use the data from King and Low (2014), and take the average world interest rate between 1970-1980. This gives us a target of 3.7%.

The data from Piketty and Zucman (2014) measure housing assets, including land, and give us the aggregate housing wealth-to-GDP target. We take an average over the 1970s for all available countries, namely Australia, Canada, France, Germany, Italy, Japan, the UK and the US. We obtain a target ratio of 145%.

Finally for debt-to-GDP we use the BIS Total Credit data, focusing on total credit to households as a percentage of GDP. Again we use the average over the 1970s for the countries available, in this case Canada, Germany, Italy, Japan, the UK and the US. The final target is 35%.

3.2.2 Calibration Procedure

We abstract from TFP growth, normalising $Z = 1$. Although changes in TFP growth over time may partly explain changes in real interest rates, we abstract from this in order to focus on the role of demographics. In the next section, instead, we will look at how demographic changes affect labour productivity. We set the parameters of the CES production function $\sigma = 0.7$ and $\alpha = 1/3$, and the annualised depreciation rate $\delta = 6\%$.

In order to set the parameters of the household's problem to match our moments, we consider a steady state of the model in which all households have the same demographic characteristics as the 1945 cohort. For a given calibration of this steady state, we run the dynamic simulations and back out the implied average life-cycle profiles in 1990-2010, and aggregate moments in the 1970s. We then adjust the calibration of the steady state until the moments that come out of the dynamic simulation match those in the data.

The steady state of the model gives us a stationary vector of relative population weights, $\tilde{\rho}$, with which we normalise the productivity profile such that aggregate labour productivity is 1, that is $\tilde{\rho}'\epsilon = 1$. We set hours worked at 0.3 throughout working life, hence $l_\tau = 0.3$ for $\tau = 1, \ldots, T^r - 1$, and $l_\tau = 0$ for $\tau \geq T^r$. Hence aggregate labour supply is $L = 0.3$, the value commonly used in the literature. The wage is then set as the marginal product of labour consistent with the firm's first order condition with respect to labour.

We normalise the life-cycle profile of assets, a, such that aggregate wealth is consistent, using the firm's first order condition with respect to capital, with the annualised interest rate target, and debt, in the first periods of life, is consistent with the debt-to-GDP target. Since assets in the final period of life are non-zero, we set $\phi > 0$ to satisfy (3) for the observed level of a_T.

Finally, we normalise housing wealth over the life-cycle such that the aggregate housing stock, $\tilde{H}$, is consistent with the housing wealth-to-GDP target. As mentioned above, we do not allow households to re-optimise their housing wealth in every period, and correspondingly, we use a step-wise function to fit the estimated life-cycle profile. Since this profile is found to be significantly above zero in both the first and last age groups, we set both $\tau = 1$ and $\tau = T$ as "move dates" in the household's problem. In order to match the observed peak in housing wealth in middle age and subsequent fall at around age 70, we allow $\tau = 5$ and $\tau = 11$ to also be "move dates". For simplicity, we set θ_1, θ_5, θ_{11} and θ_T to satisfy the first order condition with respect to housing, (4), with $\theta_\tau = 0$ for all other τ.

For the final step of the calibration, for a given life-cycle profile of labour and non-labour income, housing wealth and net assets, the steady state budget constraint gives consumption over the life-cycle

$$c_\tau = wl_\tau\epsilon_\tau + (1 + r)a_\tau - a_{\tau-1} - p^h(h_\tau - h_{\tau-1}) + \pi_\tau$$

Following Glover et al. (2014), we set $\beta_1 = 1$ and calibrate β_τ, $\tau > 1$, such that the Euler equations, in (1) and (2), are satisfied given this stream of consumption

$$\beta_\tau = \frac{\beta_{\tau-1}}{1+r} \frac{\tilde{\psi}_{\tau-1}}{\tilde{\psi}_\tau} \frac{c_\tau}{c_{\tau-1}}$$

3.2.3 Calibration Outcomes

Table 1 shows the average real interest rate, housing wealth-to-GDP ratio and debt-to-GDP ratio from the model's dynamic transition path for the period corresponding to 1970-1980. We see that the model matches the targeted moments well.

Table 1: Aggregate moments targeted for the calibration

	Model	Data
Housing wealth to GDP (1970s)	145%	145%
Household debt to GDP (1970s)	35%	35%
World natural interest rate (1970s)	3.70	3.70

Figure 5 shows the life-cycle profiles for productivity, housing and net worth from the data and the model. For the model, given that these profiles change over time in the transition, we take the equivalent of the estimates from the data, namely the average of the cross-sectional age-profile of each variable over 1990-2010.

4 Results

We now present the results of the model simulations. We begin by describing the demographic processes that we insert exogenously into the model. We show the resulting transition of the main macroeconomic variables of interest, namely the real interest rate, savings and debt, and decompose these results in terms of the changes in the age distribution of the population and changes in the savings behaviour of households as they expect to live longer. We then explore other macroeconomic implications of demographic trends: on the housing market, labour productivity,

Figure 5: Calibration of Life-cycle Profiles

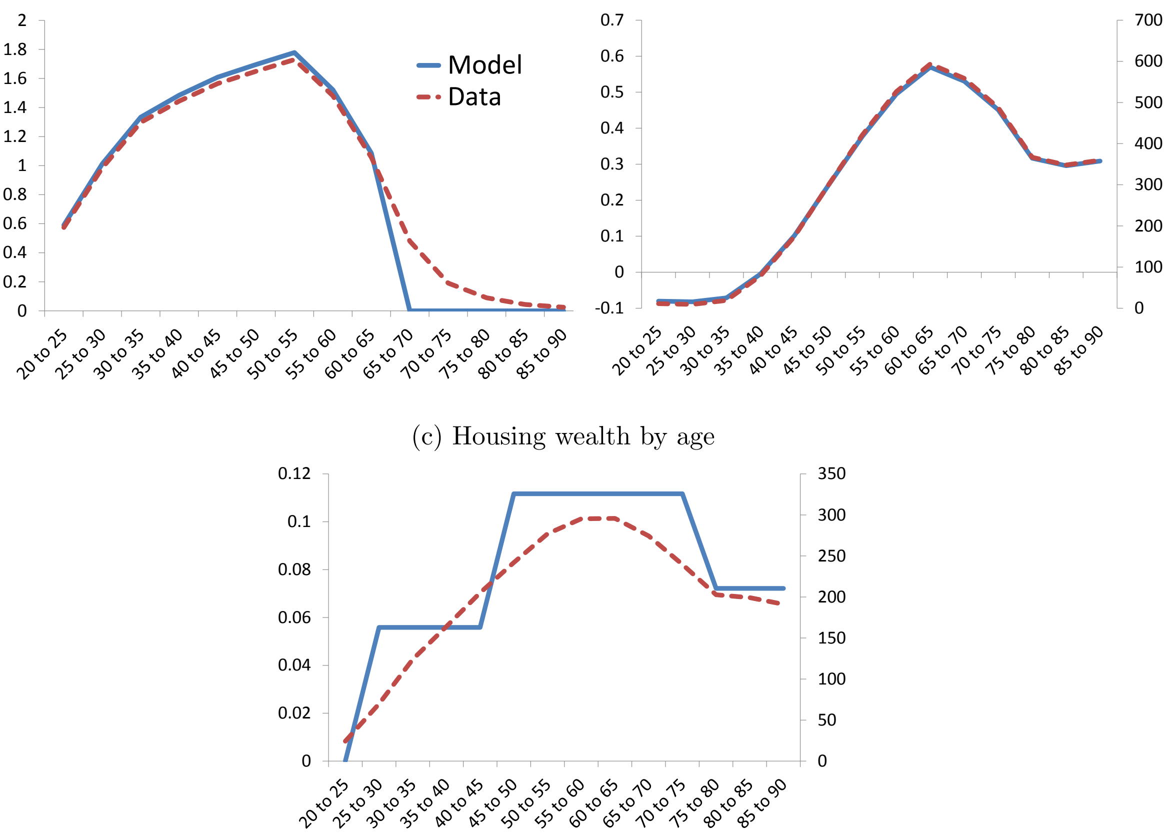

(a) Labour productivity by age

(b) Net worth (excluding housing) by age

(c) Housing wealth by age

welfare and distribution, and in an open economy context. Finally, we carry out some robustness exercises looking at the retirement age and the role of housing in the model.

4.1 Exogenous Demographic Shocks

Population growth, g_t, and the survival probabilities, $\psi_{\tau,t}$, are the exogenous demographic processes that drive fluctuations in our model. Using the data described and shown in Section 2, we set these two series so as to match the evolution of the age structure of the economy from the 1950s, and projected until 2100.

Figure 6: Demographics in the Model vs Data

(a) Old-Age Dependency Ratio (b) High Wealth Ratio

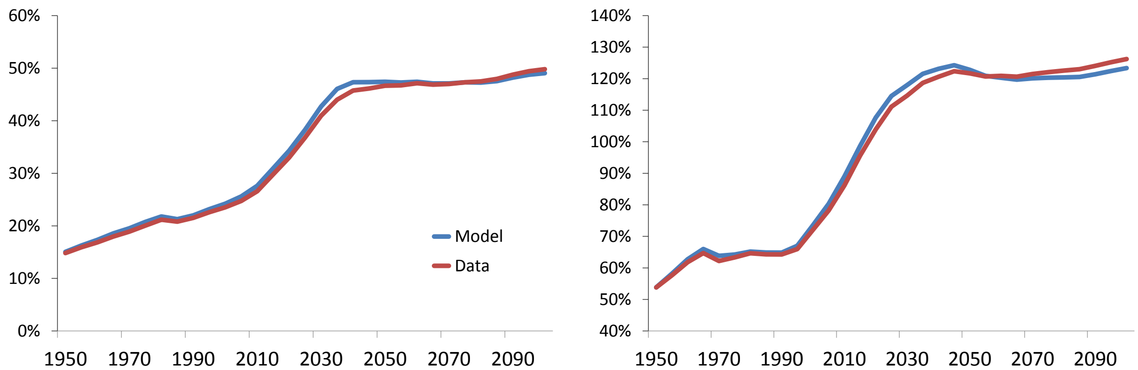

Source: UN Population Statistics and own calculations.

Specifically, we set g_t as the relative size of consecutive 20-24 year old cohorts over time. We then set $\psi_{\tau,t}$ to match the observed evolution of each cohort throughout their life, meaning that the rate of decline in the size of a given cohort from one period to the next is taken to be the death rate.[7]

To show how well we fit the demographic trends with these two exogenous series, panel a) of Figure 6 shows the OADR of the model against the data. Panel b) plots a slightly different ratio, which we call the high-wealth ratio (HWR). To define this ratio, we use the empirical life-cycle profile of assets to define the 'high-wealth' phase of an agents life. As can be seen from panel b) of Figure 5, agents have accumulated a large amount of wealth by around the age of 50-55, and maintain that level of wealth until the end of their life at the age of 90. Hence we define the HWR as the ratio of those over 50 to those below 50.

Looking at Figure 6, we see that, despite the slight simplifications that we make, both the OADR and the HWR in the model are very close to that in the data.

[7] The existence of immigration means that cohort sizes can go up as well as down over time, particularly for younger age groups. To remove this possibility, we smooth the death rate before retirement to match the overall decline of a given cohort between the ages of 20 and 64. If a cohort size is higher at the age of 64 than at the age of 20, which is the case for more recent years, we assume a zero probability of death before retirement.

Figure 7: The Real Interest Rate

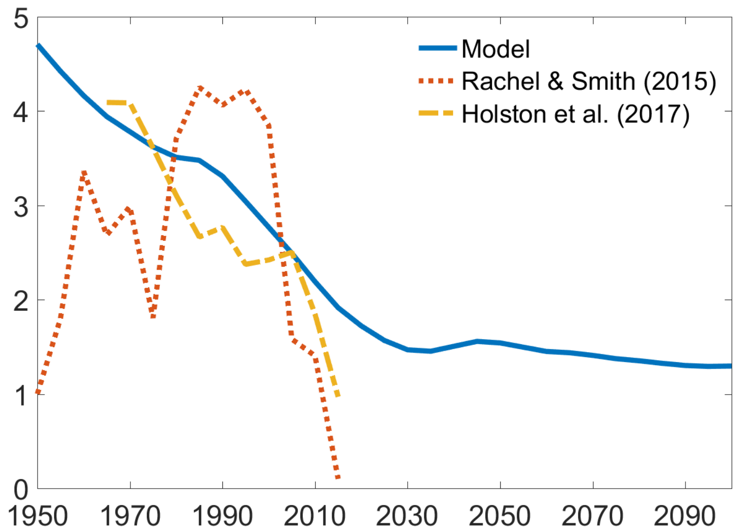

Holston et al. (2017) measures the natural interest rate, and is averaged across the US, UK, Euro-area and Canada. Rachel and Smith (2015) measures of the world real interest rate.

4.2 Baseline Results

Given the exogenous demographic changes described above, we solve for the general equilibrium transition path of the economy, assuming perfect foresight.

4.2.1 The Interest Rate, Savings and Debt

Figure 7 shows the main outcome of our model, regarding the real interest rate, compared to two empirical counterparts. The red dotted line shows the measure of the world real interest rates, taken from Rachel and Smith (2015). This measure of realised interest rates is clearly more volatile than the neutral interest rate that is captured in our model. The dash-dotted yellow line shows the model-based estimate of the natural interest rate from Holston et al. (2017). In the model, the annual interest rate decreasing by 160bp between 1980 and 2015, and forecast to decrease by a further 60bp by 2100. Compared to the average natural interest rate evolution between 1980 and 2015, demographics are able to explain 75% of the roughly 210bp drop estimated by Holston et al. (2017), and around 45% of the fall in the Rachel and Smith (2015) measure. While it is true that the actual world interest rate has fallen by more than is predicted by our model through demographic changes, leaving room for other more transitory explanations of the current low level of interest rates, it is important to note that the demographic changes themselves do not reverse, and leave the economy with a permanently lower natural interest rate. In the transition path, it is still possible to see the transitory impact of the baby-boom, slowing down the interest rate decrease in the 1990s and accelerating it between 2010 and 2040. However, in the long run, the main driver behind the transition path is the increase in life expectancy, as mentioned in Section 2, and this trend is projected to persist.

The converse of the fall in the real interest rate is a rise in the capital stock, shown in Figure 8a against two empirical measures of capital intensity: an index of the ratio of capital services to GDP for 19 OECD countries, in the red dotted line, and an index of the ratio of the capital stock to value added in the US business sector from the Fernald (2014) data for the US, in the yellow dash-dotted lines. The demographic trends alone slightly over-estimate the rise in capital, particularly for the US, as we abstract from other factors that drive investment. Nonetheless, both measures of capital intensity show a rise, particularly in the period of interest since the 1980s. This gives credence to the notion that part of the fall in the real interest rate is the standard neoclassical effect of a rise in the capital intensity of the economy. At the same time, as seen in Figure 8b, household debt to GDP has been increasing, in line with the data.[8] The investment to GDP ratio, however, does not increase monotonically: instead it smoothly follows the population growth patterns, especially for the baby-boom, as seen in Figure 9. This Figure further highlights that it is the HWR, rather than the OADR, that drives these changes.

The key mechanism triggered by the demographic transition is the following. Firstly, households anticipate that they will live longer and spend more time in retirement. They are therefore willing to transfer more of their income during working life to the future, in order to smooth their consumption. Secondly, the slower population

[8]Household debt are from the BIS databases. The household debt ratio is the ratio of household debt to nominal GDP.

Figure 8: Capital and Debt

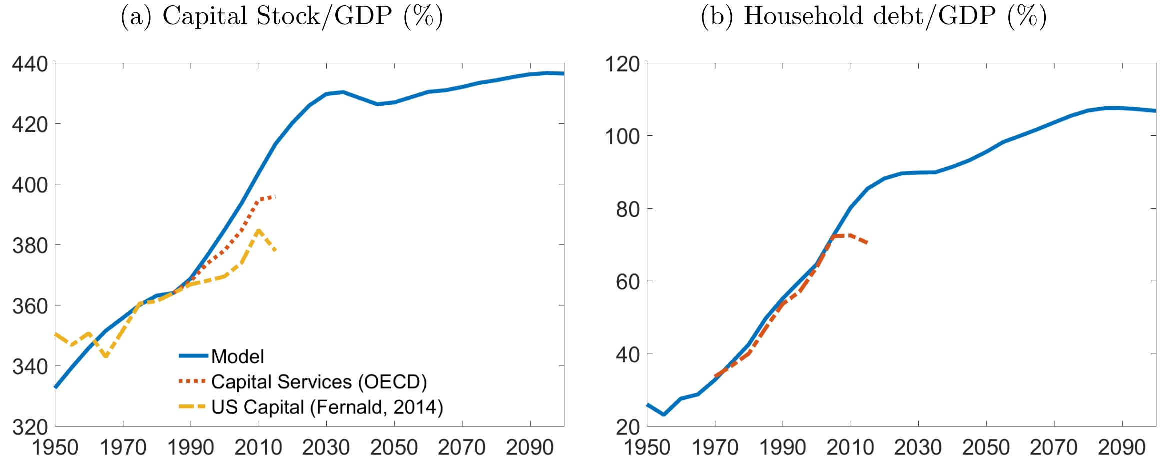

growth and increased longevity imply that older households make up a larger share of the total population alive at each period. These two changes both increase the level of aggregate savings to GDP over time. To keep the capital market balanced given this higher capital supply, the interest rate decreases. The lower interest rate also has an offsetting effect as it encourages more borrowing by the young, raising net household debt/GDP, and pushing down on aggregate savings/GDP.

We can gain a deeper understanding of the transition path by carrying out two simple exercises. Firstly, we can decompose the changes in aggregate savings into the two distinct drivers: changes in the age composition of the population and changes in the life-cycle savings of each household. Due to both the increase in life expectancy and the lower population growth, the weights of the different age groups in the total population change, with the share of older households increasing over time (cf. Figure 1). Were households not anticipating their longer life-time and adjusting their consumption and saving choices, this change in weights would still imply different aggregate savings and aggregate consumption outcomes. Second, given the perfect foresight assumption, each household cohort anticipates their longer expected life-time, and adjust their consumption, housing and saving decisions over their life-cycle.

The impact of these two distinct drivers on aggregate savings and household debt

Figure 9: Impact of Demographics on Investment

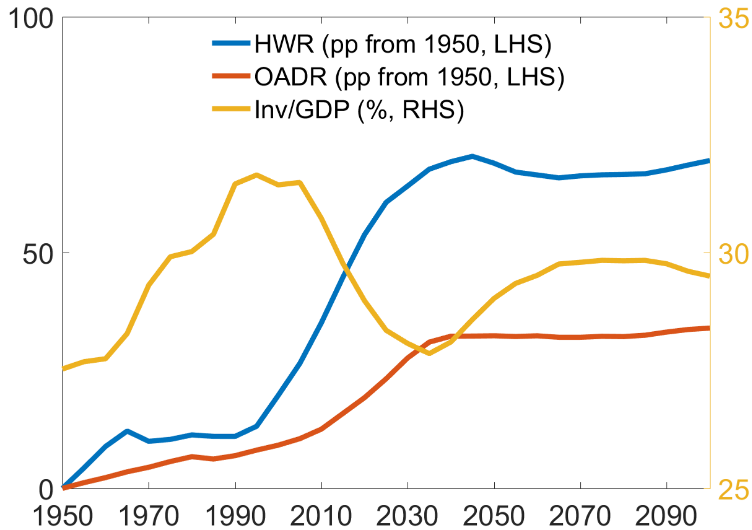

is shown on Figure 10.[9] The marginal impact of changes in the population age structure (shown in the red line) on aggregate savings per capita tends to be larger than the baseline: indeed, it only takes into account the smaller (resp. larger) share of younger (resp. older) households in total population. Since older households hold more assets, increasing their share in the economy drives up the level of aggregate savings per capita. Similarly, only younger households are indebted, so that the aggregate household debt over GDP decreases with the share of young households in the economy.

Conversely, taking only changes in optimisation over the life-cycle into account, the aggregate savings per capita actually decrease massively from 1990 onwards: since

[9] We proceed as follows to obtain this decomposition. The baseline aggregate variables are calculated as weighted sums of the alive cohorts' per capita variables. These aggregate variables are therefore driven by the interaction of the changes in the weights of each cohort in the total population and the changes in the individual housing and saving decisions of the alive cohorts. To compute the impact of population weights only, we re-calculate the aggregate variables keeping fixed the alive cohorts' per capita variables at their 1950 level. On the opposite, to obtain the impact of optimisation decisions only, we re-calculate the aggregate variables keeping the weights fixed at their 1950 level. The results shown in Figure 10 are therefore *not* the outcome of a general equilibrium transition, but an ex-post decomposition. Finally, as population weights and individual household decisions multiply each other to obtain the baseline aggregate variables, it is normal that the separate effects of these two drivers do not add up to the baseline path.

Figure 10: Decomposing the Drivers

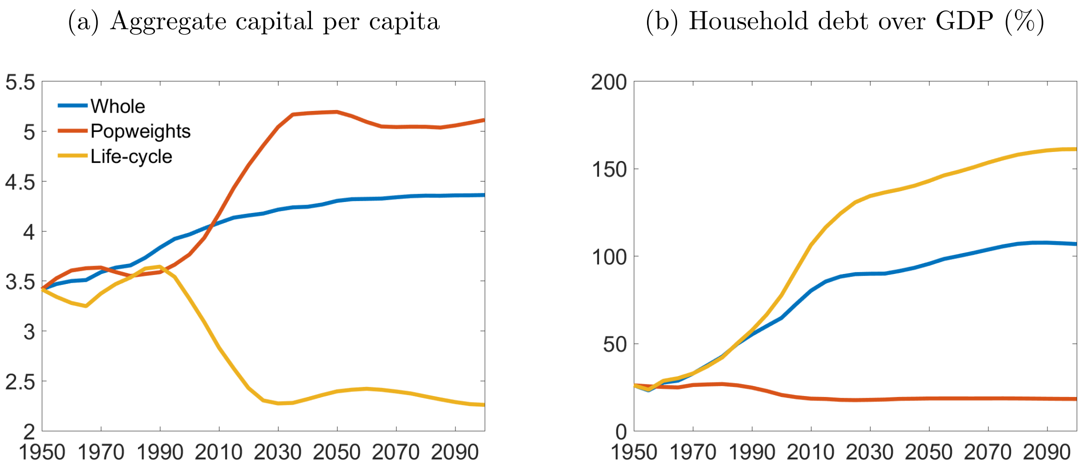

Popweights is changing only the population age structure and Life-cycle is changing only the household's optimal behaviour.

the interest rate is lower, households shift their portfolio towards consuming more, holding more housing wealth, and more debt when young. Without the offsetting effect of the falling share of the increasingly-indebted young in the population, this leads to a fall in aggregate savings.

The second exercise we can carry out is to decompose the change in the life-cycle savings profile into the partial equilibrium effect of increased longevity, holding all prices constant, and the general equilibrium impact of the change in interest rates and house prices. This is shown in Figure 11, where the blue dashed lines and yellow dash-dotted lines show the general equilibrium savings and housing-wealth profiles of the 1980 and 2015 cohorts respectively, and the red solid lines show the partial equilibrium optimal savings and housing-wealth profile of the 2015 cohort if prices were held fixed at their 1980 levels.

Without any price adjustment, the 2015 cohort, which has a higher life expectancy, decides to save more and hold less housing than the 1980 cohort (comparison between the blue dashed and the red solid curves). When prices do adjust, however, the equilibrium interest rate is lower in 2015, and for the 2015 cohort saving is less

Figure 11: Counterfactual wealth and housing profiles

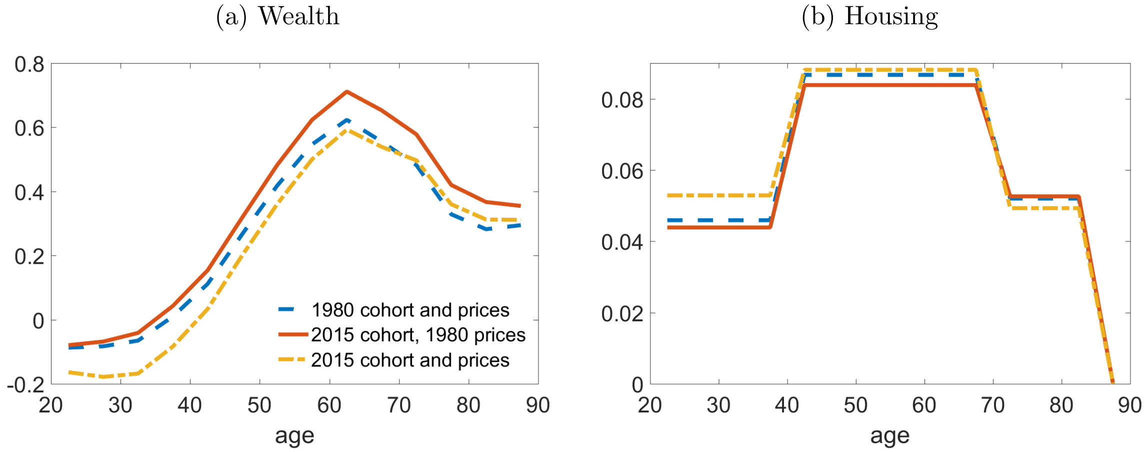

attractive and borrowing in youth is more attractive, so that the desired non-housing wealth of that cohort is lower than that of the 1980 cohort (comparing the blue dashed and the yellow dashed-dotted curves). Conversely, as house prices rise, housing wealth is higher for the 2015 cohort.

4.2.2 Housing

One important feature of our model compared to the literature is the presence of housing. Households directly derive utility from housing, but housing also serves a second purpose, as households can use it as an additional way of transferring wealth over time, in that it is durable and can be sold to fund consumption and bequests. In our framework, households have perfect foresight, which allows them to anticipate the evolution of the housing price over their life-time, and housing holdings are therefore a store of wealth over the life-cycle. As the interest rate falls, the user cost of housing falls, and so demand for housing rises. With the supply of housing per capita held fixed in our model, housing prices are pushed up, and, as a consequence, housing wealth to GDP ratio increases, as shown in Figure 12. In fact, we are able to explain

Figure 12: Housing in the Baseline Simulations

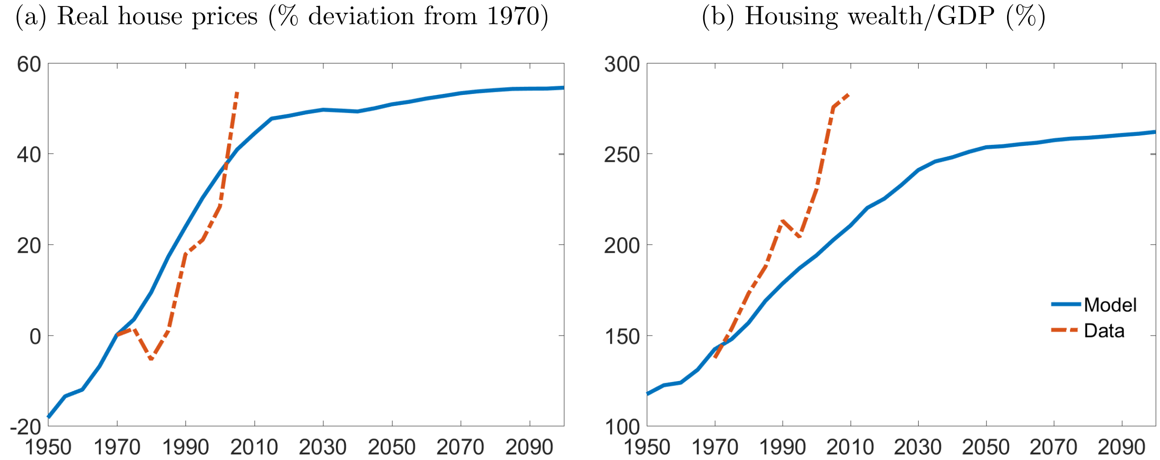

85% of the observed increase in real house prices.[10] To be able to afford the more expensive housing assets, young households have to borrow more, and so the rising house price also contributes to the rising debt-to-GDP ratio. Housing accordingly provides an alternative vehicle for the transfer of resources over the life cycle, and will raise interest rates in the same way that a bubble can in standard OLG models.

Still, the evolution of house prices does not follow the real interest rate one for one. In our model, the households are only allowed to move (change their housing wealth) at specific ages. This reflects the idea that decisions to buy housing are discrete in real life. Specifically, households buy housing at age 20 and 40, and sell housing at age 70 and 85 (cf. Figure 11).

This has two implications. Firstly, as a savings instrument, housing has a longer maturity than capital, meaning that house prices are more forward looking: households buying housing know that they will have to wait between 15 and 30 years before being able to sell it. The user-cost of housing, therefore, reflects expected future changes in prices (both the interest rate and the housing price). This is shown in Figure 13, which plots the four-period (20-year) ahead average real interest rate

[10] Data on housing wealth come from Piketty and Zucman (2014), Real house prices are from the BIS databases, and are the ratio of nominal house prices to the consumer price index.

Figure 13: User cost of housing and forward looking interest rate

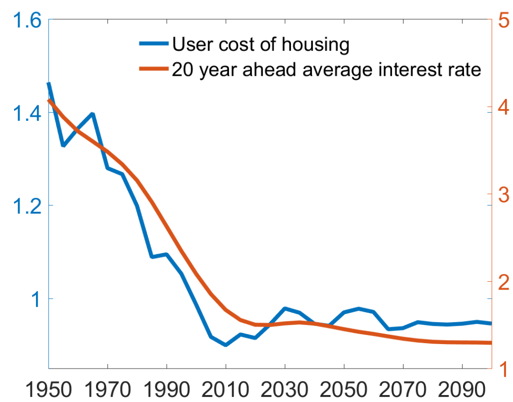

against the per-period user-cost of housing. These variables reflect the transition of house prices more closely than the real interest rate.

The second implication of the discrete move-dates is that the house price is sensitive to changes in the relative size of different cohorts as they move from buying to selling housing. This means that the baby-boom plays an important role in dynamics of house prices. In 2015, the oldest households from the baby-boom generation (born between 1945 and 1970) reach age 70, and the share of the group aged 40 to 69 in the total population starts decreasing, while the share of the group aged 70 to 84 picks up. This implies that the share of households with the highest housing demand, the 40-69 year olds, decreases, while the share of households with a lower housing demand, aged 70 and over, increases. This means that, even as the demographic transition continues, the rise in housing demand slows down around this time and this pushes down the house price. This is shown in Figure 14. The blue line is the aggregate housing demand per capita, which is equal to the aggregate housing supply per capita in general equilibrium, which is constant by assumption. The red line is obtained by keeping the optimal housing choice of each age cohort fixed at its 2015 level, and varying only the weight of those cohorts within the total population. Aggregating over alive cohorts at each date, we obtain the change in housing demand due to the age composition of the population only. Clearly, this drives up the demand for housing from 1980 until 2015, while it stabilises it and drives it down from 2015

Figure 14: Aggregate housing demand, from the general equilibrium and when only population weights are changing

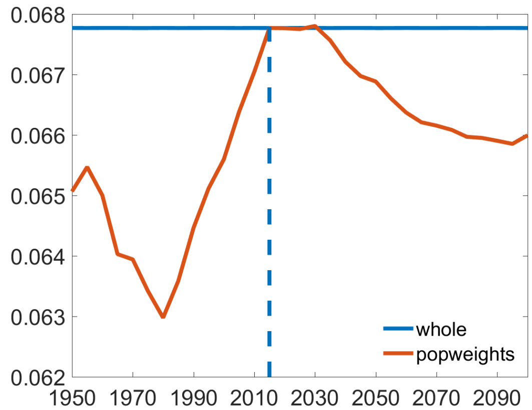

onwards.

4.2.3 Productivity

While our study focuses on the impact of demographics on the natural interest rate, our model can shed some light on other changes observed over the past 40 years. Notably, demographic changes can partially explain the slower productivity growth observed recently. As shown on Figure 5a, the productivity of young and old workers is lower, and productivity reaches its peak around age 50. Hence, a change in the age distribution of the working population implies a different level of the aggregate productivity.

The evolution of the average age of the working population in our model and the resulting productivity growth rate are shown on Figure 15. We can clearly see the impact of the baby-boom generation on the figures. From 1970 onwards, the young baby-boomers start working, bringing down the average age of the working population and the productivity growth. Until 2000, the baby-boomers age and gain in work experience, increasing for the labour force's average age and productivity. From 1990, the baby-boomers generation reaches ages 50 and above, their productivity decreases and hence the productivity growth slows down, while the average age of the

Figure 15: Implications of Demographic Trends for Productivity

(a) Average age of working population (b) Annual productivity growth (%)

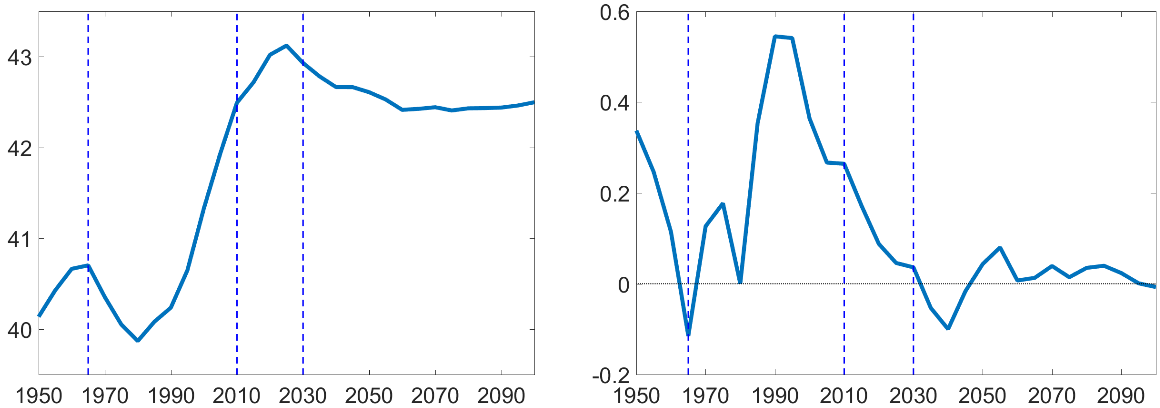

working population keeps increasing. Finally, from 2015 onwards, the baby-boom generations start retiring progressively. The average age of the working population decreases slowly, and productivity growth slows down further. While demographic changes are not the only explanation for the recent slow down in productivity growth, our model shows that the ageing workforce may have played a role in this evolution.

4.2.4 Distributional impact of demographic change

Our model includes heterogeneity across agents in terms of birth year and age, which allows us to measure inequality in two dimensions. The first is within-period: we can compute standard measures of inequality across households alive within each time period, to study how inequality across age groups has evolved over time. The second is inequality across cohorts: are households born earlier better or worse off than those born later? We examine these two aspects below. It is important to emphasise, however, that the model used in this paper is not designed to analyse inequality from either a positive or a normative angle - there is no within-cohort heterogeneity, and there are conceptual difficulties involved when assessing welfare over changing lifespans. We report these results for the sake of completeness, and to give a sense of how the changing age composition of the population and associated

Figure 16: Gini coefficients

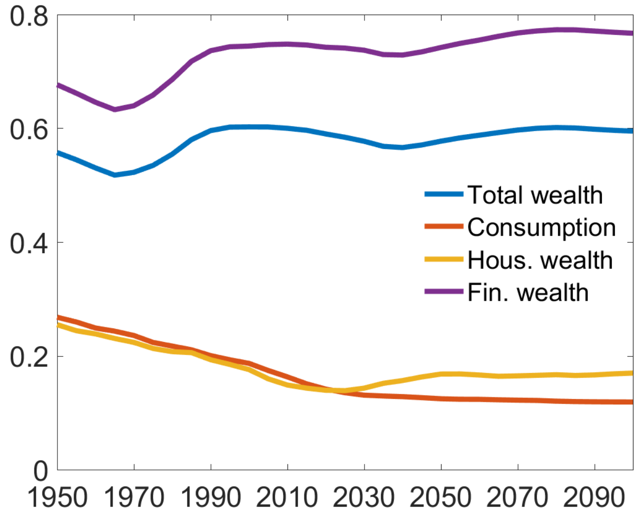

changes in relative prices may affect measured consumption and wealth inequality. Furthermore, we use predicted rather than actual relative prices in order to isolate the impact of demographic change, and ignore, as elsewhere in this paper, the impact of any changes in trend TFP. Our results accordingly should not be read as a full intergenerational welfare analysis.

Within-period inequality A standard measure of inequality is the Gini coefficient, which we compute for consumption, financial and housing wealth. Figure 16 shows the evolution of Gini coefficients implied by our model from 1950 onwards. As above, there are two main drivers of this evolution: the direct effect of the changing age structure and the equilibrium effect of the changing lifecycle profiles of each variable. For consumption, these two components work in the same direction. As agents become older, the age structure is more concentrated, lowering the dispersion of consumption and hence lowering the Gini coefficient. At the same time, the lower interest rate allows households to borrow and consume more when young, while older households tend to consume less as they expect to live for longer. Again this lowers the dispersion of consumption across agents and lowers the Gini coefficient.

For wealth inequality the trend is more muted, with the Gini coefficient for total wealth rising from 1965 until 1990, then slightly decreasing, and rising again after 2030. Financial wealth is the main driver behind the evolution of the Gini coefficient for total wealth. Here, the two drivers work in opposite directions. The initial increase, as well as the increase post-2030, are mostly due to a more unequal wealth distribution across households, with younger agents being more indebted and middle-aged agents holding higher financial wealth. These higher levels of debt and savings are the flip-side of the consumption smoothing described above. The evolution of the wealth distribution at given ages is similar between 1990 and 2030, but the change in the age composition of the population offsets it: the share of poorer, younger households in the total population decreases, while the share of richer, older households increases, thus stabilizing the Gini coefficient for a few decades.

The increase in housing wealth inequality implied by the model after 2020 is due to similar factors. The share of households aged above 70 increases strongly between 2020 and 2045. These households tend to hold less housing, as they sell back part of it to finance their retirement consumption. As a consequence, the distribution of housing in the population becomes more unequal, increasing the associated Gini coefficient.

Inequality across cohorts In our model, cohorts differ across several dimensions: their size and life expectancy according to exogenous demographic trends, but also consumption, wealth and housing levels given by general equilibrium outcomes. One simple way to compare welfare across cohorts is to compute each cohort's expected life-time utility. The red line in Figure 17a shows that the expected life-time utility by birth year is clearly increasing over time, meaning that an agent born in 1980 has a higher expected utility than one born in 1950. In part, this is the simple consequence of increased longevity and hence being able to spread consumption over a longer period of time.

To disentangle the various components at play behind this chart, first we can abstract from the change in longevity by looking at the realised life-time utility of an agent who remains alive until age 90, shown in the blue line in Figure 17a. In this case, the

Figure 17: Life-time utility

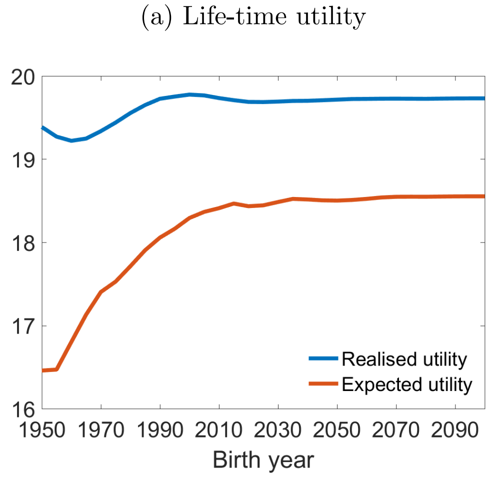

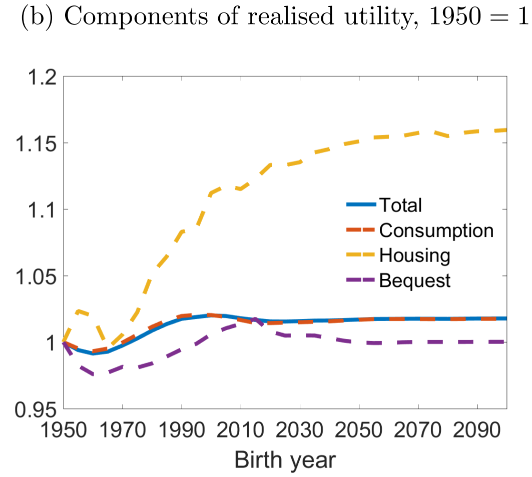

change in realised utility is much smaller and does not monotonically increase any more. Each cohort's realised life-time utility increases from 1960 onwards, with the generations born in the 1980s better off than the baby-boomers in terms of realised consumption and housing. The generations to come (born after 2015) are slightly worse off.

To understand the underlying drivers of this evolution, we examine the respective importance and changes in utility due to housing, bequests and consumption in Figure 17b. The increase in housing utility is the most dramatic (+16% between 1950 and 2050). This is due to a better distribution of housing over the life-cycle, allowed both by the increase ability of young households to borrow and buy housing, and by the changes in the age distribution in the population. The housing share in total utility is however small (less than 2%), so housing cannot be the main driver of total realised utility. Conversely, consumption utility only increases by 2% between 1950 and 2050, but accounts for about 85% of total realised utility and is therefore its main driver. Finally, utility gained from leaving bequests to the next generations represents around 13% of realised utility but changes only slightly over time and tends to balance changes in housing utility.[11]

[11]This corresponds to the additional utility from leaving bequest included in the households' utility function, and not to the utility gains obtained by the households receiving the bequests.

Figure 18: Additional life-time income to afford 1950 consumption bundle at current prices

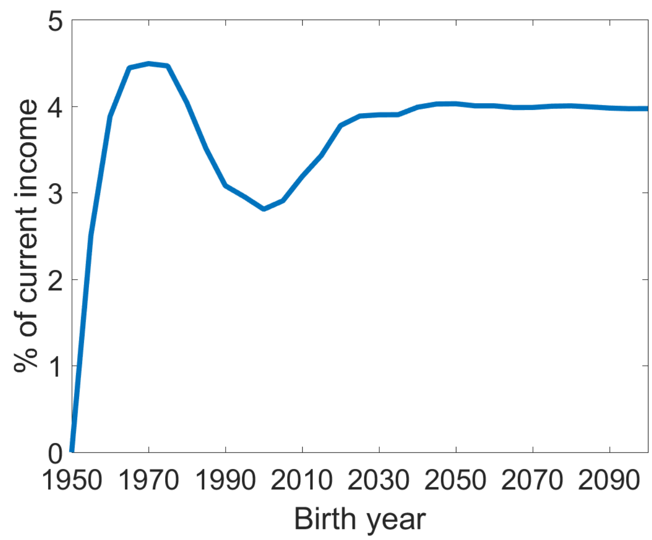

Examining income differences across cohorts is another way to understand the origins of their welfare differences. First, we compute the additional income a cohort born at time t would need to afford the same consumption bundle as the 1950 cohort, at time t prices, as a share of the 1950 cohort's income. Figure 18 shows that this additional income is the largest for cohorts born in the 1960s to 1970s, and then smaller for cohorts born between 1990 and 2010. This is one of the explanations for the bumps in realised utility of these cohorts. While total income is relatively stable across cohorts, Figure 19 shows that this is hiding an increase in labour income and a decrease in capital, housing and bequests income. Because of the decrease in interest rates and the progressive slow down in housing prices, income from housing and interest on capital drops dramatically.

Their aggregate impact remains small however, because housing and interest only account for a small share of total income. Interest income is equal to the net capital share in the economy. In our setup, there are no supernormal profits and hence the net capital share is just the product of the capital-output ratio (a number in the region of 3 to 4) and the annual risk-free real interest rate, a number between 1% and 5% in our simulation. Turning to housing, Figure 12 shows that housing

Figure 19: Life-time income and its components

(a) Absolute values　　(b) Normalised to 1 in 1950

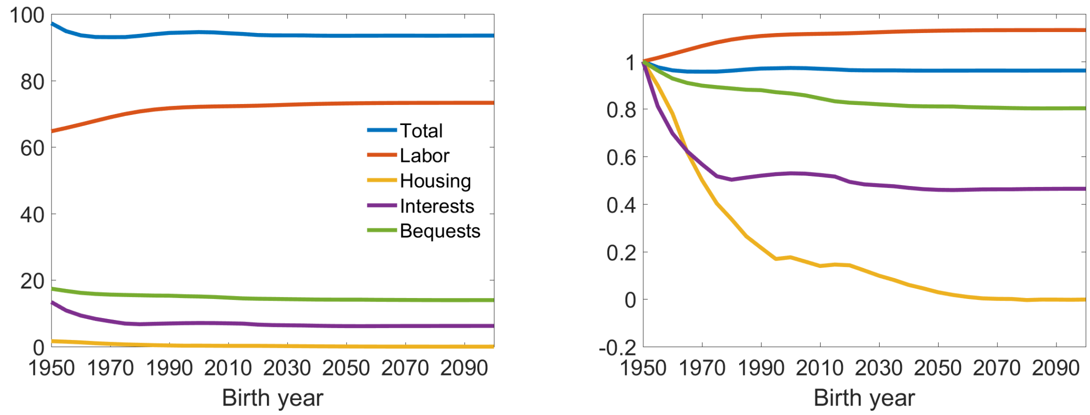

wealth increases by around 1.5% of GDP per year between 1950 and 2030, at which point gains flatten off somewhat. These capital gains accordingly cannot form a very large share of any generation's permanent income, even though they accrue disproportionately to the early baby-boomer generation.

One clear way to summarise this information is to implement a consumption-equivalent variation following Jones and Klenow (2016). We calculate the proportion by which the consumption of a household born in 1950 would need to be adjusted to equalise his welfare to that of a household born at another date t. The intuition is the following. Take a household born at time t, and change his living conditions (in terms of life expectancy, share of GDP used for consumption, inequality, and so on) for the ones of a household born in 1950. To bring this household's welfare back to its initial level, we need to adjust his consumption by a certain amount. This is the consumption-equivalent variation, measured as a proportion of the consumption of a household born in 1950. A consumption-equivalent variation smaller (resp. larger) than 1 means that the household born in 1950 is better (resp. worse) off than the one born at time t. Similarly to Jones and Klenow (2016), we are further able to decompose this measure into various components reflecting: (i) life expectancy; (ii) consumption; (iii) housing; (iv) bequests; (v) consumption smoothing; (vi) housing smoothing and (vii) bequest smoothing. All these components are measured relative to the reference cohort born in 1950. Components (i) to (iv) are clearly positively related to the welfare of the household born at time t, as they directly enter the utility function. Components (v) to (vii), instead, measure the difference at time t between the average utility from, for example, consumption, $E(u(c))$, and the utility obtained from average consumption, $u(E(c))$.[12] An improved smoothing of consumption over time corresponds to a mean-preserving contraction in consumption across age groups. It keeps the utility from average consumption fixed, but increases the average utility from consumption, thus increasing the term (v) so that the relative welfare of the household born at time t improves.

Figure 20 shows the evolution of this consumption-equivalent variation and its components over time. Figure 20a confirms the clear welfare increase over time already observed on Figure 17a. The 1950 cohort's consumption would need to be increased by 11% percent to render a household born in 1970 indifferent between being born in 1970 and in 1950. This goes up to 20% for a household born in 1990 and 25% for 2010, before stabilising towards 28%. The purple solid line in Figure 20a shows that the main driver is life expectancy. Figure 20b further shows that the improvement in consumption smoothing in our model is the second most important driver of welfare, reflecting the same mechanisms that lowered the Gini coefficient for consumption shown in Figure 16, while total consumption and bequest have been oscillating and tend to compensate each others. Here again, the impact of housing and housing smoothing is very limited.

Although the effect of life expectancy appears incredibly important, there is a difficulty in including this in the welfare of different cohorts: while everything else is scale invariant, utility in death is assumed to be fixed at zero, meaning that the relative value of being alive does depend on the calibration scale, specifically if consumption, housing and bequests are above or below one. While the life expectancy component is consistently found to be a major driver of welfare over time, the precise estimation

[12]These components correspond to the ones labeled "inequality" in Jones and Klenow (2016). Given that we have one representative agent of each age in our set-up, there is no inequality within age groups, and these components only measure variations of consumption across age groups over time, and not variations of consumption within age group. We therefore prefer to interpret them as "smoothing" rather than "inequality" components.

Figure 20: Consumption-equivalent variation and its components

(a) Total (λ) and life expectancy component (b) Other components

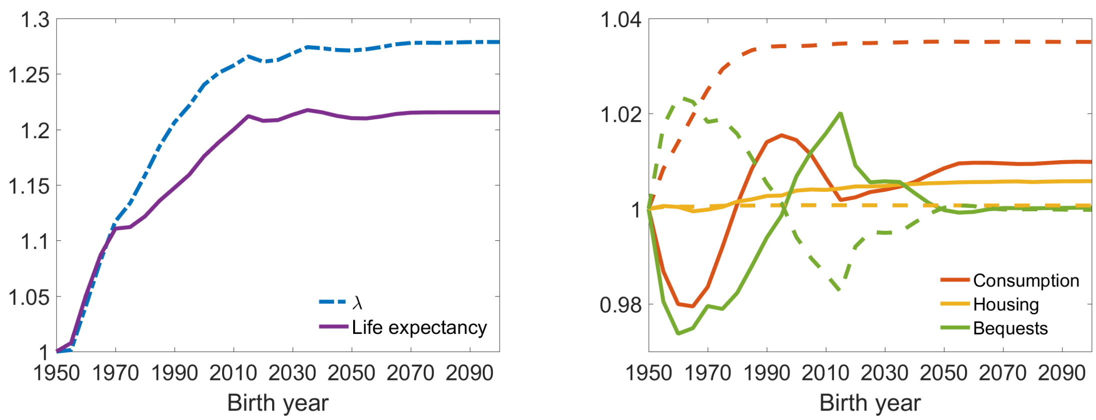

Dashed lines in panel (b) show the 'smoothing' component for each term.

of its impact depends on the normalisation of the model. We therefore prefer to focus on its qualitative implications only. None of the other components presented in Figure 20b suffer from this conceptual problem.

To conclude, demographic changes imply an increase in welfare across cohorts over time, and longevity, though difficult to measure, seems to be its most important driver. Abstracting from longevity, changes in life-time consumption are the main drivers of realised utility, and seem to be detrimental to the baby-boom generation. The effect of capital gains from housing are lower than one might expect. Finally, an improved consumption smoothing across age, thanks to lower borrowing costs, is found to have a non-negligible positive impact on welfare.

4.2.5 Open Economy implications

So far we have considered the group of OECD countries as one closed economy, and looked at the effects of the demographic trends in the aggregate population. While an ageing population is common to all these countries, different countries within this group are ageing at different speeds. Figure 21 shows the OADR for a handful of countries within our aggregate group. As can be seen, Japan and Germany, for example, are ageing much faster than the aggregate, while Australia and the US are ageing more slowly.

Figure 21: OADR Across Countries

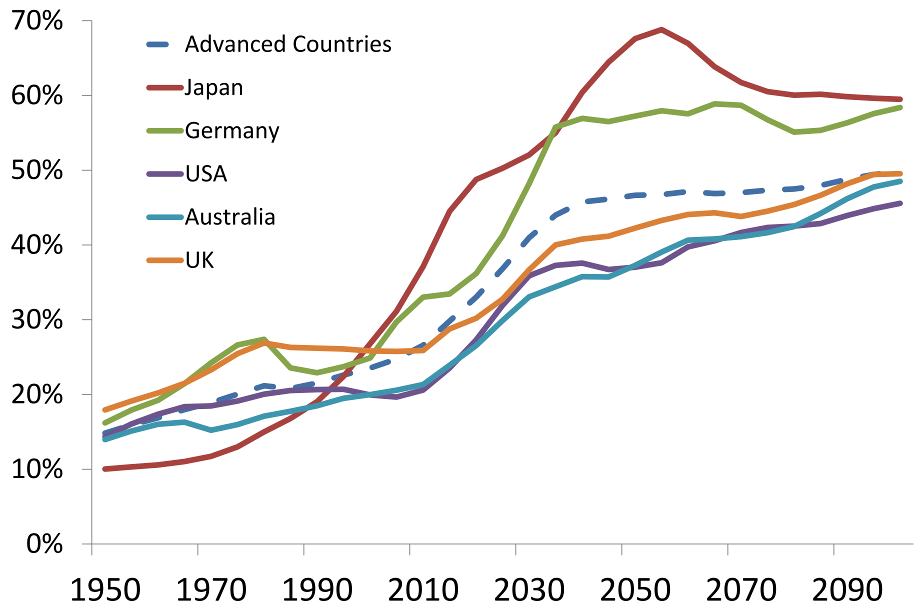

Source: UN Population Statistics (projections based on medium-fertility scenario)

How can our model account for these differences? Consider each of these countries as a small open economy trading on fully integrated global capital markets. In other words, each country takes as given the global real interest rate that arises in the aggregate, as calculated in the previous sections.[13] All else equal, this will mean that the firms in each country will demand the same level of capital relative to output, which can be seen from their first order condition. However, there will be no market-clearing condition for the domestic capital markets, meaning that household savings can be above or below the capital demanded by firms. The discrepancy between domestic savings and domestic capital will give rise to a non-zero net foreign asset

[13]Notice that, so long as we keep the parameters of the model the same for each country, taking the path of the interest rate from the closed economy exercise is equivalent to solving the model as a multi-country world economy with perfectly integrated capital markets, abstracting from country-specific differences in real interest rates.

(NFA) position for the domestic economy. In particular, if domestic savings are higher than domestic capital, this means that domestic households must place their savings into capital abroad. Conversely, if domestic capital is higher than domestic savings, this means that some of the domestic capital is owned by foreign households.

Consider a country such as Australia, which is ageing more slowly than the average. This means that demographic trends are putting less upward pressure on savings in Australia, and hence the global real interest rate is below the interest rate that would arise if Australia was a closed economy. In other words, the savings of domestic households in Australia is below the desired capital level of Australian firms. This translates to a negative NFA position for Australia, as capital flows into Australia from foreign households. Conversely, for a country such as Germany, which is ageing faster than the average, the global interest rate is above the rate that would equilibriate the domestic capital market, and this translates to capital outflows from Germany and the accumulation of foreign assets by German households.

To quantify this, we can solve for equilibrium in a small open economy version of the OLG model, where the interest rate is exogenous and instead of the capital market clearing condition, we have an equation that defines net foreign assets

$$\widetilde{NFA}_t = \tilde{A}_{t-1} - \tilde{K}_t$$

We solve this version of the model dynamically with the exogenous path of the real interest rate set as the path of the real interest rate from the aggregate exercise, as shown in Figure 8, and feeding in the demographic variables of a given country. Figure 22 shows the resulting path of the net foreign assets for the US, UK, Australia and Germany. The simulations assume that the economy is always at the dynamic equilibrium, omitting, for example, the major fiscal and physical consequences of the Second World War. The model also omits any frictions in the international movement of capital, such as capital controls or home-bias in portfolio allocations, which were an important feature of the world economy at least in the early post-war period. Nonetheless even this simple exercise can capture the dynamics of NFAs, with Australia, the UK and the US having increasingly negative NFA positions both

Figure 22: NFA/GDP (%) in the Small Open Economy Simulations

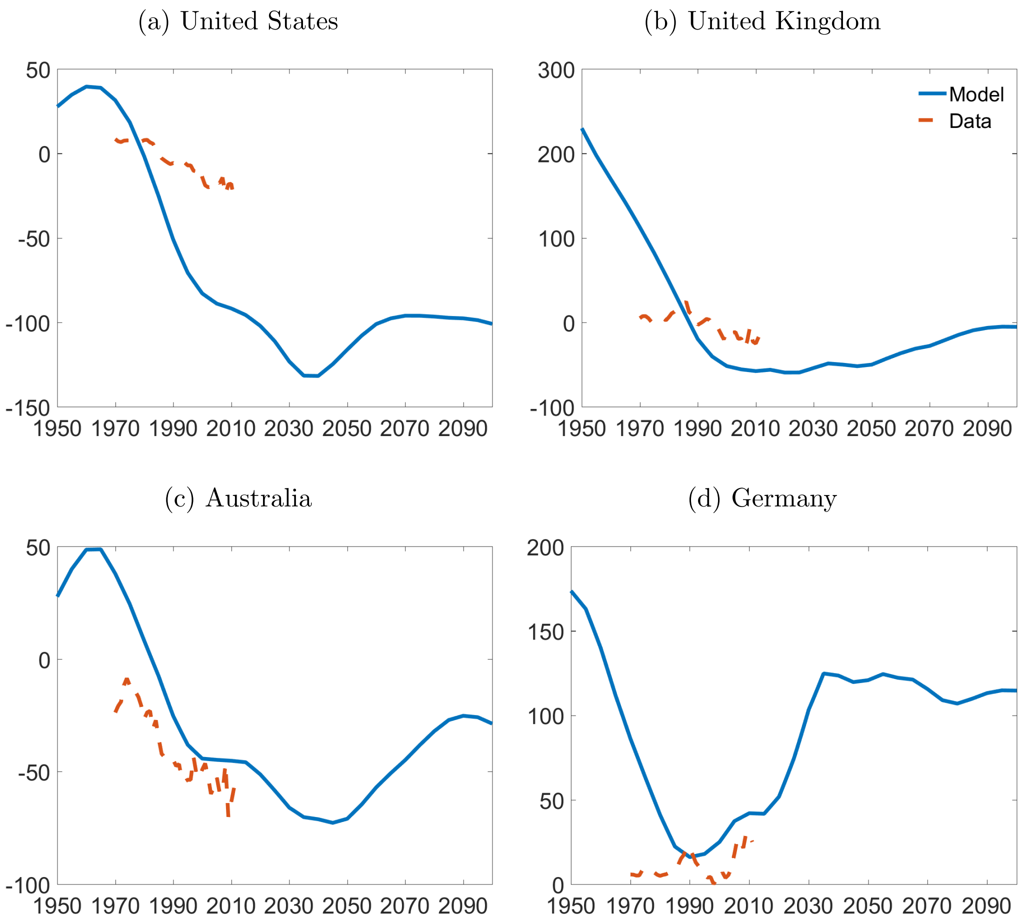

in the model and in the data, and Germany building up an increasingly positive NFA position. The model also suggests that the NFA position in the US and Germany will diverge further in the coming decades, as their demographic characteristics diverge from the aggregate of the OECD, while for the UK and Australia it will remain stable.

To get a broader idea of the cross-country fit of this exercise, Figure 23 plots the level of the NFA-to-GDP ratio in 2010 against the predicted level from the model.[14] This

[14]The NFA data is taken from the updated and extended version of dataset constructed by Lane and Milesi-Ferretti (2007).

Figure 23: NFA/GDP in the Model vs Data

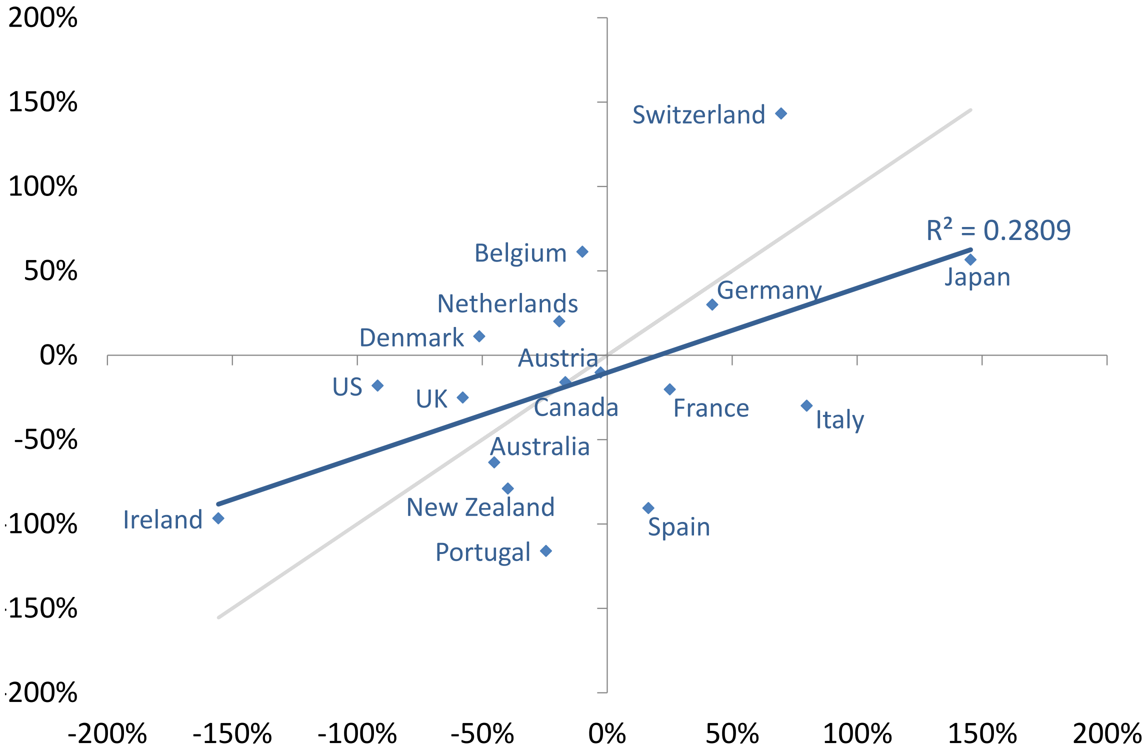

Note: Model on x-axis and Data on y-axis, grey line is the 45 degree line.

exercise can be interpreted as a test of the mechanisms of our model against the data. Again, we see that the model predicts slightly larger NFA positions than we observe in the data, with the trend line in this scatter plot being somewhat shallower than the 45 degree line. Nonetheless, a substantial part of the cross-country differences in NFAs can be explained by the model looking only at differences in demographics. This gives us greater confidence about the mechanisms underlining all of the results from our model.

Finally, Figure 24 plots the NFA position in 2010 against the High-wealth ratio in 2010, for the model outcome and the data, across all of the 17 countries in our aggregate advanced economies group. We see again that the model tends to predict a larger NFA position than observed in the data. Nonetheless, it does well to explain the cross-country pattern of NFA positions.

Figure 24: Demographic Changes and NFA accumulation

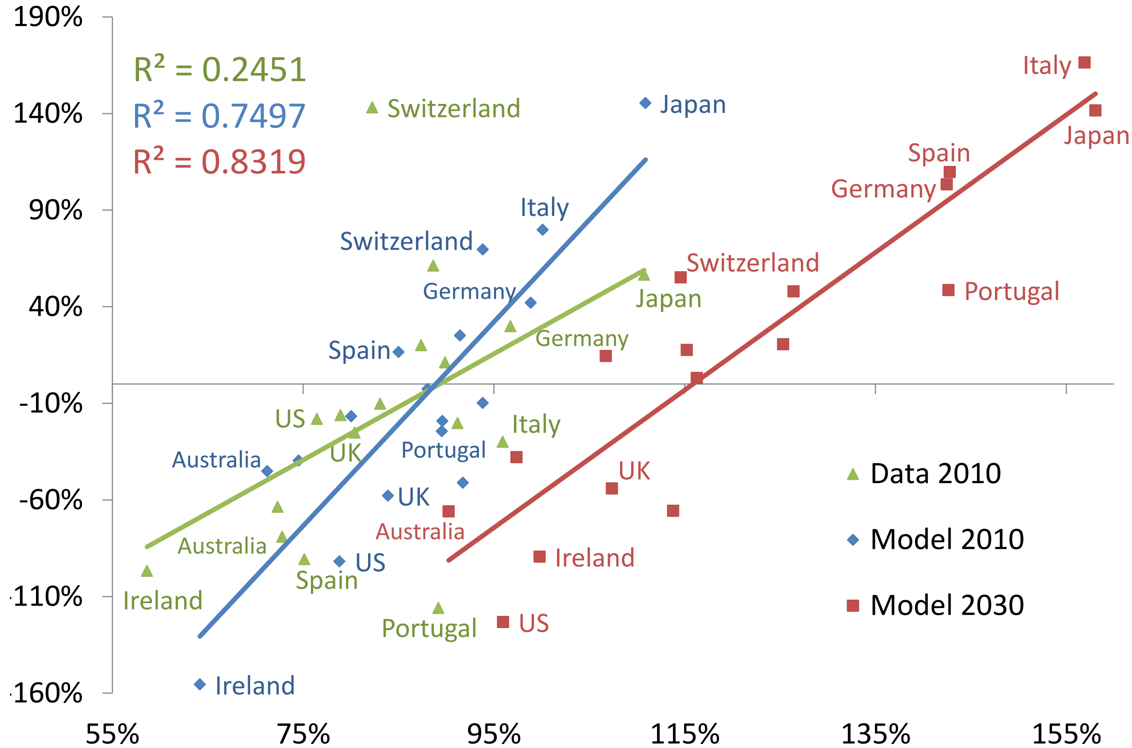

Note: HWR on x-axis and NFA/GDP on y-axis.

Figure 24 also includes the model predictions for NFA positions against the HWR in 2030. All countries move to the right on the HWR scale as they age. As this happens, the model predicts that some countries will move towards higher NFA positions, as they age faster than the average, while other countries will have increasingly negative NFA positions as they age more slowly than the average.

4.3 Robustness

4.3.1 The Role of Housing

The overall impact of housing on the model results is quantified on Figure 25, which compares the baseline results against the results from a model in which we exclude housing. To facilitate interpretation, we keep the parameter values obtained in the baseline case to solve the model without housing. Consequently, aggregate savings and the interest rate are higher (resp. lower) over the whole transition period, and aggregate variables without housing do not match the target set in the baseline case.

As expected, the level of the capital to output ratio increases more in the absence of housing, as households do not have any alternative for transferring wealth over time. Households also accumulate less debt, as they do not need to borrow to afford housing. Given the curvature of the production function, the impact of housing on the interest rate drop is smaller than on the level of capital to GDP. In terms of the marginal effect of including housing in the model, the fall in the interest rate between 1980 to 2100 is around 230bps in the model without housing, 10bps larger than the baseline. Conversely, the rise in the household debt-to-GDP ratio over the same years is 15pp lower in the model without housing.

4.3.2 Changing the Retirement Age

The retirement age in our model is fixed at age 65 during the whole transition. Increasing the retirement age may seem more realistic, as reforms in that direction have been implemented in most advanced economies. This is likely to offset the effects of population ageing, as households spend less time in retirement for a given life expectancy.

Solving the model with a change in the retirement age during the transition requires some additional assumptions, particularly in terms of the timing of the announcement and implementation of these changes. As a first step, however, solving the model with a higher retirement age throughout the simulation can give us an insight into the importance of the retirement age. Figure 26 shows the transition path of the main variables in our model, in the baseline case, and when the retirement age is fixed at 70 all the time. We still need to make an additional assumption about the productivity level of older workers: here we assume that it is the same as the 60-64 year old cohort, which we consider as an upper bound on their potential productivity.

When retiring at age 70, the households save less, and the interest rate is higher.

Figure 25: Simulations With and Without Housing

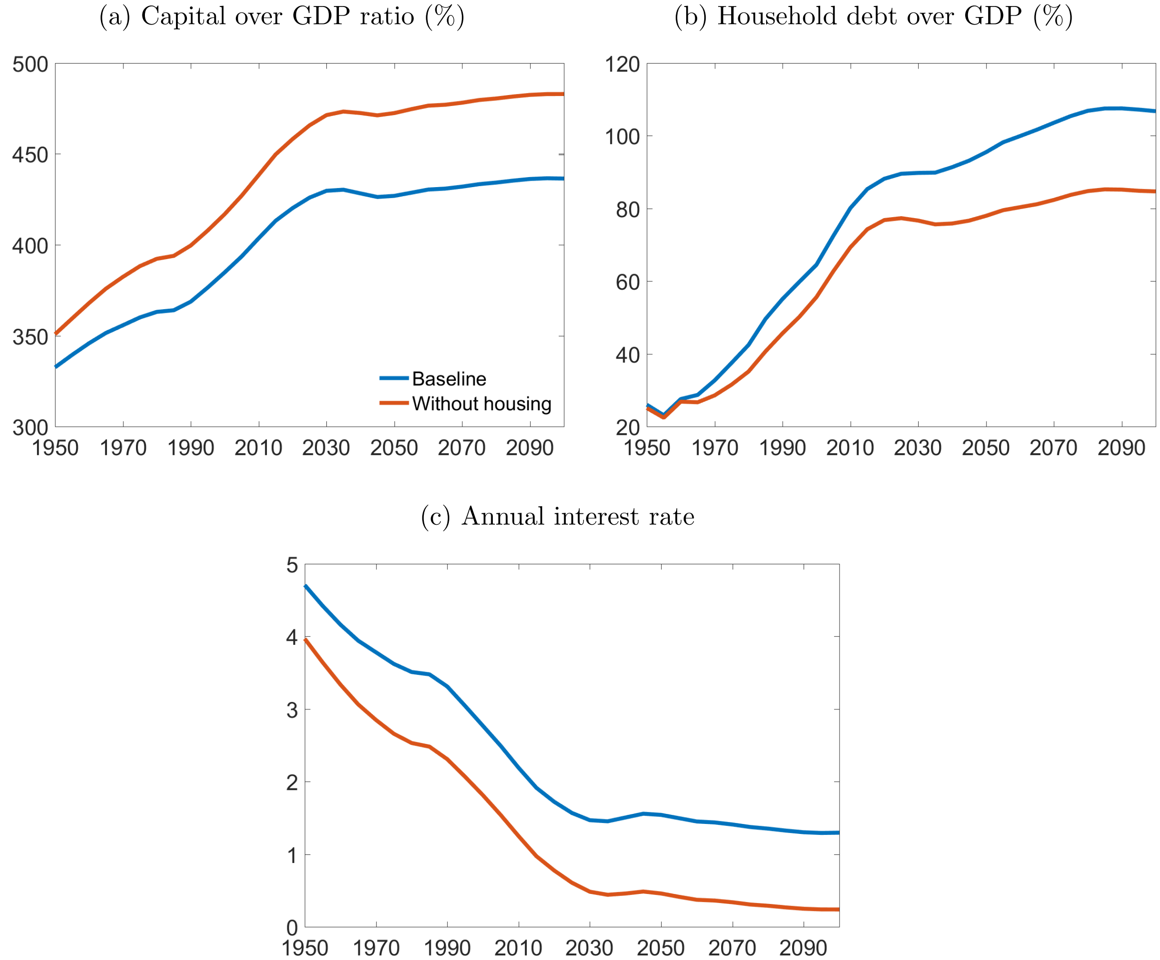

The interest rate drop between 1990 and 2025 is very similar in both cases (170bp with late retirement, against 174bp in the baseline). If the retirement age were to change unexpectedly during the transition, say in 2000, the models outcome would be identical to the baseline case until 2000. After that date, the households would start progressively adjusting their saving and housing decisions to the new retirement age, to reach a final steady state identical to the one obtained with retirement at age 70. Hence the transition would lie somewhere between these two lines.

The potential effect of raising the retirement age even as high as 70 years old is

Figure 26: Model simulations with higher retirement age

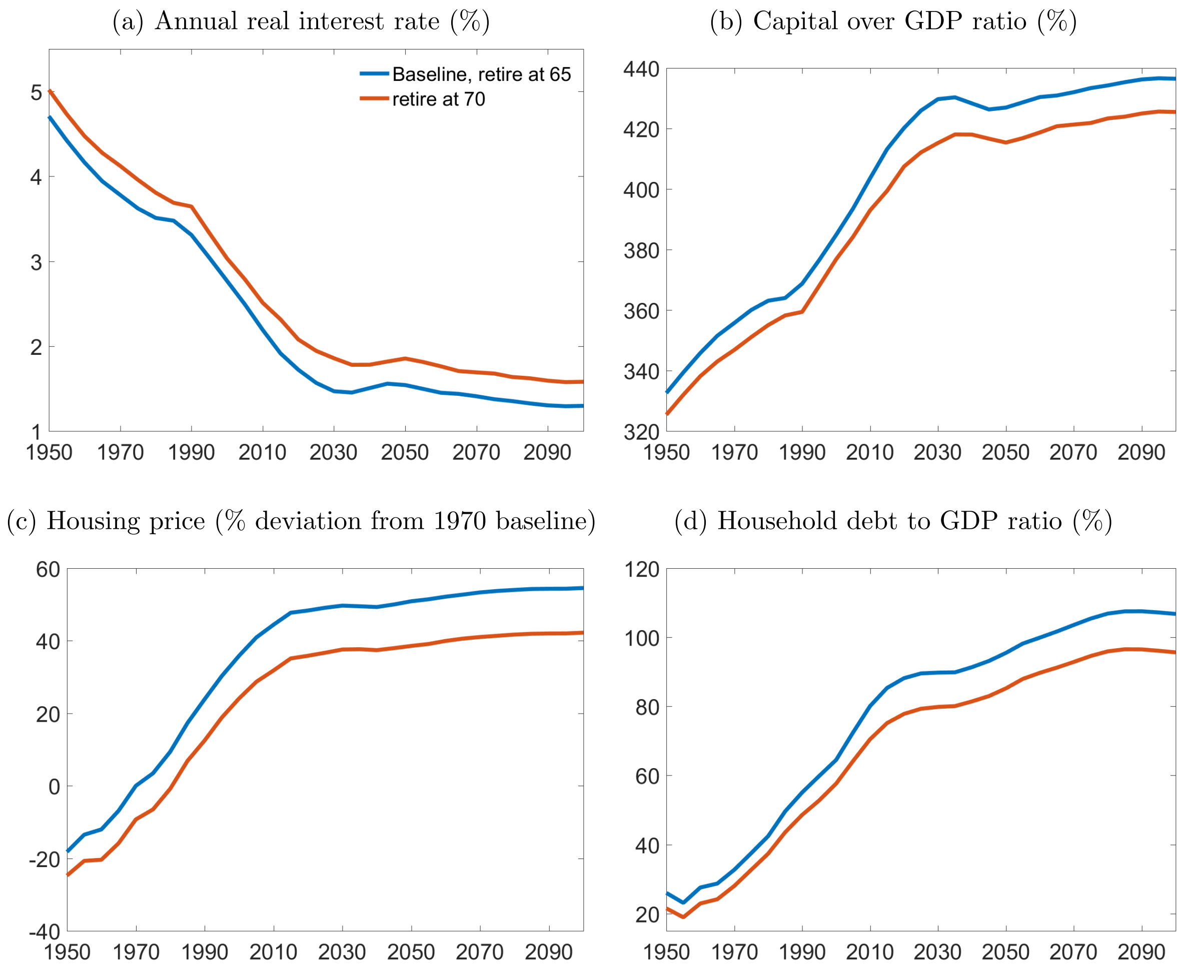

very modest in this model. Given the life cycle profile of productivity, an additional five years of labour income later in life does not offset the incentive to save in the highest productivity stage of life in order to smooth consumption. Furthermore, as life expectancy goes towards 90, five additional years of work have little effect on the overall proportion of life spent in retirement. Note that the retirement age in the UK is currently set to increase gradually from 65 to 68 by 2046, a smaller change than the one we have assumed.

Figure 27: Demographic change in the United States and the OECD

(a) Annual cohort growth

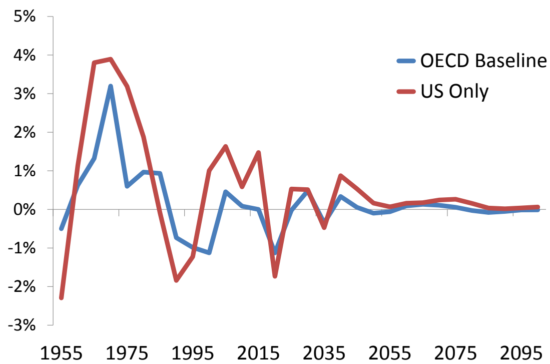

(b) OADR

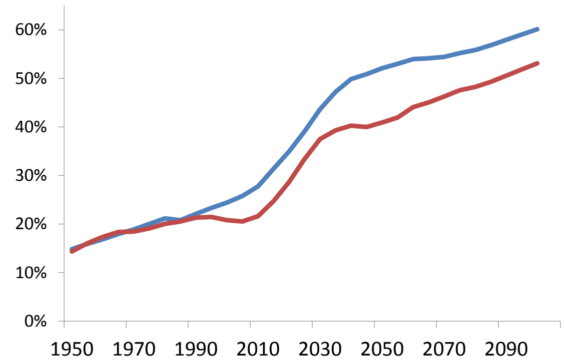

Source: UN Population Statistics

5 Extensions

5.1 Comparing with the United States

While our main results consider the aggregate evolution of OECD countries, looking at the case of the United States more specifically brings useful insights.[15] This is true not least because much of the current literature on low interest rates, and the role of demographics, has focused on the US as a closed economy. Population ageing in the United States is somewhat slower than the OECD average: population growth is more dynamic and life-expectancy at age 60 remains below that of the OECD.[16] Consequently, the old age dependency ratio doubles between 1950 and 2015 in the OECD, while it rises by only two thirds in the United States (see Figure 27).

The impact of demographic change on the interest rate is therefore smaller in the United States: 134 basis points between 1980 and 2015. As the baby-boom is stronger

[15] Here we look at the situation of the United States as a closed economy, so that the domestic savings have to equate domestic capital to reach the equilibrium on the capital market. We will turn to an open economy exercise in Section 5.

[16] The life expectancy at age 60 of the cohort born in 1980 is 84.5 years in the OECD, against 83.7 in the United States

Figure 28: Simulations for the US as a Closed Economy

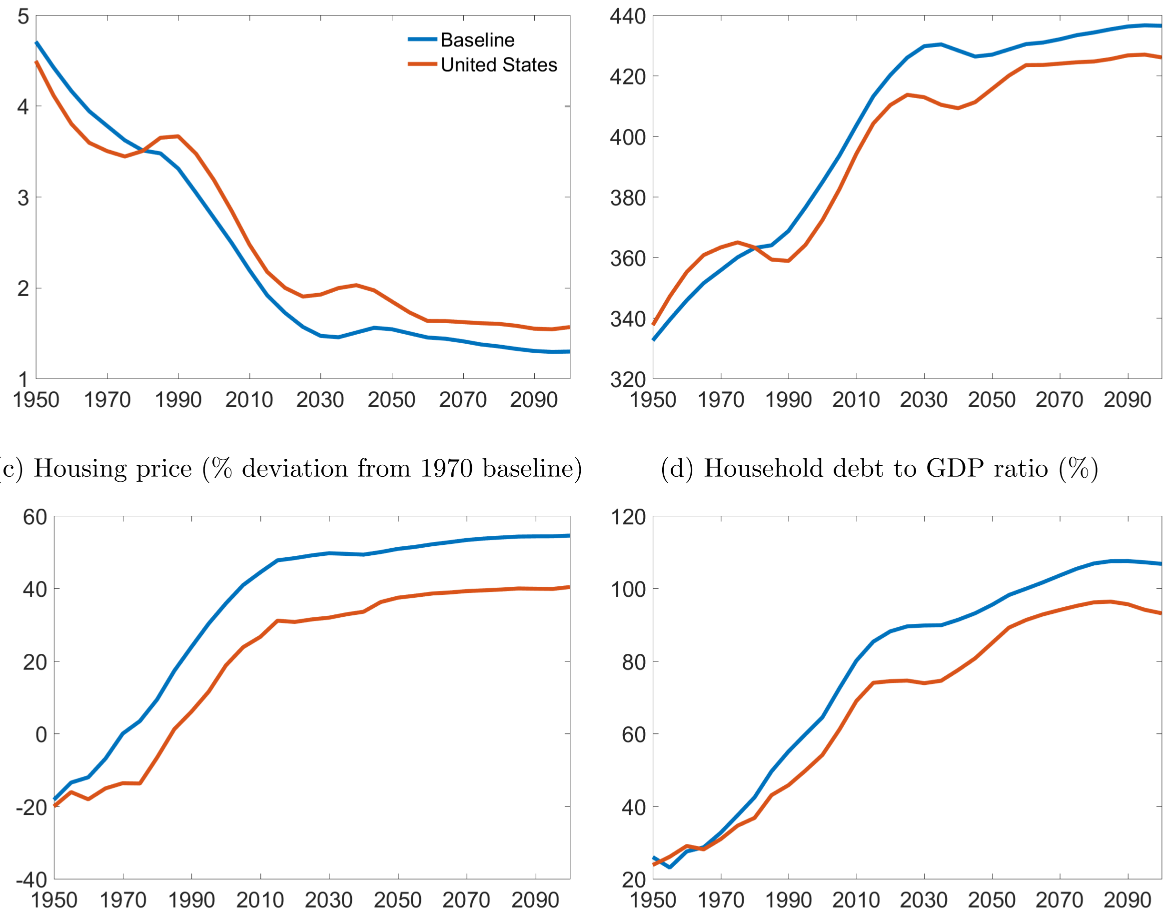

(a) Annual real interest rate (%)

(b) Capital over GDP ratio (%)

(c) Housing price (% deviation from 1970 baseline)

(d) Household debt to GDP ratio (%)

in the United States, the resulting transition path of the interest rate is also less smooth. Similarly, the capital to GDP, household debt to GDP and housing price increase slower than in the OECD case (see Figure 28). In the data, the US real interest rate starts from a higher point and decreases more between 1980 and 2015 relative to the World interest rate, meaning that demographic changes explain a smaller part of the fall in the US interest rate. Over the same period, the increase in housing prices is however slower in the United States than in the OECD, corresponding to the implications of the model (see Figure 29a). In terms of household debt to GDP ratio, the data for the US are more strongly influenced by the boom and bust of the 2000s, but it seems that the trend increase is equivalent in the US to the whole OECD (Figure 29b).

Figure 29: House prices and household debt in the US and the OECD

(a) Real housing price (% deviation from 1976) (b) Household debt to GDP ratio (%)

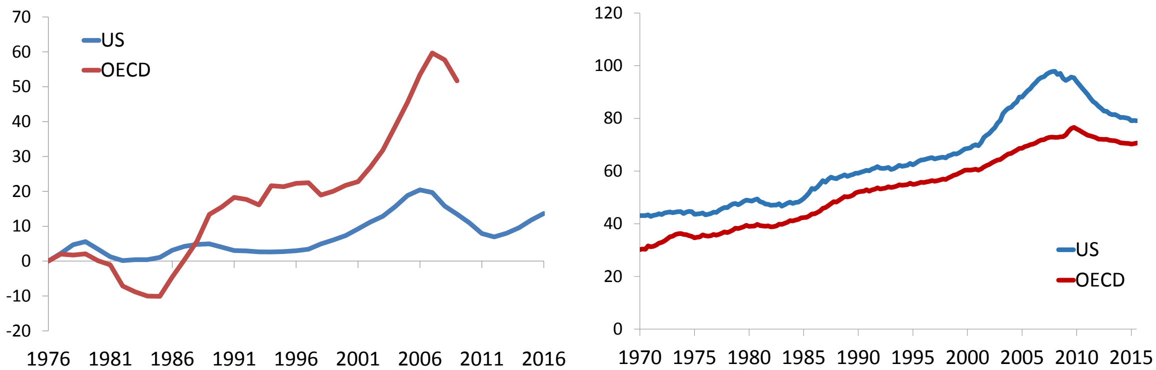

5.2 Monopoly Power

As a final extension, we consider what happens in our model if we remove the assumption of perfect competition on the firms side. We model monopoly power from the firm that leads to prices as a mark-up over marginal costs, which we denote by μ. This creates a wedge between wages and the rental price of capital and the marginal product of labour and capital, respectively

$$w_t = \frac{1}{\mu}\frac{\partial Y_t}{\partial L_t} \quad (6)$$

$$r_t = \frac{1}{\mu}\frac{\partial Y_t}{\partial K_t} \quad (7)$$

These equations imply that a rise in the degree of monopoly power, which raises the price mark-up, pushes down on the demand for capital at a given interest rate. In other words, the investment schedule is shifted inwards. However, to see the general equilibrium effects of introducing monopoly power into the model, we also need to consider how the savings schedule responds. Savings will also respond to rising monopoly profits because these profits must be paid to someone, who will then spend or save them.

To see these effects in our model, we simulate this alternative model under different assumptions about how these super-normal profits are distributed among the households. For the purpose of this illustration, we set the net mark-up to 15%, and keep all other parameters in the model unchanged from the baseline calibration.[17] We first consider the case where the super-normal profits are exogenously distributed as part of the non-labour income, in other words it is taken as given by the household. Within this exogenous case we consider both a case in which all of the profits are given to the young, specifically to households in the first 4 periods of life, and a case in which all the profits are given to the old during retirement. As an alternative to this exogenous case we also consider a case where there is a claim to the future stream of profits from the monopolistic firm, which can be traded among households and has an equilibrium price.[18]

Figure 30: Simulations With and Without Monopoly Power

(a) Capital over GDP ratio (%) (b) Annual real interest rate (%)

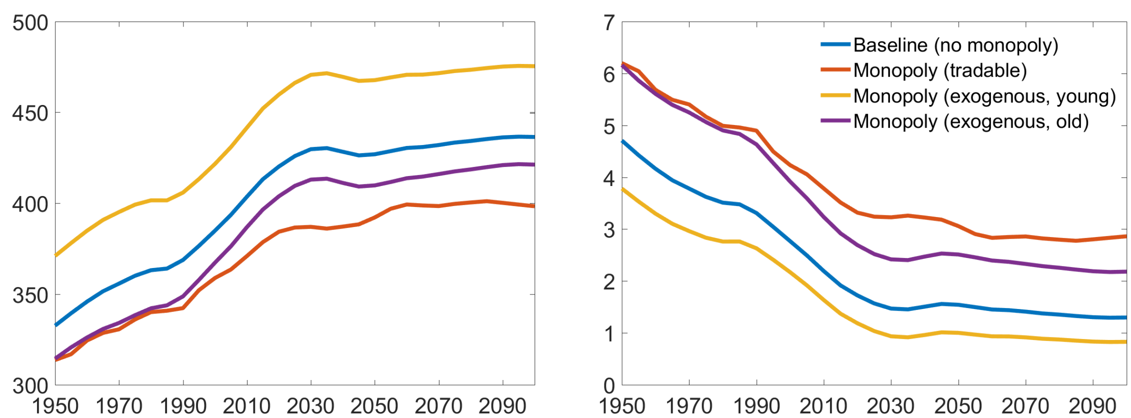

Figure 30 shows the capital-output ratio, interest rate and debt-to-GDP ratio for

[17]The mark-up of 15% corresponds to the calibration for the 1970s in Eggertsson et al. (2017).

[18]This equilibrium price will be such that the return on this claim is equal to the real interest rate. With this condition, the household's problem remains unchanged except that their choice variable will be some $\tilde{a}$ which is a composite of the two assets.

these cases, against the baseline case for comparison. The introduction of monopolistic competition, and, importantly, the assumption about the distribution of the resulting profits both have a non-negligible effect on the results, not just in terms of the magnitude but even in the direction of the effect.

In particular, introducing monopolistic profits that are given exogenously to the young raises the capital-output ratio and lowers the equilibrium interest rate. In this case, the profits are being paid to households at a time when they want to save. Hence introducing these profits increases the resources that these households have, and so pushes out the saving schedule. Assuming that the young own the rights to the monopolistic firms in each period can represent a situation, for example, where the young invent new products which then become obsolete through the creative destruction wrought by the next generation of entrepreneurs. On the other hand, if the profits are given to the households during retirement, this acts to smooth the household's income over the life-cycle, hence reducing the need to save in order to smooth consumption. This lowers the capital-output ratio and raises the real interest rate.

The case in which the claim to the future stream of monopolistic profits is tradable, where we are effectively modelling equity markets, behaves in a similar way. In this case, the claim on monopoly profits is an alternative savings vehicle to transfer income across time. This reduces the supply of capital, pushing in the savings schedule. As well as raising the interest rate at each point in time, this case implies that the demographic changes have a smaller effect on the risk-free interest rate, due to the existence of this alternative savings vehicle.[19]

6 Conclusions

In this paper we use an overlapping generations model, calibrated to advanced-country data, to assess the contribution of population ageing to the fall in real

[19]This result for the tradable case is similar to the findings of Auclert and Rognlie (2016), who show that higher monopoly profits can alleviate a liquidity trap by providing an alternative vehicle for savings.

interest rates which the world has seen over the past three decades. We find that global demographic change can explain three-quarters of the 210bp fall in global real interest rates since 1980, and larger fractions of the rises in house prices and debt. Importantly, the sign of these effects will not reverse as the baby-boomer generation retires: demographic change is forecast to reduce rates by a further 37 bp by 2050. Our model can also explain about 30% of the pattern of industrialised-country NFA positions.

Among the many uncertainties contained in our analysis, we conclude by highlighting the two most important. The first relates to individual behaviour, and in particular the prediction in our model that households will respond to higher life expectancy with increased saving. How much of these demographic changes are actually anticipated by households in reality? There is limited evidence in the literature showing that savings rise as life expectancy rises. De Nardi et al. (2009) use variations in life expectancy by gender, initial health and permanent income to show that higher life expectancy does lead to higher savings, but their focus is on the savings behaviour of retirees rather than workers. Similarly, both Bloom et al. (2003) and Kinugasa and Mason (2007) use cross-country panel regressions to show that higher average life expectancy can explain higher national savings rates, but they do not address the potential reverse causation from higher wealth to higher life expectancy due to availability of health care and sanitation.

The second uncertainty around our results relates to the global economy, and in particular to the pace and ultimate extent to which emerging markets and low-income countries integrate into world capital markets. These populations have different demographic profiles than advanced economies: they are generally much younger, although emerging markets are set to age rapidly in the coming decades (see Figure 31). Their integration into world capital markets, either directly or indirectly through migration into advanced economies, could potentially mitigate the downward pressure on real interest rates from demographic change. On the other hand, if households or institutions in these economies have a higher propensity to save than advanced economies, they could put further downward pressure on real interest rates.

Figure 31: Old-Age Dependency Ratio Around the World

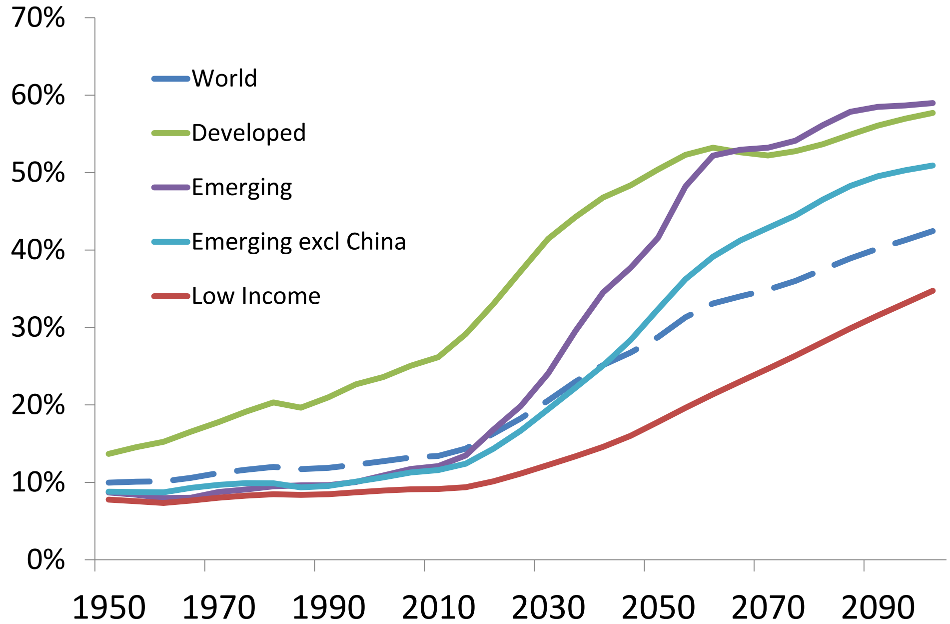

Source: UN Population Statistics (projections based on medium-fertility scenario)

References

Auclert, A. and M. Rognlie (2016). Inequality and aggregate demand. *Mimeo*.

Backus, D., T. Cooley, and E. Henriksen (2014). Demography and low-frequency capital flows. *Journal of International Economics 92, Supplement 1*, S94 – S102. 36th Annual {NBER} International Seminar on Macroeconomics.

Barkai, S. (2016). Declining labor and capital shares. *Stigler Center for the Study of the Economy and the State New Working Paper Series* (2).

Bloom, D. E., D. Canning, and B. Graham (2003). Longevity and life-cycle savings. *The Scandinavian Journal of Economics 105*(3), 319–338.

Caballero, R. J., E. Farhi, and P.-O. Gourinchas (2017, May). Rents, technical change, and risk premia accounting for secular trends in interest rates, returns on capital, earning yields, and factor shares. *American Economic Review 107*(5), 614–20.

Carvalho, C., A. Ferrero, and F. Nechio (2016). Demographics and real interest rates: Inspecting the mechanism. *European Economic Review 88*(C), 208–226.

De Nardi, M., E. French, and J. B. Jones (2009, May). Life expectancy and old age savings. *American Economic Review 99*(2), 110–15.

Domeij, D. and M. Flodén (2006, 08). Population Aging And International Capital Flows. *International Economic Review 47*(3), 1013–1032.

Eggertsson, G. B., N. R. Mehrotra, and J. A. Robbins (2017, January). A model of secular stagnation: Theory and quantitative evaluation. Working Paper 23093, National Bureau of Economic Research.

Fernald, J. G. (2014). A quarterly, utilization-adjusted series on total factor productivity. Technical report.

Gagnon, E., B. K. Johannsen, and J. D. Lopez-Salido (2016, September). Understanding the New Normal : The Role of Demographics. Finance and Economics Discussion Series 2016-080, Board of Governors of the Federal Reserve System (U.S.).

Gertler, M. (1999, June). Government debt and social security in a life-cycle economy. *Carnegie-Rochester Conference Series on Public Policy 50*(1), 61–110.

Glover, A., J. Heathcote, D. Krueger, and J.-V. Rios-Rull (2014, May). Intergenerational Redistribution in the Great Recession. Staff Report 498, Federal Reserve Bank of Minneapolis.

Goodhart, C. and M. Pradhan (2017). Demographics will reverse three multi-decade global trends. *BIS Working Papers* (656).

Holston, K., T. Laubach, and J. C. Williams (2017). Measuring the natural rate of interest: International trends and determinants. *Journal of International Economics 108*(S1), 59–75.

Jones, C. I. and P. J. Klenow (2016, September). Beyond GDP? Welfare across Countries and Time. *American Economic Review 106*(9), 2426–2457.

King, M. and D. Low (2014, February). Measuring the "world" real interest rate. Working Paper 19887, National Bureau of Economic Research.

Kinugasa, T. and A. Mason (2007). Why countries become wealthy: The effects of adult longevity on saving. *World Development 35*(1), 1–23.

Kiyotaki, N., A. Michaelides, and K. Nikolov (2011). Winners and losers in housing markets. *Journal of Money, Credit and Banking 43*(2-3), 255–296.

Knoll, K., M. Schularick, and T. Steger (2014). No price like home: Global house prices, 1870-2012. CEPR Discussion Papers 10166, C.E.P.R. Discussion Papers.

Krueger, D. and A. Ludwig (2007, January). On the consequences of demographic change for rates of returns to capital, and the distribution of wealth and welfare. *Journal of Monetary Economics 54*(1), 49–87.

Lane, P. R. and G. M. Milesi-Ferretti (2007). The external wealth of nations mark ii: Revised and extended estimates of foreign assets and liabilities, 19702004. *Journal of International Economics 73*(2), 223 – 250.

Marx, M., B. Mojon, and F. Velde (2016). Why are real interest rates so low? 2016 Meeting Papers 1581, Society for Economic Dynamics.

Philippon, T. and C. Jones (2016). The secular stagnation of investment? *Mimeo*.

Piketty, T. and G. Zucman (2014). Capital is back: Wealth-income ratios in rich countries, 1700-2010. *The Quarterly Journal of Economics 129*(3), 1255–1310.

Rachel, L. and T. D. Smith (2015, December). Secular drivers of the global real interest rate. Working Paper 571, Bank of England.

Waldron, M. and F. Zampolli (2010, March). Household debt, house prices and consumption in the United Kingdom: a quantitative theoretical analysis. Bank of England working papers 379, Bank of England.