Population Aging and the Macroeconomy*

Noëmie Lisack,[a] Rana Sajedi,[b] and Gregory Thwaites[c]
[a]Banque de France
[b]Bank of England
[c]University of Nottingham

We quantify the impact of demographic change on real interest rates, house prices, and household debt in an overlapping-generations model. Falling birth and death rates across advanced economies can explain much of the observed fall in real interest rates and the rise in house prices and household debt. Since households maintain relatively high wealth levels throughout retirement, these trends will persist as population aging continues. Countries aging relatively slowly, such as the United States, will increasingly accumulate net foreign liabilities. The availability of housing as an alternative store of value attenuates these trends, while raising the retirement age has limited effects.

JEL Codes: E21, E43, E13, J11.

1. Introduction

The population of advanced countries has aged rapidly over the past half-century, with life expectancy and the old-age dependency

*We are grateful to Pavel Brendler, Daisuke Ikeda, anonymous referees, and seminar participants at the Banka Slovenije, Danmarks Nationalbank, the Riksbank, the International Monetary Fund, the Reserve Bank of Australia, the Center for Central Banking Studies, the Bank of England, the EUI Alumni Conference, the Federal Reserve Bank of San Francisco, the Money Macro and Finance Conference, and the DEGIT XXII Conference for their feedback. Remaining errors are ours. This paper reflects the opinions of the authors and not necessarily the views of the Bank of England or its committees, or the Banque de France or the Eurosystem. This paper is an updated version of Lisack, Sajedi, and Thwaites (2017). Author contact: Lisack: Banque de France, 31 rue Croix des Petits-Champs, 75001 Paris, France; e-mail: noemie.lisack@banque-france.fr. Sajedi: Bank of England, Threadneedle Street, London, EC2R 8AH, United Kingdom; e-mail: rana.sajedi@bankofengland.co.uk. Thwaites: University of Nottingham, School of Economics, Sir Clive Granger Building, University Park, Nottingham, NG7 2RD, United Kingdom; e-mail: gregory.thwaites@nottingham.ac.uk.

ratio having already reached unprecedented highs and projected to remain high for several decades. At the same time, long-term real interest rates have been on a downward trend, while house prices and household debt have risen dramatically. This paper quantifies the link between these important trends and examines the wider macroeconomic implications of demographic change.

We build a calibrated neoclassical overlapping-generations (OLG) model to quantify the impact of these factors for advanced economies. We find that the aging of the population in advanced economies can explain around 75 percent of the fall of roughly 210 basis points (bps) in the Holston, Laubach, and Williams (2017) measure of natural real interest rates between 1980 and 2015. Demographic pressures are forecast to reduce real interest rates a further 46 bps by 2050. Furthermore, past falls in interest rates, along with the life-cycle pattern of housing demand, mean that demographics can explain more than four-fifths of the 50 percent rise in real house prices between 1970 and the start of financial crisis. 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 over that same period.

Our main findings highlight that the concerns about future rises in real interest rates as baby-boomers retire 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 the population in high-wealth stages of life, starting from around age 50 according to the data, will therefore tend to raise the capital-output ratio even in the absence of any behavioral reaction to higher life expectancy. Therefore, we find that while the retired may dissave, they still hold more wealth, and this higher stock pushes down on the interest rate. This is in contrast to the findings in Carvalho, Ferrero, and Nechio (2016) and the arguments in Goodhart and Pradhan (2017). Second, population aging is predominantly linked to a trend rise in life expectancy, rather than the transient rise in birth rates in the baby boom. 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 aging of the baby-boom generations 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 tends to raise the capital-labor 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. This pushes down on the marginal product of capital, which is a key determinant of the real interest rate.[1] 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' savings function slopes upward or downward.

We use our model to study the role that housing plays in mediating these effects. 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. 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.

Using the heterogeneity by age and birth year within our model, we examine the implications of demographic changes for the level and distribution of welfare across time. Looking at expected welfare across cohorts over time, the main driver is increased longevity, which is conceptually difficult to quantify. Other than longevity, changes in lifetime consumption are the main drivers of lifetime utility, and seem to be detrimental to the baby-boom generations. Lower interest rates tend to increase welfare by allowing agents to better smooth consumption over their life cycles.

Finally, we delve further into the open-economy dimension of our model by considering the demographic transition from the

[1] In fact, in a frictionless neoclassical model such as ours, the marginal product of capital is the only determinant of the real interest rate, while it could be one of several in a model with risk premiums, financial frictions, or price markups (Caballero, Farhi, and Gourinchas 2017; Eggertsson, Mehrotra, and Robbins 2019).

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 explain around 20 percent of the dispersion in advanced-economy net foreign asset positions.[2] 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 and other macroeconomic trends over the long run. In common with the other papers in this literature, our model is very stylized, and in particular does not include an account of the risk premiums and financial frictions that may have caused the return on capital to diverge from government bond yields (see, for example, Caballero, Farhi, and Gourinchas 2017; Marx, Mojon, and Velde 2018). The aim of this study is to isolate the effect of demographic change on savings behavior and the real return on capital and other assets.

Our paper is one of several addressing the impact of demographic change on the real interest rate, house prices, or external payments. Many papers focus only on a single economy: see, for instance, Kiyotaki, Michaelides, and Nikolov (2011), Gagnon, Johannsen, and López-Salido (2021), and Eggertsson, Mehrotra, and Robbins (2019) on the United States or Waldron and Zampolli (2010) on the United Kingdom. Gagnon, Johannsen, and López-Salido (2021) and Eggertsson, Mehrotra, and Robbins (2019) model in detail the household's life cycle, but relative to this paper they focus on the role of demographic change on the decline in the natural interest rate in the United States, while Kiyotaki, Michaelides, and Nikolov (2011), conversely, focus on the effect of interest rates and demographics as drivers of land and housing markets. By considering all advanced economies together, instead, our paper relates to Carvalho,

[2]Net foreign asset positions in the data are generally smaller in magnitude than the values predicted by our model, since the model ignores the presence of frictions in international capital flows.

Ferrero, and Nechio (2016) and Marx, Mojon, and Velde (2018) and contributes to that literature by using a fully fledged life-cycle model, examining the effects of demographics on house prices and household credit, as well as accounting more precisely for the high levels of wealth observed among older households, and highlighting the effect of this stock of wealth on the interest rate. Last, our open-economy exercise builds on Domeij and Flodén (2006), Krueger and Ludwig (2007), and Backus, Cooley, and Henriksen (2014), who study international capital flows via current account movements. It is also related to the more recent work of Barany, Coeurdacier, and Gribaud (2018), who take into account country-specific borrowing constraints and social security. Here as well, our paper contributes to the literature by including housing as an alternative store of value, and by highlighting the importance of wealth stocks—and not only flows—for evaluating the impact of demographic change on the macroeconomy.

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 concludes.

2. Demographic Trends

To document 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. Our focus is on an aggregate of advanced economies.[3]

The shares of different age groups in the total population over time, shown in figure 1A, present two main patterns. Firstly, we see a clear rise in the share of older generations, for example, with the over-80s going from 1 percent of the population in 1950 to around 5 percent in 2015 and reaching a projected 14 percent by the end of

[3]In particular, we use Western Europe (Austria, Belgium, Denmark, Finland, France, Germany, Greece, Iceland, Ireland, Italy, Luxembourg, Netherlands, Norway, Portugal, Spain, Sweden, Switzerland, and the United Kingdom), North America (Canada and the United States), Japan, Australia, and New Zealand.

Figure 1. Aging Population

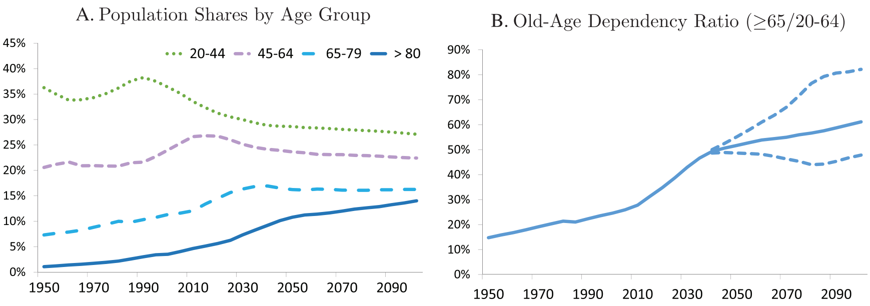

Source: UN Population Statistics (projections based on medium-fertility scenario; dashed lines in panel B show high- and low-fertility scenarios).

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.

This evolution is summarized by the old-age dependency ratio (henceforth OADR), shown in figure 1B, defined as the ratio of over-65s to 20- to 64-year-olds. Again the clear upward trend shows the rise in the share of older generations relative to the working population. The dashed lines, which show 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 percent in 1950 to almost 50 percent in 2100. In the medium-fertility scenario the final number is over 60 percent.

Having documented these trends, we now examine their causes. Figure 2A shows the growth rate of consecutive 20- to 24-year-old cohorts over time. We can see that the period 1970–85 saw elevated growth rates, reaching over 3 percent per annum, corresponding to the baby-boom generations born between 1945 and 1960. Growth rates have since fallen, and have 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 toward the older age groups. This is amplified by the smaller size of the new younger generations entering the population.

Figure 2. The Baby Boom

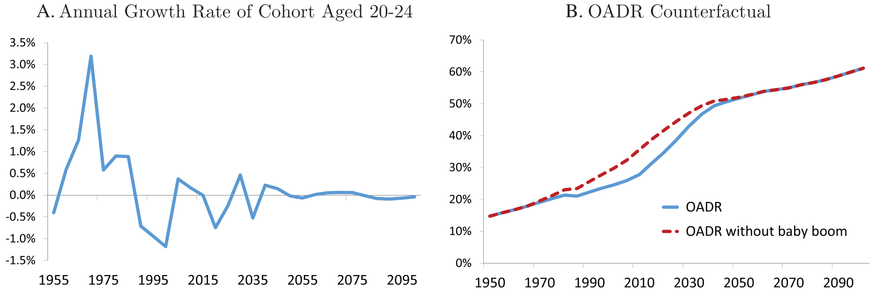

Source: UN Population Statistics (projections based on medium-fertility scenario).

To further illustrate the baby-boom effect on the aggregate demographic trends, figure 2B shows the counterfactual OADR when we assume that cohort growth in 1970–85 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 counterfactual. 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 counterfactual 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 3 shows life expectancy conditional on living to age 60.[4] While a 60-year-old in 1950 would not expect to live past the age of 77, by 2015 a 60-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.

[4]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. Life Expectancy at 60

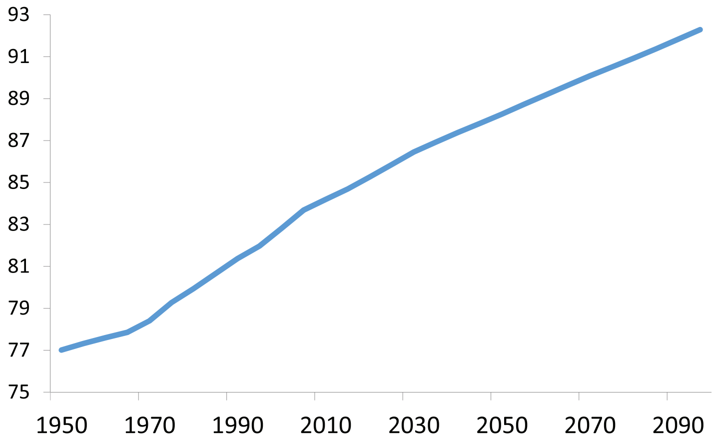

Source: UN Population Statistics (projections based on medium-fertility scenario).

As the data make clear, aging 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 overlapping-generations model to uncover the macroeconomic effects of these important trends.

3. Quantitative Model

3.1 The Model

We consider a general equilibrium setup with overlapping-generation households and a representative firm producing in a perfectly competitive environment.[5] We describe these two agents in turn, and then describe the aggregation and market clearing conditions.

[5]The underlying model is the same as that of Sajedi and Thwaites (2016), allowing for demographic change.

3.1.1 Household

Agents are born at age 1 and can live up to T periods. We denote by $x_{\tau,t}$ the value of a variable x, for a household born in period t, when they are aged τ.

Agents 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 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 labor, l, inelastically, and gain utility from a consumption good, c, and a housing good, h, which is traded at relative price p^h. Hence the T-period optimization 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

$$\begin{aligned} &c_{\tau,t} + a_{\tau,t} + p^h_{t+\tau-1}(h_{\tau,t} - h_{\tau-1,t}) \\ &\quad \le w_{t+\tau-1}\epsilon_\tau l_{\tau,t} + (1 + r_{t+\tau-1})a_{\tau-1,t} + \pi_{\tau,t} \qquad \text{for } \tau = 1, \ldots, T, \end{aligned}$$

where ϵ is the age-specific productivity level, w is the wage per efficiency units of labor, and a is a safe asset with return r.[6] We assume that $l_{\tau,t} = \epsilon_\tau = 0$ for $\tau \ge T^r$.

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 nonlabor income, taken as exogenous by the households.

There are a fixed number of periods when the household is able to "move house," i.e., reoptimize their housing wealth, and hence

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

outside of these "move dates" the household has an additional constraint $h_{\tau,t} = h_{\tau-1,t}$.[7] 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:

$$\lambda_{\tau,t} = \beta_\tau \tilde{\psi}_{\tau,t} c_{\tau,t}^{-1} \qquad \forall \tau = 1, \ldots, T$$
$$\lambda_{\tau,t} = (1 + r_{t+\tau})\lambda_{\tau+1,t} \qquad \forall \tau = 1, \ldots, T-1$$
$$\lambda_{T,t} = \beta_T \tilde{\psi}_{T,t} \phi a_{T,t}^{-1}$$
$$\sum_{j=\tau}^{\tau'-1} \beta_j \tilde{\psi}_{j,t} \theta_j h_{\tau,t}^{-1} = p_{t+\tau-1}^h \lambda_{\tau,t} - p_{t+\tau'-1}^h \lambda_{\tau',t} \qquad \forall \tau \in \text{"move dates,"}$$

where τ' in the last equation denotes the next move date after τ.

3.1.2 Firm

The firm's problem is to choose the aggregate factors of production, K_t and L_t, to maximize profit, taking as given the rental rate of capital, r_t^k, the wage per efficiency units of labor, w_t, and the production function, $Y = F(K, L)$. Note that L_t denotes the aggregate efficiency units of labor 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 constant elasticity of substitution (CES) production function $F(K, L) = \left[(1-\alpha)L^{\frac{\sigma-1}{\sigma}} + \alpha K^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}$, we have the following first-order conditions[8]:

---

[7] This is a simple way to capture the fixed costs associated with buying and selling housing, which makes housing a less liquid asset.

[8] Although changes in total factor productivity (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. Instead, below, we will look at how demographic changes affect labor productivity due to the age-specific productivity levels.

$$w_t = (1-\alpha)\left(\frac{Y_t}{L_t}\right)^{\frac{1}{\sigma}}$$

$$r_t^k = \alpha\left(\frac{Y_t}{K_t}\right)^{\frac{1}{\sigma}}.$$

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

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

3.1.3 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. Normalizing 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 $(T \times 1)$ 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 $(T \times 1)$ 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.

3.1.4 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, $\tilde{A}_t = A_t/S_{t+1} = \frac{S_t}{S_{t+1}}\tilde{\rho}_t'\mathbf{a}_t$.

Labor Market. Aggregate labor supply must equal labor demand. Let $\epsilon\mathbf{l}_t = \{\epsilon_\tau\, l_{\tau,t-\tau+1}\}_{\tau=1}^{T}$ denote the vector of efficiency units of labor 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, $\tilde{H}$.[9] 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.[10] This endowment is distributed across households through nonlabor income, along with the bequests, as detailed below.

Bequests and Nonlabor Income. At each time t, the nonhousing assets and the housing wealth of the generations that died

[9]This is in line with Knoll, Schularick, and Steger (2017), who find that the bulk of the increase in house prices is attributable to the increase in the value of residential land.

[10]The alternative would be to allow this additional housing to be produced, with a technology which transforms the consumption good into housing. This does not materially affect the results, as discussed in section 4.3 and in online appendix C.1, available at http://www.ijcb.org.

in the previous period, as well as the additional housing endowment, added in each period to maintain a stable level of housing per capita, are distributed to living households through bequests (see online appendix A.1, at http://www.ijcb.org, for more details). This nonlabor income is evenly distributed among households above a given age, T^b, while younger households are not entitled to any nonlabor income.[11]

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

$$\tilde{Y}_t = \tilde{C}_t + \tilde{I}_t,$$

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 above 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 five 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 (i) to match life-cycle profiles of labor productivity, housing wealth, and net worth, and (ii) to match 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. Hence we must target these moments at particular points in time. 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

[11] This assumption aims 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.

transition over the past few decades. Full details of the calibration procedure are provided in online appendix A.2.

3.2.1 Demographic Transition

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.

Specifically, we set g_t as the relative size of consecutive 20- to 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.[12]

3.2.2 Data

Life-Cycle Profiles. Given limited cross-country data availability, we will assume that U.S. households are representative of all advanced-economy households in terms of the life-cycle profiles of labor 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 labor 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. To calibrate housing wealth over the life cycle, we take the sum of *"Primary Residence"* and

[12] 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.

"Other Residential Real Estate" in the SCF. We then subtract this from the total *"Net Worth"* to obtain nonhousing assets a.[13]

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 Holston, Laubach, and Williams (2017), and take the average world interest rate between 1970 and 1980. This gives us a target of 3.42 percent.

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 United Kingdom, and the United States. We obtain a target ratio of 147 percent.

Finally, for debt-to-GDP we use the Bank for International Settlements (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 United Kingdom, and the United States. The final target is 40 percent.

3.2.3 Other Parameters

We set the parameters of the CES production function $\sigma = 0.7$ and $\alpha = 1/3$, and the annualized depreciation rate $\delta = 6\%$.

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 labor supply is $L = 0.3$, the value commonly used in the literature. Households are assumed to retire at age 65, corresponding to $T^r = 10$. They start receiving bequests at age $T^b = 7$, i.e., age 50.

In matching the life-cycle profile of housing wealth, we set the "move dates" in the household's problem to $\tau = 1, 5, 11$, corresponding to ages 20, 40, and 70. The final period of life will also be

[13]See online appendix A.3 for details. Note that by matching labor income and assets, consumption over the life cycle will be determined residually by the household budget constraint, and so it is not being matched to data.

Figure 4. Calibration of Life-Cycle Profiles

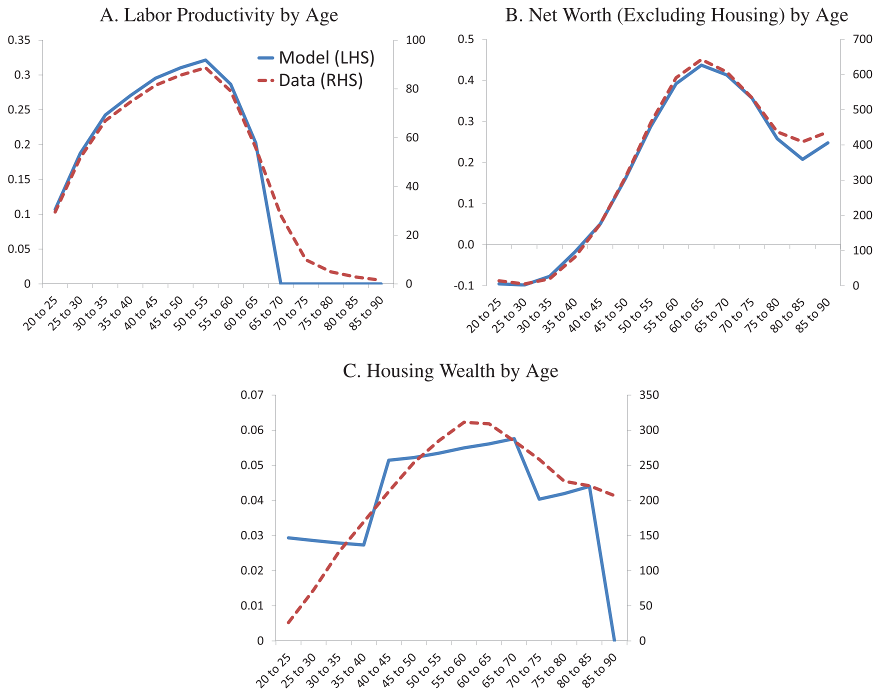

Notes: Model lines are relative to the expected lifetime labor income of the 1950 cohort. Data lines are thousands of U.S. dollars.

a "move date," as we have assumed that agents do not leave housing bequests.

3.2.4 Calibration Outcomes

Figure 4 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 1990–2015. The aggregate moments are matched exactly by construction.

To show how well we fit the demographic trends with the growth rate and death probabilities as exogenous series, figure 5A shows the

Figure 5. Demographics in the Model vs. Data

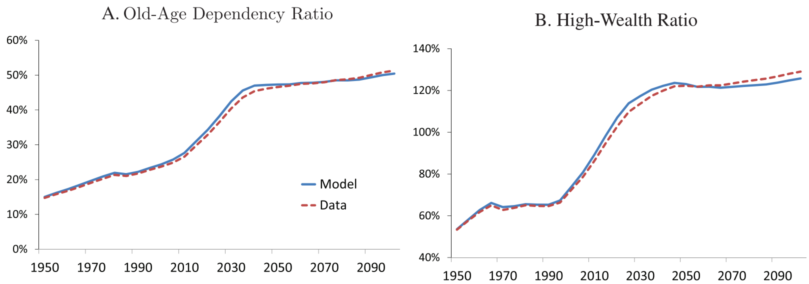

Source: UN Population Statistics and own calculations.

OADR of the model against the data. Figure 5B 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 agent's life. As can be seen from figure 4B, 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 aged 20–49. Looking at figure 5, we see that, despite the simplifications that we make, both the OADR and the HWR in the model are very close to that in the data.[14]

4. Results

We now present the results of the model simulations. Given the exogenous demographic changes described above, we solve for the general equilibrium transition path of the economy, assuming perfect foresight. We show the 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 behavior of households as they expect to live longer. We then explore other

[14]Note that the data line here ignores the population over 90, in line with the model, and so the OADR does not rise as much as in figure 1.

macroeconomic implications of demographic trends: on the housing market, labor productivity, welfare and distribution, and in an open-economy context. To facilitate the comparison with the model outcomes, all data points shown in the figures are five-year averages. Finally, we carry out some robustness exercises to further investigate our results.

4.1 Baseline Results

4.1.1 The Interest Rate, Savings, and Debt

First, we turn to the main outcome of our model, namely the real interest rate, compared with its empirical counterparts. Given that we are using a real model, the real interest rate here should be interpreted as the natural real interest rate: the interest rate that prevails in the absence of nominal rigidities or, equivalently, the interest rate that is consistent with inflation at its target and the output gap closed over the medium to long run. This object is not directly observable, but in figure 6 we show the model outcome against two empirical estimates that are available for advanced economies. Namely, the red dotted line shows the average advanced-economy long-run real interest rate, based on 10-year sovereign yields, taken from Rachel and Smith (2017), and the dashed-dotted yellow line shows the model-based estimate of the natural interest rate from Holston, Laubach, and Williams (2017).[15] These measures of real interest rates are clearly more volatile than the long-term natural interest rate that is captured by the demographic trends in our model.

In the model, the annual interest rate decreases by 157 bps between 1980 and 2015.[16] Compared with the empirical counterparts, between 1980 and 2015, demographics are able to explain 75 percent of the roughly 210 bps drop estimated by Holston, Laubach, and Williams (2017) and around 45 percent of the fall in the Rachel

[15] For figures in color, see the online version of the paper at http://www.ijcb.org.

[16] Since we have calibrated the level of the real interest rate in the 1970s, we focus here on the changes since 1980.

Figure 6. The Real Interest Rate (%)

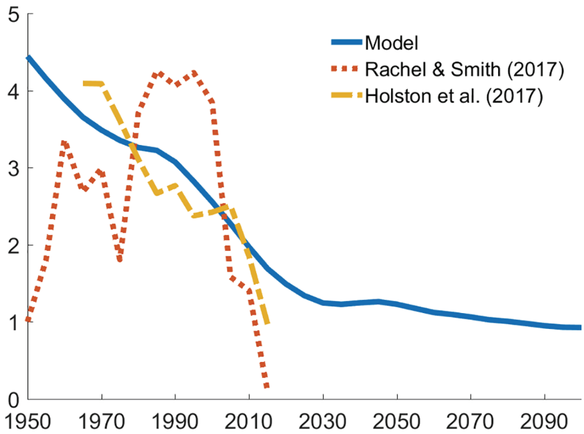

Notes: Holston, Laubach, and Williams (2017) measures the natural interest rate, and is averaged across the United States, United Kingdom, euro area, and Canada. Rachel and Smith (2017) measures the world real interest rate.

and Smith (2017) measure. These empirical measures have 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. Nonetheless, 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 real interest rate in our model is forecast to decrease by a further 46 bps by 2050 and 76 bps by 2100.

The converse of the fall in the real interest rate is a rise in the capital stock, shown in figure 7A against two empirical measures of capital intensity: an index of capital services-to-GDP for 19 advanced countries, in the red dotted line, and an index of the ratio of the capital stock to value added in the U.S. business sector from the Fernald (2012) data, in the yellow dashed-dotted lines.

Figure 7. Capital and Debt

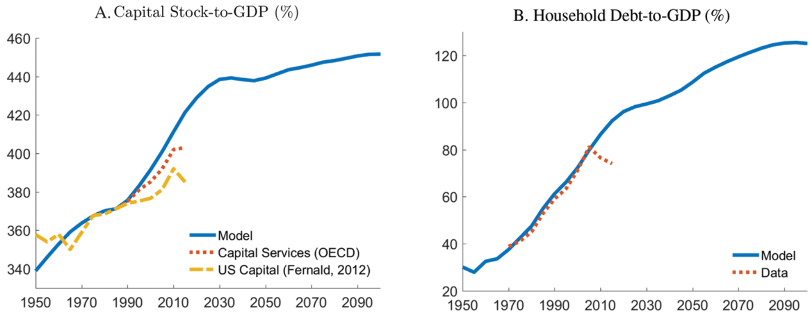

Sources: Capital services are measured by OECD data for Australia, Austria, Belgium, Canada, Denmark, Finland, France, Germany, Ireland, Italy, Japan, Korea, Netherlands, New Zealand, Portugal, Spain, Sweden, United Kingdom, and the United States. U.S. capital is from Fernald (2012). Household debt is from BIS databases; the household debt-to-GDP is the ratio of household debt to nominal GDP.

The demographic trends alone slightly overestimate the rise in capital, particularly for the United States, 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.

The key mechanisms triggered by the demographic transition are the following. Firstly, households perfectly anticipate that they will live longer and spend more time in retirement. They are therefore willing, all things equal, to transfer more of their income during working life to the future, in order to smooth their consumption. Secondly, the slower population growth and increased longevity imply that older households make up a larger share of the total population alive at each period. Were households not anticipating their longer lifetime and adjusting their consumption and saving choices, this change in population weights would still change aggregate savings and aggregate consumption outcomes. In partial equilibrium, these two mechanisms would both increase the level of aggregate savings-to-GDP over time. In general equilibrium, to keep the capital market balanced given this higher capital supply, the interest rate decreases.

Figure 8. Decomposing the Drivers

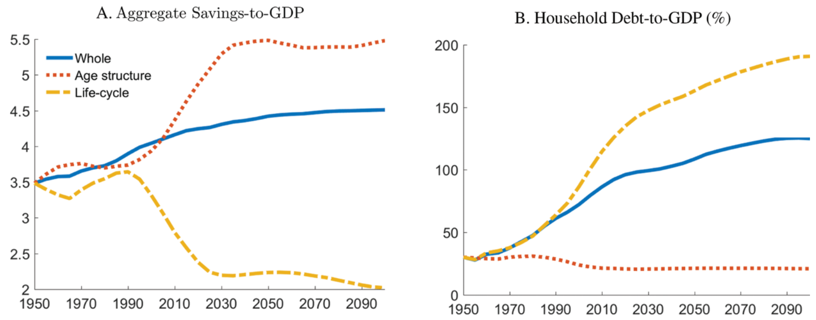

Note: Age structure is changing only the population age structure and life cycle is changing only the household's optimal behavior.

This lower interest rate has an offsetting effect, as it discourages savings and even encourages more borrowing by the young, raising net household debt-to-GDP and pushing down on aggregate savings-to-GDP. As shown in figure 7B, household debt-to-GDP rises in the model in line with the data.

We can see the effects of these two mechanisms by decomposing 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 decisions of each household, given the new general equilibrium prices. The impact of these two distinct drivers on aggregate savings and household debt is shown in figure 8.[17] The marginal impact of changes in the population age structure (shown in the red dotted 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 the 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-to-GDP decreases with the decreasing share of young households in the economy.

[17]See online appendix B.1 for the details of the construction of this decomposition and a second exercise distinguishing between partial and general equilibrium effects on household decisions.

Conversely, taking only changes in optimization over the life cycle into account, shown in the dashed-dotted yellow line, the aggregate savings per capita actually decrease massively from 1990 onwards: since the interest rate is lower, and despite their increased life expectancy, households shift their portfolio toward 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.

We conduct a similar decomposition exercise for household debt-to-GDP, as shown in figure 8B. We reach similar conclusions to the aggregate savings case: the changes in the population structure tend to stabilize or decrease the household debt-to-GDP, while the household reoptimization of their consumption, savings, and borrowing decisions given their increased life expectancy and the lower interest rate implies a higher debt-to-GDP ratio.

4.1.2 Housing

One important feature of our model compared with 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 lifetime, and therefore anticipate the return on housing as a store of wealth over their life cycle. As the interest rate falls, so does the user cost of housing, which is the opportunity cost of investing into an additional unit of housing instead of the financial asset, 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, the housing wealth-to-GDP ratio increases, as shown in figure 9. In fact, we are able to explain 85 percent of the observed increase in real house prices.[18] 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.

[18]This can be seen as an upper bound on this effect, since any increase in housing supply per capita would mitigate the rise in house prices.

Figure 9. Housing in the Baseline Simulations

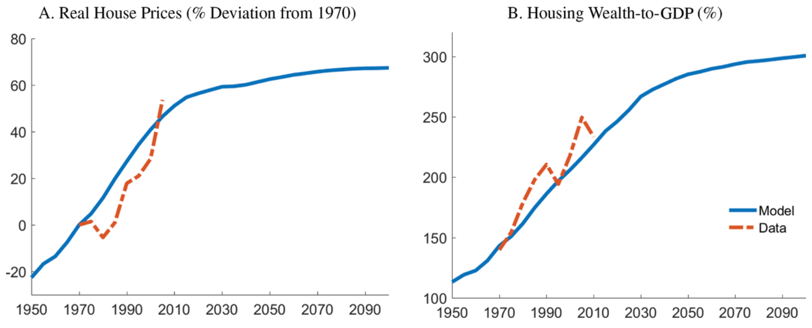

Sources: Housing wealth comes from Piketty and Zucman (2014). Real house prices are from the BIS databases and are the ratio of nominal house prices to the consumer price index.

Housing accordingly provides an alternative vehicle for the transfer of resources over the life cycle, and will raise interest rates.

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 has two implications. Firstly, it means 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 generations (born between 1945 and 1949) reach age 70, so that 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. Concretely, the share of households with the highest housing demand, the 40- to 69-year-olds, decreases, while the share of households with a lower housing demand, aged 70 and over, increases. Consequently, even as the demographic transition continues, the rise in housing demand slows down around this time, and hence house prices flatten out.[19]

The second implication of the discrete move dates is that, as a savings instrument, housing has a longer maturity than capital, meaning that house prices are more forward looking: households

[19]See online appendix C.1 for more details.

Figure 10. User Cost of Housing and Forward-Looking Interest Rate

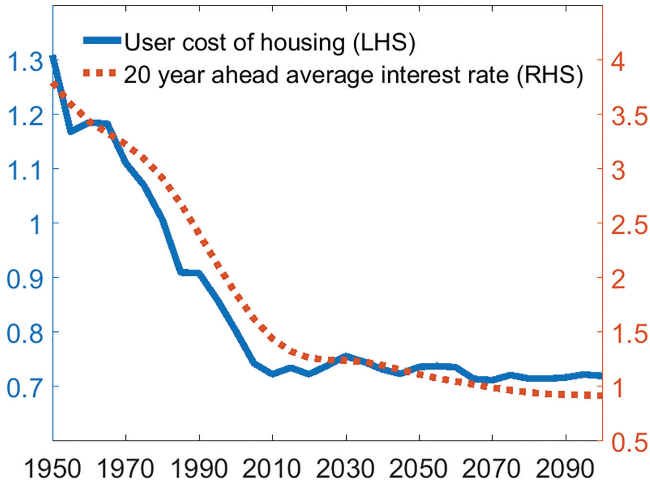

Note: The user cost of housing is the product of the interest rate and the house price, less the capital gain on the house.

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 10, which plots the four-period-ahead (20-year-ahead) average real interest rate against the per-period user cost of housing. These variables follow the transition of house prices more closely than the real interest rate.

4.1.3 Productivity

While we abstracted from trend productivity growth as a driver of macroeconomic trends, our model can partially explain the observed movements in labor productivity through the effect of demographic changes. As shown in figure 4A, 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 in

Figure 11. Implications of Demographic Trends for Productivity

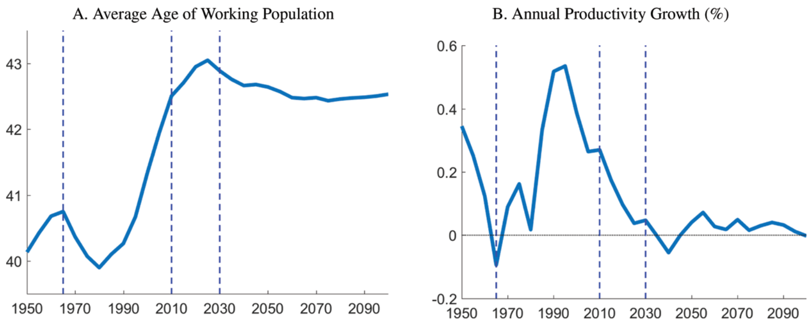

Note: Vertical dashed lines show when the oldest baby-boom generation enters the workforce, retire, and reach age 90.

figure 11. We can clearly see the impact of the baby-boom generations on the figures. From 1970 onwards, the young baby-boomers start working, bringing down the average age of the working population and hence productivity growth. Until 2000, the baby-boomers age and gain in work experience, increasing the labor force's average age and productivity. From 1990, the baby-boomer generations reach ages 50 and above, their productivity decreases, and hence the productivity growth slows down, while the average age of the 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 slowdown in productivity growth, our model shows that the aging workforce may have played a role in this evolution.

4.1.4 Distributional Impact of Demographic Change

Since our model includes heterogeneity in terms of birth year and age, we can measure inequality along two dimensions: across households within period, and across cohorts defined by their birth year. It is important to emphasize, however, that our model is not designed

Figure 12. Gini Coefficients

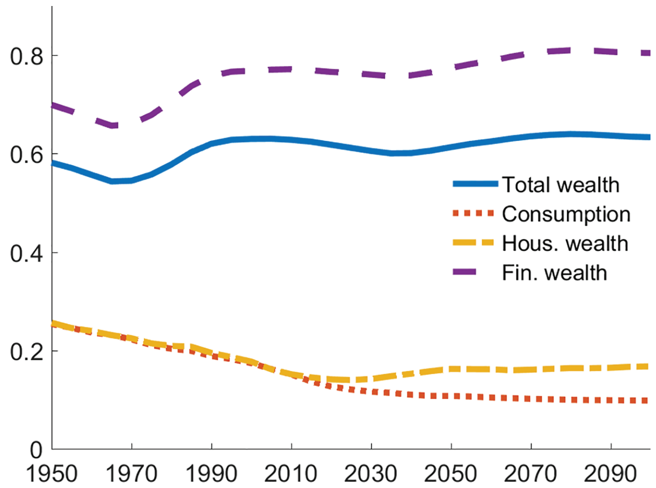

to analyze inequality from either a positive or a normative angle—there is no within-cohort heterogeneity, and there are conceptual difficulties in comparing welfare over changing lifespans.[20] Our results accordingly should not be read as a full intergenerational welfare analysis. Nonetheless, it is interesting to see what insights we can get from the model about distributional impact of demographic changes.

Within-Period Inequality. Figure 12 shows the evolution of Gini coefficients for consumption and financial, housing, and total wealth 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 life-cycle 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

[20]Furthermore, we use model-implied rather than actual relative prices in order to isolate the impact of demographic change, and ignore, as elsewhere in this paper, the impact of other drivers in the data.

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, as the two drivers work in opposite directions. The initial increase from 1965 to 1990 as well as the post-2030 increase 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. Between 1990 and 2030, however, the age composition effect offsets this: 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. If simply comparing expected lifetime utility across cohorts, we would find a clear increase over time, meaning that an agent born in 1980 has a higher expected utility than one born in 1950. This is, however, the simple consequence of increased longevity and hence being able to enjoy consumption over a longer period of time, and necessarily depends on underlying assumptions regarding the value of being alive rather than dead.

To obtain a meaningful comparison, it is necessary to disentangle the various components at play, most importantly to abstract from the change in longevity. To do this, we use the decomposition suggested by Jones and Klenow (2016), which allows us to subtract the longevity component of utility from the total utility variation.[21] The result is shown in the blue solid line in figure 13, in consumption

[21] Full details of this exercise are provided in online appendix B.2.

Figure 13. Decomposition of Expected Lifetime Utility (consumption equivalents, %)

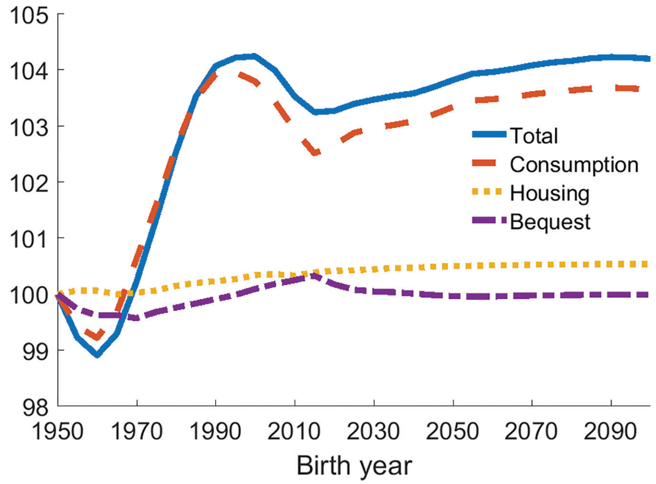

Note: Death probabilities fixed at 1950-cohort level.

equivalents with respect to the 1950 cohort. This is actually equivalent to comparing utility across cohorts, keeping the death probabilities fixed at the level of the 1950 cohort.[22] Once abstracting from longevity effects, the change in realized utility across cohorts is much smaller and does not monotonically increase, with baby-boomers and households born around 2015 appearing to be worse off.

We can further decompose this total utility across components related to housing, bequests, and consumption in the dashed lines in figure 13.[23] The consumption component is clearly the main driver, with both the magnitude and dynamics looking very similar to total utility. The contribution of housing utility slightly increases for later cohorts, due to a better distribution of housing over the life cycle,

[22]The choice of the cohort here affects the precise magnitude of the utility differentials between cohorts, but does not affect the overall dynamics.

[23]Each component is an aggregate of the level and "smoothing" component in the Jones and Klenow (2016) methodology; see online appendix B.2 for the full decomposition.

Figure 14. Lifetime Income and Its Components

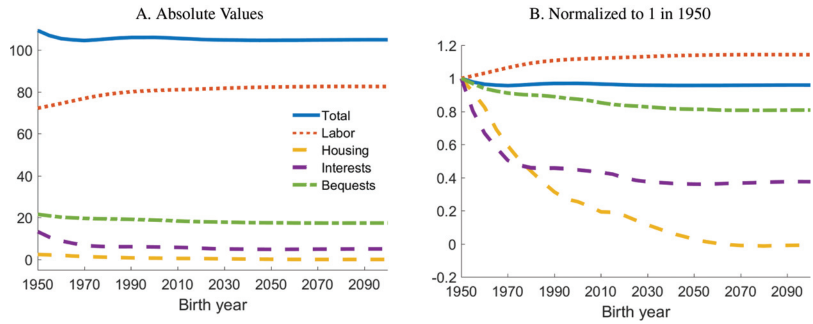

allowed both by the increased ability of young households to borrow and buy housing, and by the changes in the age distribution of the population. The overall contribution of housing utility remains very small, however, as its share in total utility is small (less than 2 percent). Finally, the contribution of utility gained from leaving bequests to the next generations does not show any clear trend over time, and is anyway very small.[24]

Examining income differences across cohorts is another way to understand the origins of their welfare differences. While total income is relatively stable across cohorts, figure 14 shows that this is hiding an increase in labor income and a decrease in capital, housing, and bequests income. Because of the decrease in interest rates and the progressive slowdown 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.[25]

[24]The bequests line corresponds to the additional utility from leaving bequests included in the households' utility function, and not to the utility gains obtained by the households receiving the bequests.

[25]Interest income is equal to the net capital share, which is the product of the capital-output ratio, which ranges between 3 to 4, times the annual real interest rate, which ranges between 1 percent and 4 percent. As shown in figure 9, housing wealth increases by around 1 percent of GDP per year from 1950 to 2030 and then flattens off. Accordingly, these capital gains cannot form a large share of any generation's permanent income, though they accrue disproportionately to the early baby-boom generations.

To conclude, increasing longevity implies an increase in welfare, but its impact is difficult to measure. Abstracting from longevity, changes in lifetime consumption are the main drivers of realized utility and seem to be detrimental to the baby-boom generations. 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 Open-Economy Implications

So far we have considered the group of advanced economies as one closed economy, and looked at the effects of the demographic trends in the aggregate population. While an aging population is common to all these countries, different countries within this group are aging at different speeds. Figure 15 shows the OADR for a handful of countries within our aggregate group. As can be seen, Japan and Germany, for example, are aging much faster than the aggregate, while the United Kingdom and United States are aging more slowly.

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. All else equal, this means that the firms in each country demand the same level of capital relative to output, which can be seen from their first-order condition. There is, however, no market clearing condition for the domestic capital markets, implying that household savings can be above or below the capital demanded by firms. This discrepancy between domestic savings and domestic capital gives rise to a nonzero net foreign asset (NFA) position for the domestic economy. In particular, if domestic savings are higher than domestic capital, then domestic households must place their savings into capital abroad. Conversely, if domestic capital is higher than domestic savings, some of the domestic capital must be owned by foreign households.

Consider a country such as the United States, which is aging more slowly than the average. There, demographic trends are putting less upward pressure on savings, and hence the global real interest rate is below the interest rate that would arise were the United States a

Figure 15. OADR across Countries

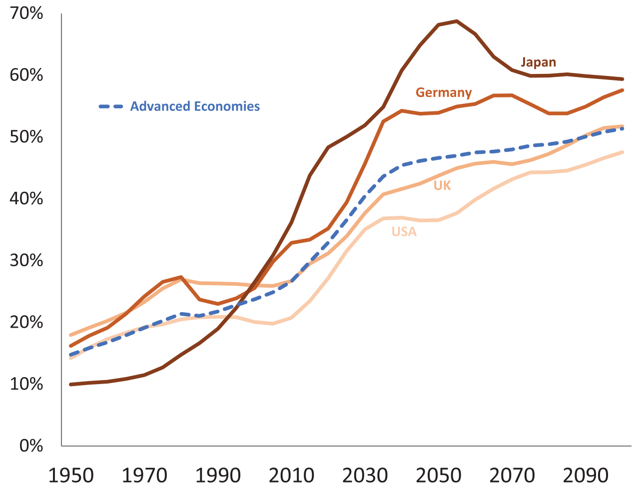

Source: UN Population Statistics (projections based on medium-fertility scenario).

closed economy. In other words, the savings of domestic households in the United States is below the desired capital level of U.S. firms. This translates to a negative NFA position for the United States, as capital flows into the United States from foreign households. Conversely, for a country such as Germany, which is aging faster than the average, the global interest rate is above the rate that would balance 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

Figure 16. NFA-to-GDP in the Model vs. Data (2015)

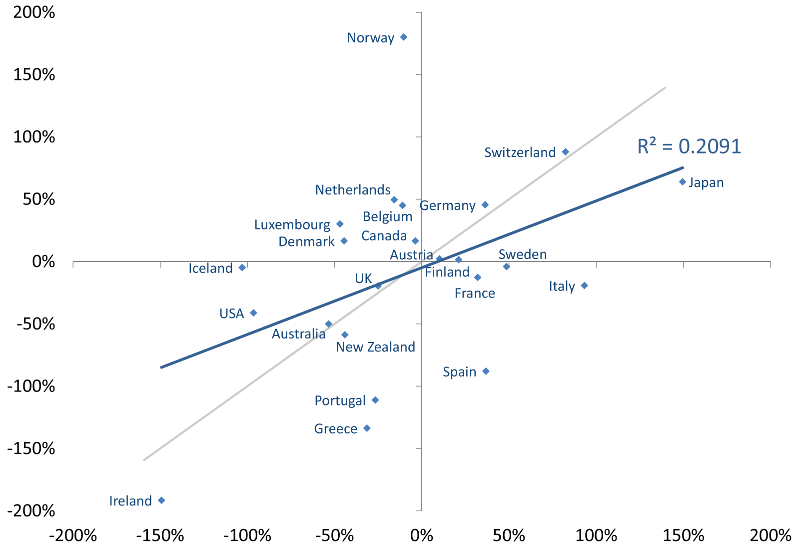

Source: International Monetary Fund's International Financial Statistics.
Note: Model on x-axis; data on y-axis; gray line is the 45-degree line.

rate from the aggregate exercise, as shown in figure 7, and feeding in the demographic variables of a given country.[26]

Figure 16 plots the level of the NFA-to-GDP ratio in 2015 against the predicted level from the model across all the countries in our advanced-economies group. This exercise can be interpreted as a test of the mechanisms of our model against the data. The model omits any frictions in the international movement of capital, such as capital controls or home bias in portfolio allocations, which were important features of the world economy at least in the early postwar period.[27] Hence we see that the model predicts slightly larger NFA positions

[26] 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.

[27] The simulations also assume that the economy is always at the dynamic equilibrium, omitting, for example, the major fiscal and physical consequences of the Second World War.

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 underlying all of the results from our model.[28]

4.3 Robustness and Extensions

Having discussed the various implications of aging population in our baseline model, we now consider some extensions to see the robustness of our results. For the sake of space, all figures are presented in online appendix C.

We begin by comparing the baseline results against a model in which we exclude housing, in order to highlight the role of housing in the baseline (see online appendix C.1). The fall in the interest rate between 1980 and 2100 is slightly larger in the model without housing: there is a 250 bps decline compared with 233 bps in the baseline. This reflects the role of housing as an alternative savings vehicle, which therefore mitigates the decline in interest rates. The rise in the household debt-to-GDP ratio is substantially lower, again reflecting the fact that purchasing housing at a young age is an important driver of higher borrowing in the baseline.

We also investigate the importance of our assumption of exogenously fixed housing supply per capita, by running an alternative exercise where new housing is "produced" from the consumption good, and hence comes out of the resource constraint. This alternative assumption yields an interest rate drop of 224 bps between 1980 and 2100, just 9 bps smaller than the baseline. If anything, this alternative assumption reinforces the role of housing in mitigating the interest rate decline. However, the small differences between this alternative and our baseline results show that the role of housing is overall unaltered by the way it is produced or introduced in the economy, and that this assumption is not important for our results.

[28]See online appendix B.3 for more details on individual countries' NFA paths and model predictions.

Second, we look at the effect of increasing the retirement age (see online appendix C.2). We compare the path of the interest rate in the baseline, in which the retirement age was held constant at age 65, with an alternative simulation in which the retirement age is held constant at age 70, holding other aspects of the calibration fixed. Since we cannot reliably calibrate labor productivity in old age, we carry out each simulation under two different assumptions about old-age productivity. The first, central, case assumes that productivity at age 65–69 is the same as at 60–64. The second, a credible upper bound, assumes that productivity remains at its peak until the end of working life, whether this is 65 or 70.

We find that the potential effect of a higher retirement age, even as high as age 70, is fairly modest in this model, when compared with the effect that demographic changes will have on the interest rate. This is because five years of additional labor income is not enough to offset the rising proportion of time spent in retirement given increasing life expectancy. The extra labor supply is offset in part by more capital, such that the capital-labor ratio, and hence the interest rate, is not affected very much. Even when we assume high productivity in old age, households still face an increasing need to save while they are working in order to finance consumption in the years that they do not have any labor income.

Finally, for comparison with the literature, we use our model to consider the case of the United States more specifically, as a closed economy (see online appendix C.3). The United States has a more dynamic population growth than the average of advanced economies and a life expectancy below average. As expected, we find that the impact of demographic change on the interest rate in the United States is smaller than for the advanced-economy average baseline.

5. Conclusions

In this paper we use an overlapping-generations model, calibrated to advanced-economy data, to assess the contribution of population aging to the fall in real interest rates and other macroeconomic trends which the world has seen over the past three decades. We find that global demographic change can explain three-quarters of the 210 bps 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-boom generations retire: demographic change is forecast to reduce rates by a further 46 bps by 2050. Our model can also explain about 20 percent of the variation of advanced-economy NFA positions.

Among the many uncertainties contained in our analysis, we conclude by highlighting the three most important. The first relates to individual behavior 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, French, and Jones (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 behavior of retirees rather than workers. Similarly, both Bloom, Canning, and Graham (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 caveat relates to the absence of government pension schemes in our model. The implications of this simplification may vary across countries and actual pension systems. Including a fully funded pension system would preserve the impact of aging on aggregate savings, with pension savings supplied by pension funds instead of being directly supplied by households. Including a pay-as-you-go pension system, by redistributing from workers toward retirees via contemporaneous transfers, could reduce incentives to save in anticipation of retirement. However, the household savings in middle age will still tend to rise in response to higher life expectancy, to the extent that public pensions are insufficient to finance a household's desired retirement. Conversely, maintaining a household's income throughout a longer retirement, by taxing a shrinking working population, would raise sustainability issues, and households could again revert to private savings in anticipation of lower state pension payouts in the future.

The third 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. 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.

References

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.

Barany, Z., N. Coeurdacier, and S. Guibaud. 2018. "Capital Flows in an Aging World." CEPR Discussion Paper No. 13180.

Bloom, D. E., D. Canning, and B. Graham. 2003. "Longevity and Life-cycle Savings." *Scandinavian Journal of Economics* 105 (3): 319–38.

Caballero, R. J., E. Farhi, and P.-O. Gourinchas. 2017. "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 (September): 208–26.

De Nardi, M., E. French, and J. B. Jones. 2009. "Life Expectancy and Old Age Savings." *American Economic Review* 99 (2): 110–15.

Domeij, D., and M. Flodén. 2006. "Population Aging and International Capital Flows." *International Economic Review* 47 (3): 1013–32.

Eggertsson, G. B., N. R. Mehrotra, and J. A. Robbins. 2019. "A Model of Secular Stagnation: Theory and Quantitative Evaluation." *American Economic Journal: Macroeconomics* 11 (1): 1–48.

Fernald, J. G. 2012. "A Quarterly, Utilization-adjusted Series on Total Factor Productivity." Working Paper No. 2012-19, Federal Reserve Bank of San Francisco.

Gagnon, E., B. K. Johannsen, and J. D. López-Salido. 2021. "Understanding the New Normal: The Role of Demographics." Forthcoming in *IMF Economic Review*. Available at https://doi.org/10.1057/s41308-021-00138-4.

Goodhart, C., and M. Pradhan. 2017. "Demographics Will Reverse Three Multi-decade Global Trends." BIS Working Paper No. 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 (Supplement 1): S59–S75.

Jones, C. I., and P. J. Klenow. 2016. "Beyond GDP? Welfare across Countries and Time." *American Economic Review* 106 (9): 2426–57.

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): 255–96.

Knoll, K., M. Schularick, and T. Steger. 2017. "No Price Like Home: Global House Prices, 1870-2012." *American Economic Review* 107 (2): 331–53.

Krueger, D., and A. Ludwig. 2007. "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.

Lisack, N., R. Sajedi, and G. Thwaites. 2017. "Demographic Trends and the Real Interest Rate." Staff Working Paper No. 701, Bank of England.

Marx, M., B. Mojon, and F. R. Velde. 2018. "Why Have Interest Rates Fallen Far Below the Return on Capital." Working Paper No. 2018-01, Federal Reserve Bank of Chicago.

Piketty, T., and G. Zucman. 2014. "Capital is Back: Wealth-Income Ratios in Rich Countries, 1700-2010." *Quarterly Journal of Economics* 129 (3): 1255–1310.

Rachel, L., and T. D. Smith. 2017. "Are Low Real Interest Rates Here to Stay?" *International Journal of Central Banking* 13 (3): 1–42.

Sajedi, R., and G. Thwaites. 2016. "Why are Real Interest Rates So Low? The Role of the Relative Price of Investment Goods." *IMF Economic Review* 64 (4): 635–59.

Waldron, M., and F. Zampolli. 2010. "Household Debt, House Prices and Consumption in the United Kingdom: A Quantitative Theoretical Analysis." Working Paper No. 379, Bank of England.