Secular Drivers of the Natural Rate of Interest in the United States: A Quantitative Evaluation*

Josef Platzer† and Marcel Peruffo‡

November 9, 2023

Abstract

We develop a heterogeneous-agent, overlapping-generations model with nonhomothetic preferences that considers the most prominent proposed explanations for the decline in the natural rate of interest (r^*) in the United States. The model accounts for a 4.3 percentage point decline between 1965 and 2015, within the range of empirical estimates. The trend is largely driven by changes in productivity growth, demographics, inequality, and the labor share. Although demographic forces will keep exerting downward pressure on r^*, we expect a trough to be reached around 2030, driven by the rise in public debt. Furthermore, our probabilistic analysis suggests that a reversion toward the levels of r^* seen in the past is extremely unlikely. Finally, we find that the impact of policy variables such as taxes and social security can be considerable.

Keywords: Demographic Change, Inequality, Natural Rate of Interest, Non-homothetic Preferences, Secular Stagnation

JEL Classification: D31, E21, E43, E60

*This paper previously circulated under the title "Secular Drivers of the Natural Rate of Interest in the United States." We are grateful to Gauti B. Eggertsson, David N. Weil, and Neil R. Mehrotra for their guidance and support. We thank Joaquin Blaum, Yann Koby, Alessandro Lin, María José Luengo-Prado, Hannes Malmberg, Pascal Michaillat, Jesse Shapiro, Ludwig Straub, Inês Xavier, Balazs Zelity, and conference and seminar participants at Aix-Marseille, Brown University, Bank of Norway, CEA, CEF, CES, ECB, FRB Boston, Federal Reserve Board, IFABS, IMF, Lubramacro, National University of Singapore, Peking University School of Economics, SEA, Universitat de Barcelona, VAMS, and Wharton Macro Finance workshop for helpful comments and discussions. We also thank Federica di Nicola for superb research assistance. All errors are our own. Josef Platzer kindly acknowledges the generous financial support of the James M. and Cathleen D. Stone Wealth and Income Inequality Project. The views expressed herein are those of the authors and should not be attributed to the IMF, its executive board, or its management.

†International Monetary Fund, email: josef_platzer@alumni.brown.edu

‡University of Sydney, email: marcel.peruffo@sydney.edu.au

1 Introduction

The natural rate of interest (r^*) has declined in recent decades in the United States and other advanced economies. This trend has important consequences not only for the conduct of monetary policy but also for government debt sustainability (Blanchard, 2023), financial stability (Heider and Leonello, 2021; Porcellacchia, 2023), and even aggregate productivity (Liu, Mian and Sufi, 2022). The recent bout of inflation and consequent interest rate hikes have once again brought the question of the current level of r^* and its future outlook to the forefront of policy debates. What are the most important drivers behind the decline in r^* over the last few decades in the United States? Is the low interest rate era over, or are we likely to return to a low interest rate environment? This paper attempts to answer these questions.

Previous research has examined the role of various factors such as demographic change (Gagnon, Johannsen and López-Salido, 2021; Auclert et al., 2021), the rise in inequality (Auclert and Rognlie, 2018; Straub, 2019), the slowdown in productivity growth (Eggertsson, Mehrotra and Robbins, 2019), and public policy variables, particularly public debt (Eggertsson, Mehrotra and Robbins, 2019). While we have a good understanding of the most promising candidates to explain the decline in r^*, there is no comprehensive study that investigates these drivers within a unified framework.[1] This paper fills this gap. We build a heterogeneous-agent overlapping-generations (OLG) model that allows us to not only evaluate the relative importance of the most important (secular) drivers of the recent decline in r^* but also project its likely future dynamics. We consider the role of eight distinct drivers: (i) demographics, (ii) income inequality, (iii) productivity growth, (iv) the labor share, (v) government debt, (vi) government consumption, (vii) out-of-pocket medical expenditures, and (viii) international capital flows. Throughout our analysis, we abstract from business cycles and always focus on the trend natural rate. That is, our study concerns the role of the proposed drivers in determining the long–run r^*.[2]

Our model suggests a decline in r^* of 4.3 percentage points (pp) over the past 50 years, which lies within the range of existing empirical estimates. This trend is largely accounted for by changes in productivity growth, demographics, inequality, and the labor share. The primary explanation for the decline in r^* from the 1960s to the 1980s is the decrease in productivity growth. However, from the mid-1980s until the 2010s, income inequality emerged as the dominating driver. The impact of demographics has materialized only in

[1] Rachel and Summers (2019) conduct such an analysis using distinct frameworks. We discuss the several differences between our paper and theirs later in this section.

[2] See Platzer, Tietz and Linde (2022) for a discussion regarding long– versus short–run natural rates.

recent decades. Looking ahead, we predict a reversion of the downward trend around 2030, driven by the dynamics of public debt. Finally, using stochastic projections for the evolution of demographics, along with reasonable future paths of productivity, inequality, and the labor share, we conduct a projection exercise regarding the future path of r^*. Our results suggest that an increase of the natural rate by more than 75 basis points from late-2010s levels is highly unlikely.

In our framework, we consider 74 generations, with individuals working while young and saving for retirement, insurance against old-age health risks, and bequests. The demographic dynamics are taken from United Nations population statistics and projections. Crucially, our model replicates both the age-earnings and age-wealth profiles observed in the data, capturing the compositional effect resulting from demographic change, as highlighted in Auclert et al. (2021). In addition, we consider two sources of inequality: idiosyncratic and permanent differences in income. By introducing non–homothetic preferences as in Straub (2019), we match the impact of permanent income changes on household savings. This moment is crucial in determining the overall effect of permanent inequality on r^*. Furthermore, we consider exogenous processes for labor productivity growth, the evolution of the labor share, government debt and consumption, international capital flows, and the relative price of medical goods. By feeding their historical paths into our simulations, we can assess their relative importance in driving the decline in r^*. Finally, the model also includes a social security system and a variety of progressive and proportional taxes and transfers, enabling us to examine the impact of those policy instruments on the natural rate.

In our first analysis, we aim to determine the relative importance of each driver over the past five decades. We calibrate the model based on data for 2015 and conduct a decomposition exercise by changing the drivers one at a time and evaluating their individual impact on r^* on a decade-by-decade basis, going back to 1965. We find that the eight drivers jointly explain a 4.3 pp decline in r^*, very close to the estimate by Holston, Laubach and Williams (2017). Quantitatively, changes in productivity growth (1.43 pp), the rise in income inequality (0.96 pp), demographic trends (0.71 pp), and the decline in the labor share (0.6 pp) can explain 87% of the decline in r^* in that period. The rise in government debt counteracts that decline but only to a limited extent (0.29 pp).[3]

Movements in productivity have an important impact from the 1960s to the 1980s (decline) as well as in the 1990s (uptick). By contrast, inequality has more importance after the

[3] In 1965, government debt still exhibited a downward trend, following the spike during World War II. This trend reverted in the 1980s.

mid-1980s. More than two-thirds of the changes due to inequality are due to changes in the permanent component of income, and we show that considering non-homothetic preferences is crucial for this result. Demographics exhibit a limited impact from the 1960s to the 2000s, driven by slowly increasing life expectancy. Starting in the 2000s, however, a rapid increase in the share of elderly individuals in the population, coinciding with the aging of the baby-boom generation, exerts strong downward pressure on r^*. The labor share, which declined from the 1960s to the 1980s and then again starting in the 2000s, is also an important factor. Finally, the remaining drivers have a limited impact, with net capital inflows, medical expenditures, and government consumption accounting for respectively 8%, 5%, and 4% of the total decline. Overall, our results suggest that a combination of drivers, not one single factor, explain the decline in r^*. This means that understanding their dynamics as a whole is crucial in assessing future movements in r^*, the focus of our next analysis.

In our second experiment, we seek to understand what the evolution of the proposed drivers means for the future of the natural rate. Recall that we are interested in the trend (or long-run) r^*, and thus our exercises abstract from business cycles. We initialize the model economy in a post–World War II setting, in 1950. By simulating the evolution of our model economy multiple times—each time using different draws for the path of drivers after 2022—we can create a fan chart for the future path of r^*. In this exercise, we rely on stochastic projections from the UN for demographic drivers and on Congressional Budget Office projections for government debt, while we impose distributions for other drivers based on their historical paths.

The transition path of r^* displays a gradual decline during the 1980s and 1990s, which accelerates thereafter. But our central projection indicates that this past downward trend will reverse in this decade. This is driven by a steady increase in the forecast level of public debt (Congressional Budget Office, 2019), which pushes the natural rate upward. While demographic forces exert further downward pressure on r^* in the future, about 0.4 pp from 2023 to 2100, they are overshadowed by the dynamics of debt. We assume that all other drivers remain unchanged at 2015 levels in the baseline scenario. Even though public debt-to-GDP ratio rises to 200% in our terminal steady state, r^* stays above current levels by only about 20 basis points in 2100. This level is relatively low compared to those observed during the post–World War II period.

Our forecast analysis relies on UN stochastic projections for demographic variables. For the other main drivers, it is difficult to obtain reliable predictions. Thus, we assume that each driver converges to a stationary value drawn from a normal distribution whose

average corresponds to our baseline calibration and whose standard deviation equals half of the change observed from 1995 to 2015. In this case, the two–standard deviation confidence interval for r^* by 2100 is estimated to be -0.2% to 1.7%. When considering demographic forces alone, the two–standard deviation interval narrows to 0% to 1.25%.

Our results also suggest that the downward trend in r^* will probably soon revert. We find that the probabilities that the natural rate will be above its 2020 level in 2030, 2040, and 2050 are respectively 10%, 56%, and 67%. On the other hand, we find that the likelihood that the natural rate will be more than 0.5% above its current level by 2050 lies at only 3%. In summary, our main takeaway from this exercise is that, despite a predicted rise in public debt, it is highly unlikely that the drivers we consider will lead to a significant rise in r^* relative to its current level.

Finally, we use the model as a laboratory to assess the impact of several policy tools on the natural rate, including tax instruments, social security benefits, universal basic income (UBI), subsidization of medical expenditures, and public debt. Our analysis broadly suggests that r^* can be regarded as a *policy choice*, given that policy changes can significantly impact its level. In particular, we find that tax reforms with a redistributive nature can have a considerable impact on the natural rate. For instance, combining an increase of 5 pp in capital income and profit taxes with a UBI policy corresponding to a $2,500 yearly transfer, a 25% tax credit (subsidy) on current out-of-pocket medical expenditures, and an increase of the marginal earnings tax rate at the 90th percentile of income by 3.5 pp raises the natural rate by 0.5 pp in the long term. Additionally, a 1 pp increase in the debt ratio from its 2015 level would raise the steady-state value of the natural rate by approximately 1 basis point. We provide individual estimates of the local impact on r^* for all the policy instruments considered in our model economy.

This paper connects to the literature studying the evolution of r^*. Broadly speaking, this literature can be divided along two dimensions. The first dimension is methodological, with studies either employing econometric techniques (e.g., Laubach and Williams (2003), Holston, Laubach and Williams (2017)) or explicitly modeling economic agents' behavior. The second dimension relates to the estimation horizon, either short or long run. Our paper utilizes a structural model to examine movements in the long-run natural rate of interest. We make two key novel contributions. Firstly, we comprehensively consider all significant drivers within a unified framework. Secondly, we introduce a novel projection exercise based on probabilistic projections of the drivers' evolution.

Closely related to our study is Rachel and Summers (2019). They investigate the decline

in r^* and explore many of the drivers that we also include in our analysis. The crucial difference is that they conduct the analysis using two separate stylized models: one without intragenerational inequality but with three overlapping generations, and the other with idiosyncratic earnings risk but no OLG structure. In contrast, our analysis takes place within a unified framework.[4] In addition, the demographic structure in our model allows us to match the exact composition of the population at any given time, a prerequisite for obtaining an accurate estimate of the impact of demographics on aggregate savings (Auclert et al., 2021). Furthermore, our model–implied age–wealth profiles and the elasticity of consumption to changes in permanent income closely align with the data, which is essential for capturing the effect of both demographic changes and inequality.

Auclert et al. (2021) investigates the impact of demographic change on macroeconomic variables, including r^*, on the world economy. While our focus is on the United States, we draw valuable insights regarding the influence of demographics from their study. Another related paper is Straub (2019), which emphasizes the significance of non-homothetic preferences in explaining the observed relationship between savings rates and permanent income. We incorporate those insights in our unified framework, and we also show that non–homotheticities are important in determining the overall impact of inequality on r^*. Furthermore, we also consider the role of changes in idiosyncratic risk.

Finally, several other papers have analyzed changes in drivers in isolation: Gagnon, Johannsen and López-Salido (2021), Carvalho, Ferrero and Nechio (2016), and Jones (2018) investigate the role of demographics; Eggertsson, Mehrotra and Robbins (2019) considers demographics, public debt, and productivity growth; Auclert and Rognlie (2018) examines the role of idiosyncratic risk; Mian, Straub and Sufi (2021) examines the role of different savings rates across the income distribution using an empirical approach; Cesa-Bianchi, Harrison and Sajedi (2022) analyzes different assumptions regarding household forward-looking behavior for the path of r^*; Coeurdacier, Guibaud and Jin (2015), Lisack, Sajedi and Thwaites (2019), and Barany, Coeurdacier and Guibaud (2018) consider international capital flows; finally, Caballero and Farhi (2014) and Kopecky and Taylor (2020) discuss the role of risk, both internationally and domestically. We acknowledge that the risk premium (potentially driven by a scarcity of safe assets) can also be an important driver of r^*, but we do not incorporate it in our model because of the computational challenges involved.

The structure of the paper is as follows. Section 2 describes the drivers we include in our analysis, their historical evolution, and why they are expected to have an impact on r^*.

[4] A unified framework is arguably better suited to capturing the incentives to save due to different factors (consumption smoothing, precautionary savings, non-homotheticities, retirement, bequests, and others).

Section 3 develops the model, and Section 4 is dedicated to functional-form assumptions and model calibration. Section 5 shows the results concerning the evolution of r^* in the past, while Section 6 looks at the future. Section 7 evaluates the impact of policy on the natural rate, and Section 8 conducts a sensitivity analysis. Finally, Section 9 concludes.

2 The Drivers

In this section, we briefly present the recent evolution of the eight drivers considered in our framework, along with the proposed mechanisms through which they influence the natural rate of interest. Note that their order of presentation does not imply order of importance.

The first driver is demographics. Individuals in the US—as well as in other developed countries—are living longer. Figure 1, in the bottom panel, shows that life expectancy at birth has grown from 71 years in 1970 to nearly 79 years at the onset of the COVID pandemic. Projections by the United Nations suggest that this growth is expected to continue in the foreseeable future. Concurrently, lower population growth rates (top panel) contribute to a shift in the population composition, characterized by a higher proportion of elderly individuals, as depicted in the bottom panel.

Trends in demographics can affect r^* through two primary channels. First, as individuals anticipate longer retirement, they tend to increase their desire to save during their working years. Second, elderly individuals typically hold more wealth in comparison to their younger counterparts. Consequently, an increase in the proportion of older people leads to a rise in aggregate savings, even if individual savings decisions remain unchanged. This mechanism is known as the *composition effect* (Auclert et al., 2021). Both forces exert downward pressure on r^*.

The second driver featured in our framework is income inequality, which has increased notably in the United States. Figure 2 illustrates this trend, indicating a rise of approximately 9 pp in the top decile's earnings share from 1970 to the 2010s. The dashed line shows that most of that increase was driven by the top 1%.

Income inequality's impact on aggregate savings and, consequently, on r^* occurs through two mechanisms. First, rises in inequality stemming from changes in idiosyncratic earnings risk lead to an increase in precautionary savings (Auclert and Rognlie, 2018). Second, a rise in inequality in permanent income, which is unrelated to risk—for instance, a rise in the college wage premium—will impact aggregate savings if people in different parts of

Figure 1: Evolution and Projection of Selected Demographic Indicators

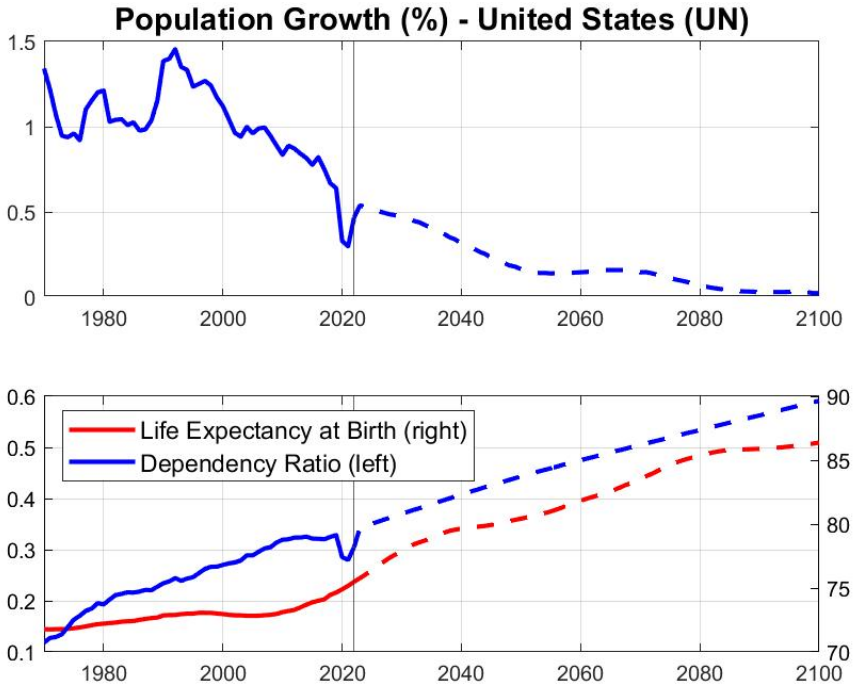

Notes: Source: United Nations (2022*b*), authors' own calculations. Life expectancy at birth refers to both sexes. The dependency ratio corresponds to the ratio of individuals aged 18 to 65 to those 65 or older. Future estimates are based on "Medium (fertility) projection scenario," which includes median estimates for fertility and mortality rates, as well as a continuation of recent levels of net migration into the future.

the distribution of permanent income exhibit different savings rates throughout their life cycle. Recent work (Straub, 2019), which we discuss in more detail below, shows that this is indeed the case. In particular, high permanent income individuals exhibit higher savings rates. Consequently, a rise in both idiosyncratic risk and the permanent component of income inequality increases aggregate savings, once again implying downward pressure on r^*.

The third driver is productivity growth, shown in Figure 3, top-left panel. Total factor productivity, though highly volatile, grew strongly during the 1960s and early 1970s, declined from then till the 1980s, and picked up again during the 1990s and early 2000s, only to return to a sluggish pace by the 2010s. Productivity growth affects interest rates by determining the demand for investments as well as households' consumption-smoothing behavior. Broadly speaking, if earnings are expected to rise rapidly, present consumption will be high relative to *current* income. However, if productivity growth slows down, the

Figure 2: Top Labor Income Shares

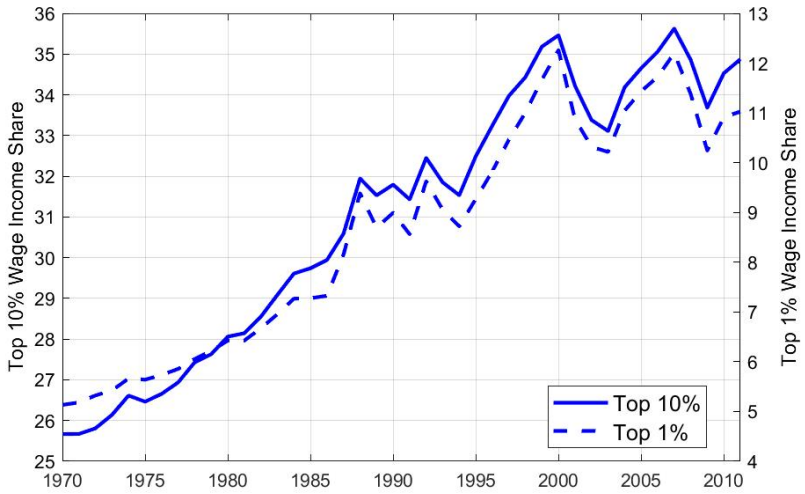

Notes: Source: Piketty, Saez and Zucman (2018). The lines depict the share of income accruing to the top 1% and top 10% of pre-tax wage earners.

desire for current consumption subsides, leading to higher savings rates and downward pressure on r^*.

The fourth driver in our framework is the labor share, which represents the portion of total income that remunerates labor. As illustrated in Figure 3 (central top panel), the labor share has declined in recent periods, with a slight interruption during the 1990s and 2000s. The flip side of the labor share is the capital share, including profits and compensation for renting out capital. Accordingly, a lower labor share implies higher returns on wealth, making savings more appealing. This mechanism, in turn, exerts downward pressure on r^*.

Next, we turn to two drivers that are primarily influenced by policy decisions. The first is government debt, which has significantly increased, particularly in the last decade, as shown in the top-right panel of Figure 3. Furthermore, the Congressional Budget Office (2022) projects that national debt will continue to rise as a percentage of GDP. Higher government-debt-to-GDP ratios translate into higher demand for private savings and thus, all else equal, exert *upward* pressure on r^*. Finally, government consumption expenditures, often referred to simply as G, have declined since 1970. An increase in government consumption, with consequent increase in taxes, crowds out private savings by reducing disposable income, thus exerting upward pressure on $r*$.

The seventh driver in our framework relates to international capital flows, as depicted in

Figure 3: Evolution of Selected Drivers

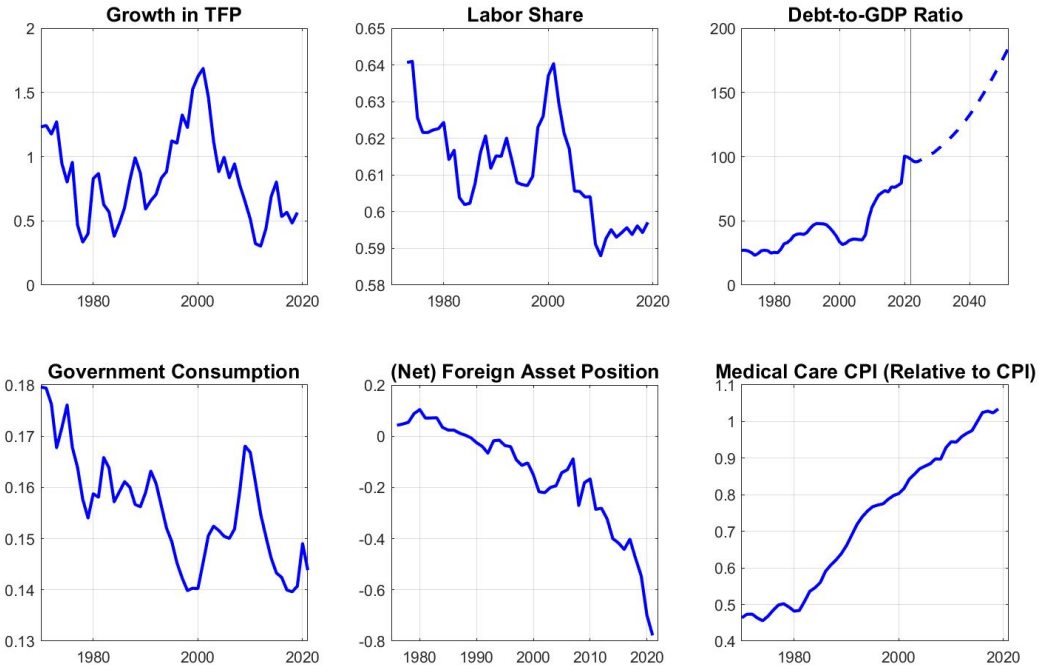

Notes: Sources: Fernald (2014) (productivity), Feenstra, Inklaar and Timmer (2015) (labor share), Federal Reserve Economic Data (FRED) (labor share, government consumption, CPI of medical expenditures), Bureau of Economic Analysis (government consumption, foreign asset position), Congressional Budget Office (2022) (debt-to-GDP ratio and projection; Figure 1-8).

the bottom central panel of Figure 3. Since the 1990s, the United States' net-foreign-asset position has significantly deteriorated. The higher supply of dollar-denominated savings contributes to falling interest rates.

Finally, our framework also considers late-life out-of-pocket medical expenditures. The bottom-right panel of Figure 3 shows that the increase in the price of medical goods since 1970 has been roughly 2.5 times as large as the increase in the Consumer Price Index (CPI). This movement, especially in conjunction with increased longevity, leads to both increased savings throughout one's working life, as retirement becomes more costly, and higher (precautionary) savings in retirement, given the additional risk faced by elderly individuals.

3 Model

In this section we outline the model. Functional-form assumptions and the calibration are discussed in the next section. Our modeling choices are meant to capture the most important aspects of equilibrium determination in savings markets. On the savings-supply side, households[5] are subject to uninsurable idiosyncratic risk and permanent productivity differences as well as changes in labor productivity throughout their life cycle. Individuals' savings decisions are driven by life-cycle, precautionary, and bequest considerations. On the savings-demand side, a representative producer invests in capital, and a central government issues debt and taxes its citizens to fund its activities, including government consumption and a social security system. Finally, we also include international capital flows. In what follows, we describe each of these features. In the interest of space, further details are relegated to Appendix B.

3.1 Demographics

We model a closed economy in which time is discrete ($t = 0, 1, 2, ...$) and there are $G^d = 74$ overlapping generations at any time. Individuals enter the model at biological age 26 ($g = 1$). Individuals face an age-specific profile of survival probabilities and die with certainty at biological age 100 ($g = 74$). An individual of age g at time t will be of age $g + 1$ at time $t + 1$ with probability $p_{g,t}$ or will die with probability $1 - p_{g,t}$.

The growth rate of the population N_t is given by n_t such that $N_{t+1} = (1 + n_t)N_t$. In addition, $\phi_{g,t}$, $g = \{1, 2, ..., G^d\}$ denotes age-specific population shares. The demographics driver consists of these three objects—growth rate of population, survival probabilities, and age-specific shares. In Appendix B, we explain how these three sets of parameters imply an age-specific profile of migration flows, and we provide further details on their treatment.

3.2 Households

An agent is indexed by i, but we often suppress this notation (as well as time subscripts) for readability. Households have offspring, but we do not track their lineage across generations, and all agents in the model are adults.[6] Instead, we calibrate the distribution of inherited bequests to US data, as we explain in the next section.

[5] We henceforth use the words "households" and "individuals" interchangeably.

[6] Gagnon, Johannsen and López-Salido (2021) explicitly account for dependents' consumption. Their results on the quantitative impact of demographic change on r^* are insensitive to this assumption.

At birth, individual expected utility is given as follows:

$$V = \mathbb{E} \sum_{g=1}^{G} \beta^{g-1} s_{g-1} \Big(p_{g-1} u_g(c) + (1-p_{g-1}) v_g(a) \Big)$$

Here β is the discount factor and s_g represents the expected probability of surviving from birth until age g, so that $s_g = \prod_{i=1}^{g-1} p_i$. Agents derive instantaneous utility from consumption c according to function $u(\cdot)$. Upon death, wealth a is transferred to surviving generations. The function $v(\cdot)$ represents a warm-glow bequest motive. The instantaneous-utility functions depend on age g for reasons explained below.

Agents in this economy move through two stages of life: work and retirement. Working individuals supply their labor to firms in exchange for wages. Retired individuals can still supply some labor, and they obtain additional income from their savings and from the social security system. Individuals enter retirement at age $g = \bar{g}$.

Each agent i is endowed with an individual productivity state z and a health state h. We cover these in more detail in the following subsections. Similarly to Auclert and Rognlie (2018), we model households as having access to a single savings instrument, which we denominate interchangeably by savings, assets, or simply wealth. This is intermediated by a mutual fund, which invests in government bonds and physical capital, owns claims to pure profits, and distributes the returns to households on a period-by-period basis.

3.3 The Consumer's Problem

The timing of events is as follows: Productivity realization (z) and current health status (h) are revealed at the beginning of the period. Whether a household survives until $t+1$ is revealed after all decisions at time t have been made. Assets a_{t+1} of a household alive at t but deceased at $t+1$ will be available for households in period $t+1$ inclusive of after-tax returns.

The household's problem in recursive form is as follows:

$$V_g(z,h,a) = \max_{a' \geq 0} u_g(c) + \beta p_g \mathbb{E}_{z',h'} V'_{g+1}(z',h',a') + \beta(1-p_g) v_{g+1}(a')$$

In the equations above, primed variables denote time $t+1$. For workers, the budget

constraint is as follows:

$$(1+\tau_c)c + a' = \big(1+ret\big)a + (1-\tau^{beq})beq + wz - T_y(wz)$$

Here a denotes wealth, ret is the after-tax return on savings, w is the wage, beq represents bequests received, τ^{beq} is the tax on bequests, z is individual productivity, and T_y is the labor income tax. There are no nominal rigidities, so all variables are real. Workers face a borrowing constraint, so $a \geq 0$ at all times. Note that h has no contemporaneous impact on the household budget for workers, but it affects individuals at old age.[7]

At age $g = \bar{g}$, individuals begin to accrue social security benefits and their health status h determines their medical expenditures. The budget constraint is then as follows:

$$(1+\tau_c)c + p_m m_g(h) + a' = \big(1+ret\big)a + (1-\tau^{beq})beq + wz - T_y(wz) + \xi + T^M$$

Here ξ is the retiree's social security income, $m_g(h)$ and p_m are the quantity and relative price of out-of-pocket (OOP) medical goods and services, and T^M is a means-tested transfer to ensure a minimum level of consumption, as in Kopecky and Koreshkova (2014).

Retirees face OOP medical expenses m_g. Several studies—for example, De Nardi, French and Jones (2010), and Kopecky and Koreshkova (2014)—highlight the importance of medical expenses for savings decisions of households, especially in later stages of life. To the best of our knowledge, we are the first to examine the role of OOP medical expenses in the decline of r^*.

We relegate the details of the transfer T^M to Appendix B.3. In practice, this transfer plays a minimal role in our analysis, as only 0.03% of individuals receive a positive transfer of this type in our baseline calibration.

3.4 Earnings Process

The productivity $z_{i,g,t}$ of individual i of age g at time t is given by

$$\log z_{i,g,t} = e_i + h_g^z(e_i) + \zeta_{i,t}(e_i). \tag{1}$$

Productivity is the combination of three terms. First is a permanent productivity term, $e_i \in \{e^L, e^H\}$. Even though the values of e^L and e^H are allowed to change over time,

[7] Health status h is still a state variable for workers because it consists of a persistent process, as we detail in the calibration.

households draw their *type* at birth and keep it until death.

Second, $h_z^g(e_i)$ is a deterministic term that changes with age and thus generates an age profile of productivity. It depends on the permanent-productivity draw e_i, as we explain in the calibration section. Third, $\zeta_{i,t}$ consists of a persistent idiosyncratic component.

If $e_i = e^L$, the process for $\zeta_{i,t}$ is simply an autoregressive process:

$$\zeta_{i,t}(e^L) = \rho\zeta_{i,t-1}(e^L) + \epsilon_{i,t}^{\zeta} \tag{2}$$

Here $\epsilon_{i,t}^{\zeta}$ is distributed normally with mean zero and variance σ_ϵ^2.

Alternatively, if $e_i = e^H$, we augment the autoregressive process described above with a star-earner state, aimed at capturing the top 1% earnings share. We implement it as follows: First, we discretize the autoregressive process in equation 2 using the method described in Tauchen (1986). We then assume that, in every period, individuals with the highest value in the grid for $\zeta(e^H)$, henceforth $\bar{\zeta}$, can become star-earners with probability p^*, in which case $\zeta = \zeta^* > \bar{\zeta}$. Similarly, in every period, star-earners' labor productivity can revert to $\bar{\zeta}$ with probability p_*.

The permanent productivity term e_i can be interpreted as an ability endowment or skill endowment obtained through education. Individuals with high permanent productivity still face idiosyncratic risk, but their expected lifetime income is permanently higher than that of low-e_i individuals. The deterministic age-dependent term $h_z^g(e_i)$ matches a hump-shaped labor-income profile as found in the data. The persistent-shock term $\zeta_{i,t}$ introduces idiosyncratic risk to earnings. Finally, the fact that the star-earner state is risky is motivated by the data in Guvenen, Kaplan and Song (2014), which finds that households in the top 1% of the earnings distribution have a high probability ($p_* = 25\%$) of not being in the top 1% in the following year.

In our counterfactual analyses, when accounting for the impact of *inequality* on r^* we alter three sets of parameters: the permanent income parameters e^L and e^H, the top 1% parameter ζ^*, and the variance parameter σ_ϵ^2. We explain the implementation details in the calibration section.

In the remainder of this paper, we often refer to the inequality induced by the permanent-income and top 1% parameters as *permanent inequality* and to the inequality induced by the variance parameter as *inequality due to idiosyncratic risk* (or due to earnings risk).[8]

[8]Our distinction is motivated by the fact that the top 1% parameter ζ^* has a strong impact on the top 1% earnings share, as does e^H on the top 10%. We thus bundle these two parameters together into *permanent*

3.5 Production

The production function of intermediate goods assumes a Cobb-Douglas form:

$$Y_t^I = K_t^{\alpha} \left(A_t L_t\right)^{1-\alpha} \tag{3}$$

In the expression above, L_t is labor, given in efficiency units, K_t represents capital, A_t is productivity, and $A_{t+1}/A_t = (1+\gamma_t)$, with γ_t being the productivity growth rate.

Monopolistically competitive retailers purchase the intermediate good and differentiate it into varieties using a linear production function. These are then combined into the final good via a standard CES aggregator such that output $Y_t = \left[\int_j y_{jt}^{R\frac{1}{\mu_t}} dj\right]_t^{\mu}$. The demand for each variety is given by $y_{j,t}^R = \left(\frac{p_{j,t}}{P_t}\right)^{\frac{\mu_t}{1-\mu_t}} Y_t$. Normalizing the price of a unit of the output Y_t to $P_t = 1$ and imposing a symmetric equilibrium, the profit-maximization problem of retailers yields the price for the intermediate good as

$$p_t^I = \frac{1}{\mu_t}.$$

Thus, retailers are the source of pure profits in our economy. These are distributed to individuals in proportion to their shares in the mutual fund (wealth).

Given the structure above, the real wage is as follows:

$$w_t = \frac{1-\alpha}{\mu_t} \frac{Y_t}{L_t} \tag{4}$$

The marginal revenue product of capital is given as follows:

$$MRPK_t = \frac{\alpha}{\mu_t} \frac{Y_t}{K_t} \tag{5}$$

Empirically, rates of return on real capital typically exceed rates of return on government debt owing to, among other things, risk and liquidity premia, real intermediation costs, and regulatory capital requirements in the banking system. To capture these in a simple way, we introduce an intermediation wedge, denoted by ν, which is considered deadweight loss.[9] Further, we assume that capital depreciates at an annual rate δ. Thus, the net (pre-tax)

income.

[9] In Section 4, we explain why our results are unchanged if we assume that the proceeds are paid to the mutual fund.

return on capital is as follows:

$$r_t^K = MRPK_t - (\delta + \nu)$$

Finally, pure profits are given as follows:

$$d_t = \left(1 - \frac{1}{\mu_t}\right) Y_t \tag{6}$$

3.6 Mutual Fund

The mutual fund intermediates total household wealth and external capital flows NFA_t, invests in capital and government bonds, and own the retailers. Every period, the returns to these investments and retailers' profits are distributed to households in proportion to their fund shares. Along any perfect-foresight path (no aggregate risk), there is a no-arbitrage condition between the (net) rate of return on capital, r_t^K, and the return on holding government bonds, the latter being the natural rate of interest r^*:

$$r_t^* = r_t^K \tag{7}$$

The after-tax return on wealth ret_t can then be determined by aggregating the total rents from government bonds, capital returns, and pure profits in the economy and distributing them proportionately to wealth (i.e., mutual fund holdings), including the wealth held by foreigners. Total after-tax rent income D_t is given as follows:

$$D_t = (1 - \tau^k)(r_t^* B_t + r_t^k K_t) + (1 - \tau^{corp}) d_t$$

The total supply of savings in the economy is given by the sum of individual savings and the net-foreign-asset position. These savings are demanded by the government and by the intermediate-goods producer. The following equation, expressed in terms of resources, determines the market–clearing condition for savings in the economy at time $t - 1$:

$$B_t + K_t = \sum_g \sum_{z,h,a} g_{t-1}^a(z,h,a,g) \lambda_{t-1}(z,h,a,g) - NFA_{t-1} \equiv \mathcal{W}_t \tag{8}$$

Here $g^a(\cdot)$ is the policy function for savings, NFA_t is the net-foreign-asset position at time t, and $\mathcal{W}_t$ denotes aggregate wealth.

The return per unit of wealth ret_t can then be obtained by dividing D_t by total wealth:[10]

$$\frac{D_t}{B_t + K_t} \equiv ret_t = (1 - \tau^k)\left(r_t^* \frac{B_t}{\mathcal{W}_t} + r_t^k \frac{K_t}{\mathcal{W}_t}\right) + (1 - \tau^{corp})\frac{d_t}{\mathcal{W}_t} \tag{9}$$

The return on savings is thus a composite of the return on government bonds, the return on capital, and pure profits. Consequently, changes in the profit share in the economy will not only have a direct impact on the return on capital and thus on r^* (expression 7) but also have an indirect effect through its impact on the return on savings via dividends and, thus, on individuals' consumption choices.

3.7 Government

The government has four roles in our model economy. First, its debt determines the supply of bonds, thus having a direct impact on the demand for savings and, consequently, on the equilibrium interest rate. Second, it establishes a social security system. Third, it consumes G, which includes other (non-interest) government spending we observe in the data. Fourth, it runs a system of taxes and transfers.

The government's flow budget constraint is given by the following equation:

$$B_{t+1} = (1 + r_t^*)B_t + (\Xi_t - \Lambda_t) \tag{10}$$

Here B_t is the stock of government debt in period t, Ξ_t corresponds to total expenditures, and Λ_t corresponds to total tax revenues, so that the expression in parentheses is the government's primary deficit.

The revenue side of the government consists of the proceeds from the following tax instruments: consumption tax τ_c, capital income tax τ_k, estate tax τ_{beq}, tax on dividends τ_d, and labor income tax $T_y(\cdot)$. Finally, we assume that the government and the social security system are consolidated. In reality, the US social security agency is its own entity. Our treatment can be seen as implicitly assuming that the government would bail out social security in the event of solvency concerns.

[10] We assume that the mutual fund assigns individuals the same portfolio, in terms of the division between capital and government bonds.

3.8 Equilibrium

We reformulate the model in terms of effective labor. For any variable X, let $\hat{X}_t \equiv \frac{X_t}{N_t}$, $\check{X} \equiv \frac{X_t}{A_t}$, and $\tilde{X} \equiv \frac{X_t}{A_t N_t}$. In the steady-state equilibrium (or balanced growth path), variables expressed in tilde notation are constant. The full set of equilibrium conditions can be found in Appendix B.5.

4 Functional-Form Assumptions and Calibration

Our baseline economy is aimed at capturing features of the US economy in the 2010s. A model period corresponds to one year. Unless otherwise stated, the moments computed from the data for our calibration are based on averages from 2010 to 2019. We categorize the parameters to be calibrated into two groups: externally and internally set parameters. The former are calibrated without requiring the solution of the model. They either have a one-to-one correspondence with objects in the data or consist of deep structural parameters for which we rely on estimates from the literature. In contrast, internally set parameters are calibrated using a moment-matching procedure. Section 4.7 provides further details.

We run two main experiments. The first is a comparison between hypothetical steady states. We consider one steady state as reflecting one particular decade. For instance, a counterfactual steady state for the 1970s will consider the values of the drivers to be at their average during the period 1970 to 1979. We typically refer to each decade by a single year—2015 for the 2010s, for instance.

The second experiment is a transition path analysis. The starting date is 1950, and we set the terminal date to 2200, ensuring that the economy has fully converged. Until 2022, we feed the model the realized value for each of the drivers.[11] Looking ahead, we make different assumptions regarding the evolution of each driver, as we explain below. The remaining parameters are taken from the calibration at the 2015 steady state.

We describe the calibration of the 2015 steady state in this section. An overview of all calibrated parameters can be found in Tables 1 and 2. Changes made when simulating past steady states are discussed and summarized in Table D7 of Appendix D. Details on the path of the driving variables along the transition path can be found in Appendix E.

[11] If data for 2022 are not available we keep the latest available data point constant until 2022.

4.1 Utility Functions

Consumption. Instantaneous utility from consumption c for an individual of age g is given by the following:

$$u_g(c) = \frac{(c/o)^{1-\sigma_g}}{1-\sigma_g} \tag{11}$$

We choose consumption preferences of the direct-addilog form as introduced in Houthakker (1960) and recently employed in Straub (2019). $\sigma_g > 0$ are age-dependent coefficients of relative risk aversion. o is a scaling term.

Assuming $\sigma_g = \sigma \; \forall \; g$ implies—ignoring for now the role of bequests—the familiar case of homothetic consumption preferences. We now briefly provide the intuition behind why addilog preferences generate non-homotheticities.

Let us interpret consumption at different ages as separate goods. In the absence of credit constraints and ignoring transfers, bequests, and the progressivity of the earnings tax rate, expenditure shares of each good are independent of permanent income (Straub, 2019). As a result, permanent income elasticities of consumption are equalized among all goods and, in particular, are equal to 1. Consequently, in this world, rises in income inequality due to changes in the permanent component of income e_i would not have any effect on aggregate consumption and, thus, on aggregate savings.[12]

Under the same assumptions, heterogeneous σ_g imply non-homothetic preferences. Expenditure shares and permanent income elasticities now depend on income. In particular, there is an inverse relationship between σ_g and a good's permanent income elasticity: relatively low-σ_g goods have a higher income elasticity. Thus, a profile of σ_g that decreases with age means high-income individuals display a relatively high expenditure share on the good "consumption at old age." Consequently, these individuals will display a higher savings rate over their working life, relative to low permanent income individuals. As a result, in this world, changes in inequality due to the permanent income component will affect aggregate savings and thus will affect the natural rate.

In our model, heterogeneity in σ_g is not the only source of nonhomotheticities. In fact, luxury bequests (De Nardi, 2004) as well as the bequest risk-aversion parameter also

[12] See Straub (2019) for a comprehensive discussion. Proposition 3 of this paper states that, under the conditions proposed, the distribution of permanent income is irrelevant for aggregate consumption and savings in the economy.

introduce nonhomotheticities due to preferences.[13] To those, we add addilog preferences, which allow us to calibrate the profile of σ_g to ensure that the model-implied elasticity of consumption with respect to permanent income, denoted ϕ_{PI}, matches that of the data.

The path of $\{\sigma_g\}$ is determined by two parameters: first, $\bar{\sigma}$, which represents the median value of σ_g; second, slope parameter σ_{slope}, which determines the decay of σ_g. Following Straub (2019), we assume the following:

$$\frac{\sigma_{g+1}}{\sigma_g} = \begin{cases} \sigma_{slope}, & \text{if } g \leq \bar{g} \\ 1, & \text{else} \end{cases}$$

That implies exponential decay until retirement age $\bar{g}$ and constant σ_g thereafter. Appendix Figure C15 plots the profile of σ_g under the baseline calibration.

Finally, the parameter o regulates the speed at which expenditure shares and income elasticities converge from their values for low-income individuals to their values for high-income individuals. Following Straub (2019), o is set to 30% of average income in the 2015 steady state, which corresponds to around \$20,000 in 2015 dollars. For a balanced growth path to exist we assume $o_t = A_t o_0$; that is, the o term grows with productivity.

Bequests. Following De Nardi (2004), the utility function for bequests is given by the following:

$$v_g(a) = b_0 \frac{1}{1 - b_{1,g}} \left[k_b + (1 - \tau_{beq}) \frac{a}{o} \right]^{1 - b_{1,g}}$$

Here b_0 governs the overall desire to leave a bequest, b_1 determines the curvature of the bequest utility function, and k_b introduces a luxury-bequest motive, which is another source of non-homotheticity. With these preferences, richer households have on average a stronger desire to leave bequests. Once again, o is time dependent and ensures that a balanced growth path exists in the case of $k_b > 0$. We set $b_1 = 2$ throughout our experiments, a standard value in the literature.

We assume that bequests received depend on the permanent level of productivity e_i and on age g; thus $beq \equiv beq(e, g)$. We calibrate the incidence function $beq(e, g)$ using data on inheritances received from the Survey of Consumer Finances. We provide further details in Appendix C.5.

[13] Even if preferences were homothetic, other features of our economy such as the redistributive nature of the transfer and tax system would also nullify the neutrality result in Straub (2019).

4.2 Demographics

In all our analyses, we feed our model the values for population growth rate n_t, survival probabilities $p_{g,t}$, and age-specific shares $\phi_{g,t}$ obtained from the 2022 Revision of the United Nations World Population Prospects (United Nations, 2022*b*). This data set provides estimates of demographic indicators at annual frequency from 1950 to 2021, and projections from 2022 to 2100.

4.3 Labor Productivity Process

We discretize the distribution of permanent productivity draws using two states, denoting them e^L for the low state and e^H for the high state. We choose an ergodic distribution of 90% of households with e^L and the remaining 10% with e^H. Note that the 10% of households with the e^H draw do not necessarily constitute the households with the top 10% labor income, since we also include an age-productivity profile and an idiosyncratic shock in the model. Since the population-weighted sum of individual productivity is a normalization, only the value of e^H has to be set, and the value of e_i will follow.[14] We add e^H to the numerical calibration routine. It is a crucial parameter to match the top 10% pre-tax earnings shares, which we take from updated data from Piketty and Saez (2003).

Similarly, we also include the parameter ζ^*, referring to the labor productivity of the star-earner state, in the numerical calibration routine. This parameter is identified by the top 1% share of earnings in the economy. In practice, given the life-cycle earnings profile (explained below), the joint calibration of e^H and ζ^* can be obtained in a separate routine, without requiring the computation of the model equilibrium.

To calibrate the profiles $h_z^g(e_i)$, $e_i \in \{e^L, e^H\}$, we use data from the Luxembourg Income Study Database (LIS).[15] From this source, we obtain the *average* life-cycle profile of earnings. We then use information from Guvenen et al. (2021) to account for the fact that earnings growth is much steeper for high-income individuals relative to low-income individuals. We relegate the details to Appendix C.2. For the persistent shock, we follow Guvenen et al. (2019) and set $\rho = 0.9$ and $\sigma_\epsilon^2 = 0.04$. We discretize using the Tauchen (1986) method and four productivity states ζ in the baseline experiment. For individuals with $e_i = e^H$, we supplement the process with $\zeta = \zeta^*$, as previously explained.

To compute counterfactual changes in income inequality over time, we proceed in two

[14] The normalization ensures that output per effective unit of labor ($\tilde{Y}$) equals 1 in the 2015 steady state.

[15] LIS is an income database of harmonized microdata collected from about 50 countries. The underlying source for data for the United States in 2015 is the Current Population Survey.

steps. First, we compute changes in the variance of the idiosyncratic shock σ_ϵ^2, using relative changes in residual wage dispersion obtained from the PSID. We relegate the details to Appendix A.3. We then recalibrate the parameters e^H and ζ^* to match the top 10% and top 1% earnings shares from Piketty and Saez (2003), ensuring that the average productivity at each age is kept constant by adjusting the life-cycle profiles $h_z^g(e)$ accordingly.

4.4 Production

We set the value of the intermediation wedge to $\nu = 0.02$, or 2%, according to data on the average spread between US 10-year government bond and investment-grade corporate bond yields. Because we treat the wedge as a deadweight loss, in the model it is indistinguishable from depreciation (as it is also not taxable). [16]

The parameters α and μ are set to match the average labor share from 2010 to 2019, which we find to be a labor share of 59.4% (Feenstra, Inklaar and Timmer, 2015). The capital-to-output ratio is set to $\frac{K}{Y} = 3.7$ in the 2015 steady state. This is between the decade average of, respectively, the capital-to-output ratio from National Income and Product Accounts (NIPA) tables and wealth-income ratios from the World Inequality Database (WID), where the latter is adjusted for public debt and net foreign assets. We allow the labor share to change between steady states and on the transition path. To do this, we leave α unchanged but adjust μ. Thus, we map the labor share indirectly, via the profit share. A falling labor share and a corresponding increase in the profit share are documented in Barkai (2020) and Eggertsson, Robbins and Wold (2018).

4.5 Medical Expenses

Our assumptions regarding medical expenses are a simplification of Kopecky and Koreshkova (2014). At age g for a household with productivity draw z and health status h, expenses $m_g(z,h)$ are as follows:

$$\log m_g(z,h) = h_m^{z,g} + h \tag{12}$$

[16] One of our calibration targets consists of the return on savings. Thus, an increases in ν would be exactly offset by reductions in δ. An alternative would be to rebate intermediation proceeds to the mutual fund and, ultimately, to households. However, the calibration procedure would again adjust δ to keep the return on savings constant. The adjustment would then exactly offset the increased returns due to the additional intermediation proceeds, so long as the tax rate on intermediation returns equals τ^k, rendering the assumption of whether intermediation represents a deadweight loss irrelevant.

Here $h_m^{z,g}$ is a deterministic component and h is a medical expense shock. $h_m^{z,g}$ generates an age profile for medical expenses. h is two-state Markov, and we can interpret the two states as being in good or bad health. The term h implies idiosyncratic risk for retirees. In the steady-state-comparison and transition path experiments, we keep medical-expense shock m unchanged but adjust the relative price of medical goods p_m according to the data.

4.6 Government

We choose labor income tax $T_y(\cdot)$ as in Heathcote, Storesletten and Violante (2017), implying post-tax labor income given by equation (13) below:

$$wz - T_y(wz) = A\lambda_0 \left(\frac{wz}{A}\right)^{1-\lambda_1} \tag{13}$$

Parameter λ_1 governs the progressivity of the income tax system, and λ_0 determines the overall level of the income tax. This functional form has been suggested as a parsimonious yet accurate characterization of the US income tax system by Heathcote, Storesletten and Violante (2017). The tax function is indexed to productivity growth, ensuring a balanced growth path. Figure C18 in Appendix C.4 shows post-tax labor income as a function of pre-tax labor income.

The social security system is a piecewise linear schedule following the Old Age and Survivor Insurance component of the US Social Security program (Huggett and Ventura, 2000; De Nardi and Yang, 2014):

$$\begin{aligned} T^{SS}(x; W_t) = \phi_{SS}\Big\{ & 0.9 \min(x, 0.2W_t) + 0.32 \max[0, \min(x, 1.24W_t) - 0.2W_t] + \\ & + 0.15 \max[0, \min(x, 2.47W_t) - 1.24W_t] \Big\} \end{aligned} \tag{14}$$

Here x is the transfer base and W_t is average labor income in the economy at time t, and ϕ_{ss} is calibrated to ensure the total expenditures in social security (relative to output) match the data in our baseline calibration.[17] For tractability, we assume the transfer base depends on the realization of z immediately prior to retirement ($g = \bar{g} - 1$), so $x_{i,g,t} = w_t z_{i,g,t}$.[18]

[17] Note that ϕ_{ss} does not have a time subscript. In our model, changes in total expenditures on social security over time are the result of a change in the drivers, in particular demographic variables, and the resulting general equilibrium effects.

[18] In the US Social Security program, pre-tax benefits $T^{SS}(\cdot)$ are calculated using a measure of average lifetime income as the benefit base.

There is no income risk in retirement: agents keep the realization of z from the last period before retirement until they die. Note that making the benefit formula dependent on average labor income W_t implies that payments are indexed to productivity growth, again ensuring the existence of a balanced growth path. Finally, we assume benefits are taxed according to the same system as labor income.

The government's budget constraint (10) puts restrictions on the joint path of debt B_t, total spending Ξ_t, and total revenue Λ_t. Spending and revenues are partially determined by endogenous objects in the model, such as social security spending and tax revenue from various sources. Throughout, we consider government debt and expenses as drivers, and thus we either feed their past values or make assumptions, or use forecasts from the Congressional Budget Office, about their future evolution. The government budget constraint is then cleared by adjusting $\lambda_{0,t}$. We leave a comprehensive analysis of how the interaction between social security policy and government debt with different tax instruments affects r^* to future research, but we provide estimates of the individual effect of changes in each instrument in Section 7.

4.7 Numerical Calibration Procedure

Internally set parameters are jointly determined to match data moments. These include (i) the discount factor β; (ii) the permanent productivity gap between low and high type $e^H - e^L$; (iii) the star-earner parameter ζ^*; (vi) the slope of the profile of age-dependent risk-aversion parameters σ_{slope}; (v) the relative price of medical expenses p_m; (vi) a shifter, k_b, that determines to what extent the bequest motive is stronger as income increases; (vii) the overall strength of the bequest motive, b_0; (viii) the median risk-aversion coefficient $\bar{\sigma}$; (ix) the social security shifter ϕ_{ss}; and (x) the depreciation rate.

Even though all 10 parameters are jointly selected, each of them is more strongly associated with one particular moment. We now briefly describe the chosen data moments and their values, along with the rationale behind the choices.

We target $r^* = 0.53\%$ since it corresponds to the average of eight estimates of the level of the natural rate in 2015.[19] The parameter most closely associated with r^* is the discount factor β. The parameters related to top earnings shares are chosen to match the top 10% and top 1% earnings shares, respectively equal to 34.7% and 11.0%, following updated data in Piketty and Saez (2003).

[19] Seven of the estimates are taken from Bauer and Rudebusch (2020). In addition, we construct a trend by smoothing a short-term real rate using a moving average over 30 years. For comparison, the yield on a 10-year Treasury Inflation-Protected Security averaged 0.45% in 2015.

To calibrate σ^{slope}, we follow Straub (2019) and match the elasticity of consumption out of permanent income ϕ_{PI} to 0.7. To compute it, we use a Monte Carlo procedure to simulate the model-generated data. We then apply the same estimation approach to estimate ϕ_{PI} with real-world data. The procedure involves a regression of consumption on income and age controls, where income is instrumented by income leads to eliminate bias, especially from idiosyncratic income shocks. This procedure is done within the numerical calibration routine. Further details on the estimation of ϕ_{PI} can be found in Appendix C.6.

We set the relative price of OOP medical goods and services, p_m, to match the total of OOP medical expenditures, which amount to 1.5% of GDP (Kopecky and Koreshkova (2014)). With regard to the bequest parameter, we choose to match the 30th percentile of the distribution of bequests (19.3% of average bequests, from Hurd and Smith's (2002) Table 1) and a total amount of bequests (relative to GDP) of 7.5%. The latter is within the range of existing estimates (see, e.g., Hendricks (2001) or Alvaredo, Garbinti and Piketty (2017)) and helps us replicate the profile of wealth across age, shown in Figure 4.

The median risk-aversion coefficient, $\bar{\sigma}$, is set so that the model-implied *semi-elasticity of savings with respect to the return on savings* is in line with the data (Auclert et al. (2021)). We select the midpoint of the boundary values reported, a value of 18.1, and conduct robustness exercises in Section 8. This elasticity determine how strongly shifts in the demand for savings affect rates of return and thus r^*.

We set ϕ_{SS} to ensure that the total social security expenditures amount to 5.8% of GDP, as in Congressional Budget Office (2020). Finally, we set the depreciation rate δ to match the estimates for returns to savings in the United States from Auclert et al. (2021), equaling 3.9%.

The complete set of parameters is displayed in Tables 1 and 2 below. Table 1 shows that we are able to match all empirical targets well. Notably, the elasticity of consumption out of permanent income, denoted as ϕ_{PI}, is matched precisely. In the next section we show that non-homothetic utility is important to achieve this outcome. A noticeable discrepancy is only observable for the distributional target for bequests, where we target 19% and the model moment is about 22%. We limit parameter k_b to a maximum value of 50, as we find increasing the parameter further does not materially improve the fit.

To conclude this section, Figure 4 shows data together with model outcomes for life-cycle wealth profiles. This includes (normalized) average wealth holdings per age group (upper section) and the proportion of total wealth attributed to each age group (lower section). The model replicates average wealth profiles very well. Specifically, it matches the empirical

Table 1: Calibration Results: Internally Set Parameters

Panel A: Numerical Procedure					
Description	Target - Data	Target - Model	Parameter	Value	Source
r^*	0.53%	0.53%	β	0.996	Bauer and Rudebusch (2020) and own calculations, see text
Top 1% Earn. Share	11.0%	11.0%	ζ^*	$\bar{\zeta}$+0.32	Piketty and Saez (2003) updated
Top 10% Earn. Share	34.7%	34.7%	e^H	3.8	Piketty and Saez (2003) updated
Beq. Distribution	0.19	0.219	k_b	50	Hurd and Smith (2002)
Bequest-to-GDP	0.075	0.075	b_0	69.6	Hendricks (2001), Alvaredo, Garbinti and Piketty (2017)
OOP-to-GDP	0.015	0.015	p_m	0.009	Kopecky and Koreshkova (2014)
ϕ_{PI}	0.700	0.700	σ^{slope}	0.993	Straub (2019)
SSY	0.058	0.058	ϕ	1.43	Congressional Budget Office (2020)
Elast. Sav. to Returns	18.125	18.125	$\bar{\sigma}$	3.134	Auclert et al. (2021)
Return to Savings	0.039	0.039	δ	0.050	Auclert et al. (2021)
Panel B: Implied by Equilibrium Conditions					
Description	Target - Data	Target - Model	Parameter	Value	Source
$\frac{K}{Y}$	3.700	3.700	α	0.302	BEA, WID
Labor Share	0.594	0.594	μ	1.175	Feenstra, Inklaar and Timmer (2015)

Note: OOP refers to out-of-pocket expenditures, and ϕ_{PI} corresponds to the elasticity of consumption with respect to changes in permanent income.

fact that households only slowly deplete their wealth in old age. The model slightly overstates wealth holdings of households aged 35 to 50 years old and underestimates wealth holdings around the age of 60.

In Appendix D.1, we explain why replicating the wealth profile is important to accurately capture the total impact of demographic change. In particular, our model can capture with precision the change in the aggregate wealth-to-income ratio induced solely by changes in the age composition of the population—the so-called compositional effects of demographic shifts—for all the years in our analysis. Also in Appendix D.1, following Auclert et al. (2021), we formalize this concept and show the quality of our model fit.

5 Steady-State Analysis

The first set of results looks at the change in the natural rate between steady states. We compare a 2015 steady state to successive steady states in the previous decades, going back until 1965. We are interested in the total change in the natural rate and the contribution coming from individual drivers. We also show a more detailed decomposition for demographics and inequality.

Our approach is the following. To calculate the total change to 2005, we change the values of all underlying drivers to their respective levels in 2005 and solve the model. The only

Table 2: Calibration Results: Externally Set Parameters

Parameter	Sign	Value	Source
Demographics:			
No. generations	G^d	74	
Retirement age	$\bar{g}+25$	66	
Survival probabilities	$p_{g,t}$	(vector)	United Nations (2022*b*)
Population masses	$\phi_{g,t}$	(vector)	United Nations (2022*b*)
Population growth rate	n	0.77%	United Nations (2022*b*)
Production:			
Productivity growth parameter	γ	0.49%	Fernald (2014) (updated)
Fiscal:			
Public debt/Y	B	72.4	Congressional Budget Office (2019)
Exogenous gov. spending	G	14.9	Congressional Budget Office (2019)
Consumption tax	τ_c	0.05	Kitao (2014)
Capital income tax	τ_k	0.3	Kitao (2014)
Profit tax	τ_{corp}	0.25	OECD and BEA (see note)
Bequest tax	τ_{beq}	0.1	De Nardi (2004)
Curvature income tax function	λ_1	0.181	Heathcote, Storesletten and Violante (2017)
International:			
Net foreign asset position	NFA	-37.4	BEA (2020)
Earnings Process:			
Persistence earnings shock	ρ_w	0.9	Guvenen et al. (2019)
Var. earnings shock	σ^2_ϵ	0.04	Guvenen et al. (2019)
Health:			
Deterministic med. expenditure	$h_m^{e,g}$	vector	Kopecky and Koreshkova (2014)
Bad health state	h_{bad}	log 13.48	Kopecky and Koreshkova (2014)
Transition matrix health	Λ_{hh}		Kopecky and Koreshkova (2014)
Minimum consumption floor	$\underline{c}$	16.5% W	Kopecky and Koreshkova (2014)
Utility:			
Scaling parameter consumption utility	o	$0.3\tilde{Y}$	Straub (2019)

Note: For the profit tax, we take the midpoint between the statutory rate (35%, OECD.Stat Table II.1) and an effective tax rate of 16%, based on Bureau of Economic Analysis estimates (Tax Receipts on Corporate Income/(Corporate Profits after Tax (without IVA and CCAdj)+Federal Government: Tax Receipts on Corporate Income), code A055RC).

exogenous variables to be changed going from the 2015 steady state to the 2005 steady state are the ones pertaining to the drivers, summarized in Appendix Table D7. All other parameters are left unchanged. This gives us a value for r^* in 2005, and the difference from the 2015 steady state gives us the total change over the period. To isolate the effect of each individual driver, we change the exogenous variables connected to each driver while leaving all other exogenous variables at their 2015 steady-state value. The sum of the effects when changing drivers one at a time does not have to be the same as the sum when changing all drivers together. We denote the difference between the two as a residual. We repeat the same calculations moving from 2005 to 1995 and prior decades successively to 1965. The change in r^* between 2015 and a specific decade is the sum of the individual decade effects. Figure 5 and Appendix Table D8 display the results.[20]

[20]In Appendix D.4 we consider a different way of performing this decomposition, in which we keep all

Figure 4: Wealth over the Life Cycle

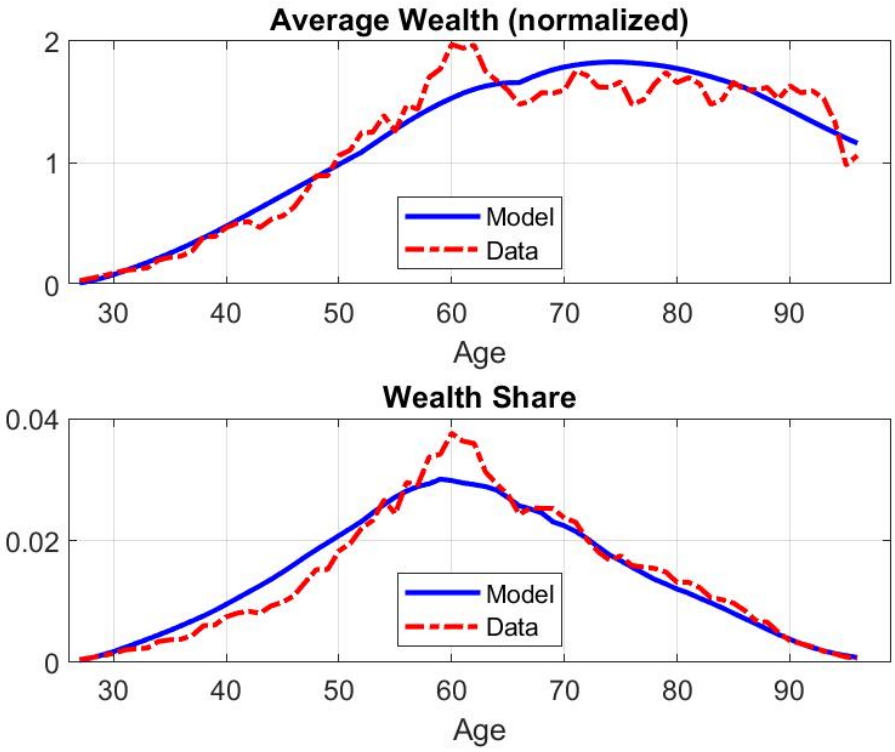

Notes: Sources: Survey of Consumer Finances 2016, authors' own calculations. The top panel shows average wealth (defined as net worth) for individuals at each age, in the model and in the data. The values are normalized relative to the average wealth across all ages (weighted by the share of population at each age). The bottom panel displays the share of wealth held by individuals at each age, in the model and in the data.

Our model accounts for a decline in r^* of 4.26 pp when moving from the 1965 steady state to the 2015 steady state. There was a relatively large decline between 1965 and 1985 (2.64 pp), followed by a slight increase to 1995. This short intermission is due to not only positive productivity dynamics during this time but also a slight positive contribution from demographic forces and an increasing labor share. Finally, there was about a 1.9 pp decline in r^* between 1995 and 2015.

There are four large contributing drivers explaining the decline since 1965: productivity growth (-1.44 pp), inequality (-0.96 pp), demographic dynamics (-0.72 pp), and the labor share (-0.6 pp). The rise in public debt was an important counteracting force, in particular

drivers at their 2015 levels and then change only one driver to its assumed decade value. We repeat this process decade by decade, driver by driver, but we always keep all other drivers at their 2015 level. The main results, in terms of drivers' relative contributions, are unchanged, but the residual term is large and positive, suggesting that there are reinforcing interactions among the drivers.

Figure 5: Decomposition of the Change in r^*

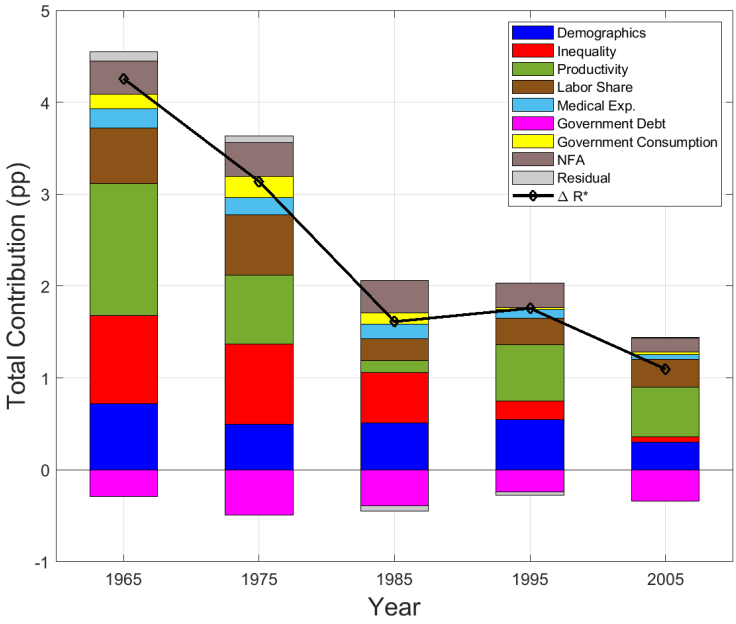

Notes: The figure shows the decomposition of change, in percentage points, into various drivers for decades since 1965. The black diamonds show total change in r^* in respective years relative to 2015. The stacked bars show the contribution to the change relative to 2015 from various drivers.

since 1975, about the time when US debt started its ascent from a postwar low. This contributed to an increase in r^* of 0.49 pp. We find that inequality is the most important driver when comparing a 1975 to a 2015 steady state, explaining 27.7% of the total decline. This is considerably larger than the contribution coming from demographic forces (15.9% of the total). The contribution from the change in the net foreign asset position is still significant in size, at about -0.36 pp since 1985, and a similar magnitude when starting in earlier decades. The contributions from medical expenditures (-0.21 pp) and government consumption (-0.16 pp) are limited. The residual, capturing interactions between the drivers, is only 0.15%.

In terms of the time pattern, we see that the contribution from declines in productivity growth came in two phases: from 1965 to 1985 (1.38 pp), and from 1995 to 2015 (1.08 pp). The intermediate increase in total factor productivity growth, and consequently a positive effect on r^*, could be attributed to IT adoption. Inequality's contribution was small from 1965 to 1975 (-0.09 pp) but substantial between 1975 and 1995 (-0.67 pp). The impact of demographic change has been most pronounced in recent decades: while the contribution was modest from 1965 to 1995 (-0.17 pp), it became one of the most important drivers

after 1995 (-0.55 pp). Notably, the large generation born after World War II (the baby boomers) entered the labor force during the 1970s, leading to strong labor force growth and low dependency ratios. Finally, the quantitative impact of the decline in the labor share was concentrated between 1975 and 1985 as well as within the last decade.

In sum, our analysis shows a substantial decline in the natural rate over the last five decades. This decline is not explained by one single underlying trend but by a confluence of factors, operating in different periods. The next subsections give more details on the underlying mechanism of two important drivers: demographic change and the rise in inequality.

Figure 6: Decomposition of Changes in r^* - Demographics

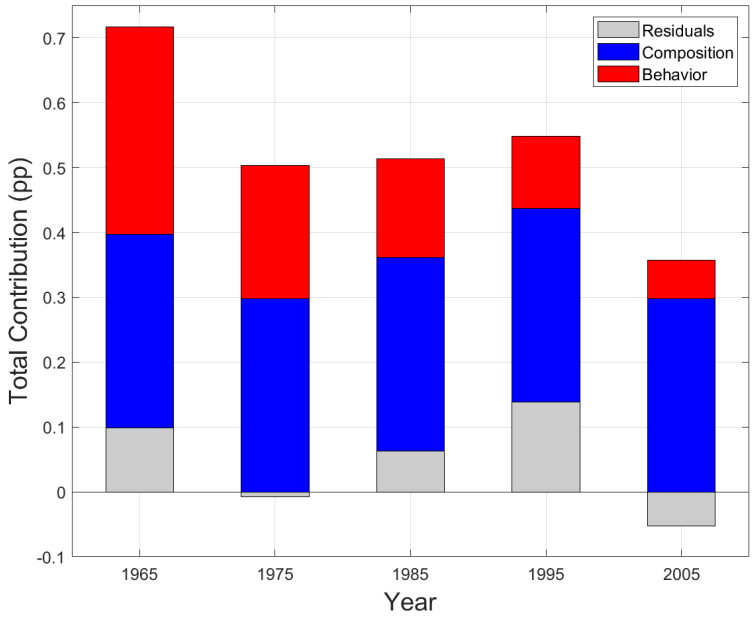

Notes: The stacked bars show the contribution, in percentage points, of subcomponents of demographic change to the change in r^* relative to 2015.

Decomposition of Demographics. Figure 6 shows a decomposition of the contribution of demographic change into subcomponents: composition, behavior, and a residual term. The composition term is calculated by holding asset profiles as well as labor supply profiles by age constant at their 2015 level, but changing the population distribution (or composition of the population) over time. The benefit of this approach is that the compositional effect can easily be calculated with data on the evolution of the population distribution alone. Since our model matches empirical asset profiles well (Figure 4), and the evolution of the population distribution is, by construction, identical to the data, the model matches the

compositional effect well.

Concretely, to isolate the compositional effect on aggregate savings, we introduce an additional term $\Delta^{comp}_{\tau,2015}$ in the savings market clearing condition and recompute the equilibrium given this term:

$$B_{t+1} + K_{t+1} = \sum_{g} \sum_{z,h,a} g^a_t(z,h,a,g)\lambda_t(z,h,a,g) - NFA_t + \Delta^{comp}_{\tau,2015} \cdot (K_{2015} + B_{2015})$$

The term $\Delta^{comp}_{\tau,2015}$, with $\tau \in \{1965, 1975, 1985, 1995, 2005\}$, represents how much the aggregate savings-to-income ratio (as a share of its baseline value) would change if savings behavior (i.e., policy functions $g^a(\cdot)$) and the age-earnings profile remained constant but the population composition (i.e., $\{\phi_{g,\tau}\}_{g=1}^{\bar{g}}$) corresponded to that in year τ. That is, it captures the impact of changes in the composition of the population, holding savings behavior fixed. For a detailed explanation of the compositional effect and how $\Delta^{comp}_{\tau,2015}$ is computed, refer to Appendix D.1 (equation (25)).

The behavior term is calculated by holding the population distribution ($\{\phi_{g,t}\}_{g=1}^{\bar{g}}$ constant at its baseline (2015) level but changing the survival probabilities and allowing individuals' choices to adjust. This counterfactual thus alters households' savings behavior: as people expect to live longer, they save more to sustain the extended retirement period. Finally, the residual term captures the interactions, the effect of population growth n_t (conditional on population shares), and the changes in tax levels required to fund social security.[21]

We find that both terms are of equal importance when looking at the demographic transition starting in 1965, with a contribution of about 0.3 pp coming from each. However, the contribution from the behavior term is gradually realized over time, mostly materializing by 2005. In contrast, almost all the contribution from the composition term comes in the last decade, from 2005 to 2015. Life expectancy increased steadily over the last decades in the US and remained fairly constant over the decade from 2005 to 2015, in line with the behavior term. Regarding the large contribution of the composition term since 2005, this reflects the large baby-boom cohort reaching the age of peak savings over this period. The contribution of the residual term is limited but volatile.

Decomposition of Inequality. Figure 7 shows a decomposition of the contribution of the inequality driver into subcomponents: transitory, permanent, and a residual term. The transitory term shows the contribution from changing idiosyncratic income risk. The

[21] We find the impact of population growth rate n_t by itself to be rather small and thus relegate it to Appendix D.5.

permanent term shows the impact from changing the two parameters focused on matching the top 10% and top 1% earnings shares in the data, e^H and ζ^*, and correspondingly adjusting e^L to keep average productivity constant.

Figure 7: Decomposition of Changes in r^* - Inequality

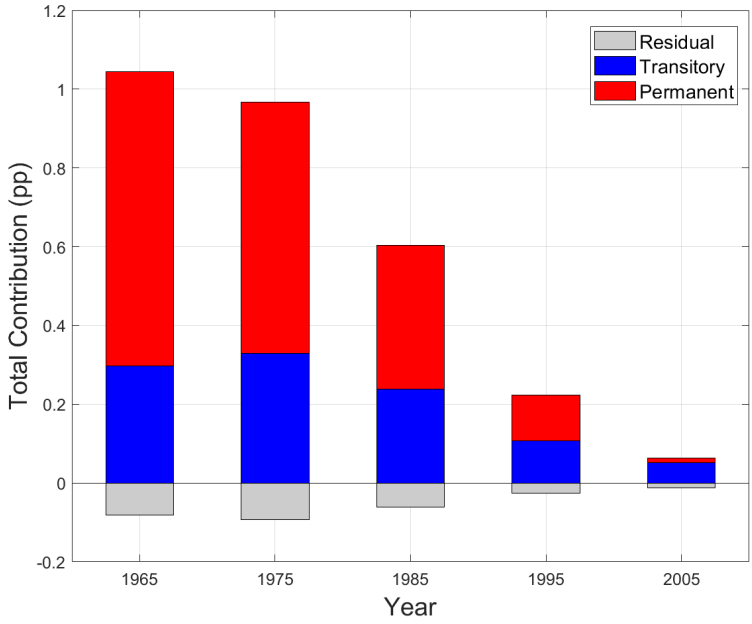

Notes: The figure shows the decomposition of change in r^*, in percentage points, into the subcomponents of inequality for decades since 1965. The stacked bars show the contribution to the change relative to 2015.

We find that the permanent term is quantitatively more important. It explains about 0.74 pp, or 77.7% of the total inequality effect, since 1965. The transitory term only accounts for about 0.3 pp since 1965. Most of the effect of the permanent term is realized between 1975 and 1995, in line with the rise of the top earnings shares we observe in the data. Finally, the residual term, capturing interactions, is small.

This result suggests that the change in permanent inequality is important to fully account for the impact of a rise in inequality on the natural rate. To capture this effect in its entirety, including non-homothetic preferences turns out to be crucial. We now turn our discussion to this issue.

Homothetic Utility Results. Table 3 reproduces our baseline findings together with the change in r^* from a simulation of our model with *homothetic* utility. Concretely, we make the following three changes to the baseline model: (1) homothetic consumption

preferences, such that $\sigma_g = \bar{\sigma} \; \forall \; g$, (2) homothetic bequest utility, such that $k_b = 0$,[22] and (3) an equal risk-aversion parameter for consumption and bequest, $b_1 = \bar{\sigma}$. We recalibrate the model to match the targets of Table 1, excluding the target for the elasticity of consumption out of permanent income and the distributional target for bequests. Table 3 shows the change in r^* between 1965 and 2015. Appendix D.6 shows the full set of results in Figure D24 and Table D9.

Table 3: Comparison of Results: Benchmark versus Homothetic Utility

Driver	Benchmark Calibration	Homothetic Utility
Demographics	0.72	0.36
Inequality	0.96	0.23
Productivity	1.44	1.22
Labor Share	0.60	0.62
Medical Exp.	0.21	0.08
Government Debt	-0.29	-0.32
Government Consumption	0.16	0.10
NFA	0.36	0.34
Residual	0.10	0.00
Total	4.26	2.63

Notes: The table displays the decomposition of changes in r^* from 1965 to 2015 in the benchmark economy and the homothetic economy. Benchmark calibration refers to the economy calibrated as in Tables 1 and 2. The homothetic economy repeats a similar calibration procedure but sets $\sigma_g = \bar{\sigma} \; \forall \; g$, $k_b = 0$, and $b_1 = \bar{\sigma}$.

The total decline in r^* is reduced from 4.26 pp to 2.63 pp. The largest change among individual drivers is for inequality, whose contribution drops from 0.96 pp to 0.23 pp. This does not come as a surprise, as we introduced non-homothetic utility to match the empirical target for the elasticity of consumption out of permanent income, ϕ_{PI}, which is a key determinant of the impact of permanent income inequality on savings. Without this feature, the elasticity of consumption with respect to permanent income rises from $\phi_{PI} = 0.7$ to $\phi_{PI} = 0.93$. Delving into the components of inequality, we see that the contribution coming from permanent income collapses from 0.74 pp to 0.11 pp, a decline of 85%, while the contribution from idiosyncratic risk falls from 0.3 pp to 0.12 pp, a decline of 60% (see Figure D26 in Appendix D.6). The contribution of demographics also falls in the case of homothetic utility, from 0.72 pp to 0.36 pp, which is almost entirely due to the vanishing of the behavioral effect on r^* (see Figure D25 in Appendix D.6). Other drivers remain fairly stable.

[22] In practice we set $k_b = 0.2$ for computational reasons.

These results show that the impact of permanent-productivity changes on r^* is sensitive to the homotheticity assumptions in the model. Working in a model with homothetic utility, one would conclude that permanent-productivity differences and consequently inequality cannot account for much of the decline in r^*. But such a model fails to capture differences in savings behavior across permanent income types, delivering a counterfactually high ϕ_{PI}. For this reason, we believe the results from the baseline version of the model deliver a more accurate representation of the impacts of inequality on r^*.

6 The Future of r^*

What can we expect of the future evolution of the natural rate in the United States? In this section we address this question. Specifically, we ask: what does the likely evolution of the most important secular drivers mean for the future of the natural rate? Again, we emphasize that we abstract from the impact of business cycles. In practice, as explained below, we always consider smoothed versions (moving averages) of our data series.

To better assess the uncertainty in the future path of r^*, we undertake two different approaches. First, we conduct a *scenario analysis*, in which we consider a pessimistic and an optimistic future path of the most important drivers, namely productivity growth, inequality, demographics, and the labor share. We then conduct a *probabilistic experiment*, in which we consider the future path of these four drivers as random. In this case, we conduct a Monte Carlo simulation of our model, which allows us to provide confidence intervals.

In all experiments, we initialize our economy at a hypothetical 1950 steady state. At that point, agents are made aware of the full trajectory of parameter changes and display perfect foresight. Eventually, all shocks achieve their final value, but it can take several additional years for the economy to converge to the new steady state. We truncate our transitions so that convergence is achieved after 250 years, in 2200.

We can distinguish between two separate periods for changes in exogenous parameters (drivers): past and future values. For the former, we use data corresponding to the realization of the various drivers listed in Table D7. The challenging part concerns the future values of the drivers, for most of which—with the exception of demographics and public debt—there are no prompt or reliable forecasts available. For these, we make different assumptions, which we detail shortly.

For the demographic projections, we use data from the United Nations World Population

Prospects (United Nations, 2022*b*). In addition, we assume that public debt increases linearly beginning in 2019, culminating in a new steady-state value of 200% of GDP in 2080, in line with estimates from the Congressional Budget Office (2022), as shown in Figure 3, top-right panel. Throughout our simulations, we assume that the labor income tax level parameter λ_0 adjusts to ensure the government's budget constraint is met at all times.

Scenario Analysis. We consider a central, a "low" (pessimistic), and a "high" (optimistic) scenario. The path of nominal debt is the same across these scenarios. We later provide a additional simulations considering alternative paths of the debt-to-GDP ratio.

In the central scenario, for demographics we consider the "central fertility" scenario from the UN estimates. This corresponds to the UN's median estimate for the evolution of fertility and mortality rates. Note that fertility and mortality rates determine the evolution of population shares $\{\phi_{g,t}\}_{g=1}^{G^d}$.[23] All other drivers converge from their 2022 values to their 2015 values within 40 years.

For the low and high scenarios, we consider different paths for the four most important drivers: demographics, inequality of the permanent type (the parameters mapping the top 10% and top 1% shares), productivity growth, and changes in the labor share. For demographics, the high and low scenarios correspond to the high- and low-fertility scenarios obtained by the United Nations (2022*b*). These correspond to the 5th and 95th percentiles of its model simulation. For the high case, we assume that productivity growth, the labor share, and the top 10% and top 1% earnings shares revert to their 1995 values. For the low scenario, we extrapolate the change in their values over the last 20 years for another 20 years into the future. In either scenario, these three drivers stabilize at their implied values two decades into the future, while demographic variables can still change up until 2100. Figure E27 in the appendix shows the path of all drivers across the different scenarios.

In all, the low scenario is meant to represent a pessimistic case, in which trends that tend to pull r^* down continue to operate. The opposite is true for the high scenario. At this point, we are silent as to whether those scenarios are likely to materialize. Yet this exercise can give us a good idea about (i) what to expect if the drivers continue their trend and (ii) what a full reversion of trends—while admittedly unlikely—would mean for future rates.

Figure 8 illustrates the three transition paths for the natural rate, while Table 4 displays the terminal values for the drivers together with their individual contribution for the terminal r^*. The plot shows that, following a substantial decline over prior decades, r^* stands at

[23] To back out the population shares using the implied fertility and mortality rates, the UN considers constant migration flows.

about 0.41% in 2020 according to our estimates. Under the baseline scenario, the natural rate continues to fall until 2029, reaching a low of 0.27%.

Figure 8: Future Paths of r^* - Central Simulation and Scenarios

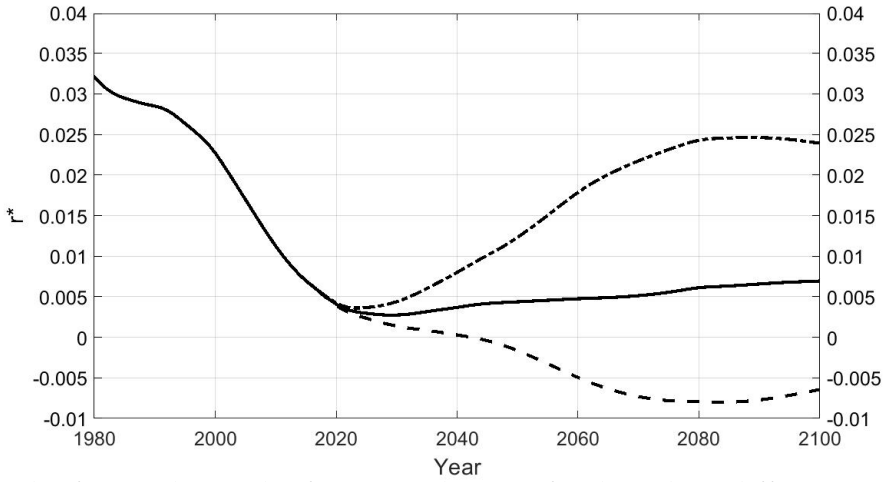

Notes: The figure shows the future evolution of r^* based on different scenarios for the evolution of the drivers. The solid, dashed, and dot-dashed lines represent respectively the central, low, and high scenarios. See text for further explanation.

Subsequently, a steady rise ensues, with values attaining 0.44% in 2050 and 0.69% by 2100. Although we foresee an uptrend in r^* going forward, the rate will only match its 2010 level by 2100, and it will stay significantly below the highs of decades prior to 2010.

The post-2030 surge is only due to demographic and debt dynamics, as the other drivers stabilize at their 2015 values. What is the individual contribution of each driver? Appendix Figure E32 provides the answer. In this figure, we repeat the scenario exercise but assume that public debt stabilizes instead at 100% of GDP by 2080. In this case, we observe a gradual, yet consistent, decline in r^*, reaching 0.16% in 2031 and -0.13% in 2100.

The crucial lesson is that demographics alone are expected to exert downward pressure on the natural rate of interest in the future. This result resembles the one in Auclert et al. (2021), with the key difference that we consider the US in isolation. That is, even for the US alone, where the demographic transition might be considered to be ahead of large developing countries (notably China and India), demographic forces tend to pull r^* down.

Recall from Section 5 that the impact of demographics can basically be decomposed into two forces: the compositional effect, representing the impact of changes in the composition of the population, and the behavioral effect, representing the fact that people change their savings behavior in response to changes in their expected longevity. In the future,

both impacts play a role. First, as shown by Figure E27, in the bottom-left panel, the UN predicts increasing longevity up until 2100, leading to an increase in savings supply and consequently to a negative impact on the natural rate, which we dubbed the "behavioral effect."

At the same time, the average age of the (adult) population increases from roughly 52 years in 2020 to approximately 59 in 2100. Some analysts have proposed that, by itself, this mechanism will lead to a rebound in r^* in the near future as the baby-boomer generation ages and depletes its savings. In fact, we find the opposite: as shown in Appendix Figure E30, the compositional effect is largely responsible for the impact of demographics on r^* into the future. Crucial for this result is the fact that the wealth profile is relatively flat at later stages of life, as shown in Figure 4. In other words, individuals do not really deplete their savings in later stages of life. Our model matches this empirical fact, which in turn is expected to drive the total impact of demographic change in the future.

Turning our attention to the scenarios, we see that under the high scenario, the natural rate promptly reverses its declining trend, stabilizing just below 2.5% by 2080. Conversely, in the low scenario, the decline persists for several more decades, culminating in a permanently negative rate of below -0.5% after 2070. Using this approach, our 2100 forecast interval spans 3 pp, ranging from about -0.5% to 2.5%, a fairly wide range. Notably, though, in the high scenario the natural rate only approximates values observed around 2000 and remains beneath peaks from decades prior. This result is the first hint suggesting that, given the evolution of the drivers considered, a strong reversion toward past values is unlikely. We return to this issue in the probabilistic analysis below.

To conclude, Table 4 shows individual drivers' contributions to r^* in the terminal steady state. The largest contribution to the alternative scenarios (both high and low) comes from a change in productivity growth, while the smallest comes from inequality dynamics. These results are naturally consistent with our assumptions made regarding the evolution of these trends in each scenario—which are based on their movements over the past two decades—together with our results from Section 5: in the 20 years leading to 2015, there was a strong decline in productivity growth, with a consequent strong impact on r^*, while the impact of inequality over the same period was relatively muted.

Table 4: Drivers and r^* in Scenarios (Terminal Steady State)

-	Low Value	r^* - Low	High Value	r^* - High
Population Growth (%)	-0.64	0.56	0.65	0.98
Labor Share	57.0	0.45	62.0	1.13
TFP	0.01	0.24	1.17	1.52
Top 10%, Top 1%	36.5, 13.2	0.63	33.3, 9.0	0.90
All	-	-0.45	-	2.38

Notes: The table shows individual contributions to the values of r^* by the drivers considered in the scenario analysis, together with their combined contribution in each scenario in the last row. The terminal natural rate in the central scenario is 0.77%, while the benchmark (2015) r^* is 0.53%.

Probabilistic Analysis. The good and bad scenarios arguably depict extreme cases that are, at best, unlikely to materialize. To address the concern, in an attempt to better quantify the range of (im)plausible values r^* can attain in the future—again, given the likely evolution of the drivers considered—we conduct a probabilistic analysis based on a Monte Carlo simulation of our model.

We proceed as follows: First, as explained below, we assume that the values the drivers will follow in the future are random. Second, we draw paths of each driver from its respective probability distribution. For each given path, we then use our model and simulate the transition path of the natural rate, similarly to the previous section. We maintain the perfect-foresight assumption throughout. In practice, our results are based on 400 joint draws for drivers pertaining to demographics, permanent-income inequality, productivity growth, and the labor share. We now proceed to explain our assumptions regarding the probabilistic evolution of each driver.

We use probabilistic projections of the United Nations Population Prospect (United Nations (2022*b*)) for demographic variables.[24] The UN uses probabilistic methods to project future fertility and mortality levels until 2100. This gives us, for each draw, a full path for population distributions, population growth rates, and mortality rates.

For the other main drivers, in the absence of external probabilistic forecasts, we make the following assumption. First, we assumed that the parameters γ, μ, e^h, and ζ^* will converge (linearly) to their terminal value in roughly 40 years.[25] We then draw these

[24]We thank the World Population Prospect team for providing us with draws from their probabilistic model. The underlying methodology is explained in United Nations (2022*a*).

[25]See Appendix E.1 for details. Throughout, we adjust e^L to keep the average z at each age constant.

terminal values from three independent normal distributions, assuming that the latter two are perfectly correlated. In other words, the evolution of productivity, the labor share, and inequality of the permanent type are independent. The normal distributions are centered on each driver's 2015 value, while the standard deviations correspond to half the change between 1995 to 2015. Appendix Table E10 details the parameters of the resulting normal distributions.

We keep debt on a path to 200% of GDP by 2080 and constant thereafter for all draws, as under the baseline. All other drivers stay unchanged at their 2015 level. An essential caveat is the assumption of driver independence. If, in reality, drivers exhibit systematic correlations, our analysis may under- or overstate the variance in future r^* values.

Given the uncertainty in our assumptions concerning the other drivers, we first present results for draws for demographic variables alone.

Figure 9: Probabilistic Evolution of r^* only Due to Demographics

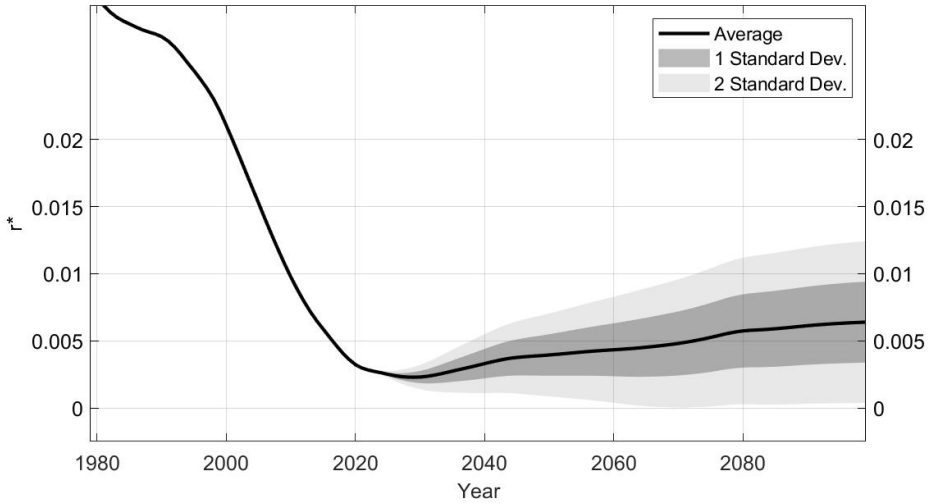

Notes: The figure shows the one– and two–standard deviation confidence intervals for the Monte Carlo simulations of the path of r^* only due to demographics. The ratio of government debt to GDP converges to 200% in 2080. For results with a debt-to-GDP ratio converging to 100%, see Appendix Figure E33.

Demographics. Figure 9 presents a fan chart based on demographic draws, highlighting one– and two–standard deviation confidence intervals. Moreover, Appendix Figure E28 displays the r^* distribution across various years and the distribution of the year when r^* reaches a trough, provided it is attained before 2050.

The two–standard deviation confidence interval of r^* in 2100 ranges from 0.04% to 1.24%,

giving a range of 1.2 pp. Given demographic forces alone, r^* is unlikely to fall to negative territory but also not likely to rise above the level we last saw near 2010. The modal level of r^* in 2030 is about 0.21%, increasing to about 0.32% by 2050 and 0.7% by 2100. By 2100, the distribution of the natural rate approaches a bell curve. Finally, the distribution of the trough period is heavily skewed to earlier years, with its mode being 2026.

Considering only the uncertainty surrounding demographics, we can draw a few conclusions. First, the demographics driver alone is very unlikely to cause a strong reversion in the downward trend. By itself, and given the future path of government debt, it is also extremely unlikely to pull r^* into negative territory. Furthermore, a trough is expected in the next few years, but—conditional on the 2015 values of the other drivers—r^* will most likely stay below 0.5% until 2040. As we will see below, although the confidence intervals are larger, the overall conclusions remain valid when we consider the other drivers.

Multiple Drivers. Figure 10 shows the fan chart for the future path of r^* once we draw realizations for all of the main drivers. Appendix Figure E29 shows the distribution of r^* in various years and the distribution of the years when r^* reaches its trough.

Figure 10: Probabilistic Evolution of r^* Due to Multiple Drivers

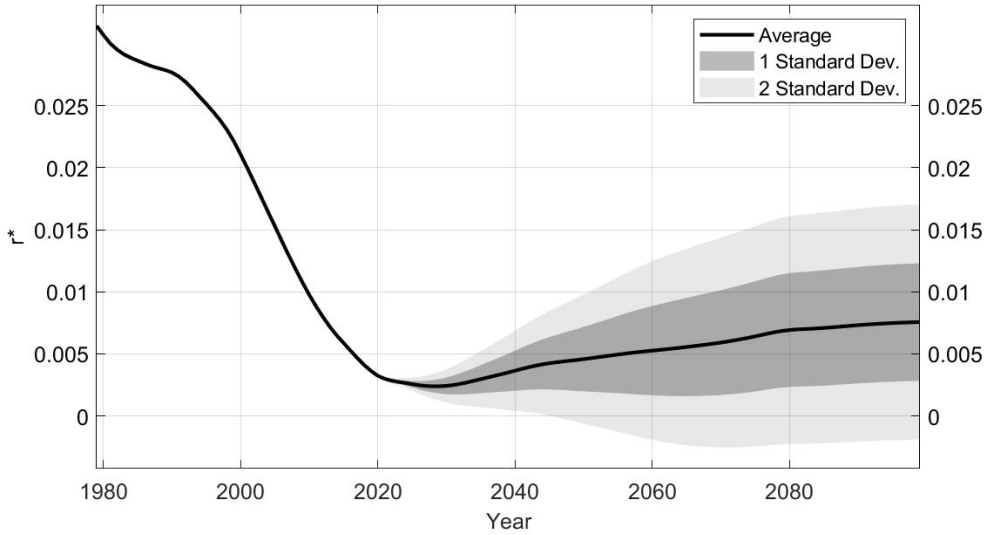

Notes: The figure shows the one– and two–standard deviation confidence intervals for the Monte Carlo simulations of the path of r^* due to changes in demographics, productivity, permanent inequality, and the labor share. The ratio of government debt to GDP converges to 200% in 2080.

Naturally, the two–standard deviation confidence interval widens and now ranges from -0.18 to 1.70 in 2100, a range of close to 2 pp. The upper bound is still below the model-

implied natural rate in 2000, or about 2%, and considerably below values for r^* before 2000. The one–standard deviation confidence interval ranges from 0.29% to 1.23%—that is, two-thirds of the probability mass lies within this 0.94 pp interval. The modal value of r^* in 2100 is 0.7%, and a considerable mass for the trough year of r^* is still around 2030 (Figure E29). Looking ahead, while there is a high probability that r^* will increase going forward, a rise above 2% remains extremely unlikely.

The results of this section suggest that, given our assumptions regarding the evolution of the most important drivers, a reversion of the downward trend is likely. For instance, our projection exercise including all drivers indicates that the probabilities that r^* will be above its 2020 level in 2030, 2040, and 2050 are respectively 10%, 56%, and 67%.

On the other hand, we find that the likelihood that the natural rate will be more than 0.5% above its current level by 2050 lies at only 3%. Our conclusion from this exercise is that, despite a predicted rise in public debt, it is highly unlikely that the drivers we consider will lead to a significant rise in r^* relative to its current level.

7 Policy Analysis

In this section we study the impact of policy on the natural rate. The government is an important actor when it comes to equilibrium determination in the market for savings. Our model includes several variables under direct influence of policy. Adjusting the tax and transfer system, including social security, affects the incentives to save. The level of public debt directly enters the demand for savings. The results in this section suggest that policy can have a significant impact on the level of r^*. We therefore propose treating the level of the natural rate as a *policy choice*.

We use our model to conduct a positive analysis and calculate the sensitivity of the natural rate to policy variables. The question of the optimal policy choice is normative and has to take into account all relevant costs and benefits. Our model is purely real and abstracts from important costs related to low levels of r^*, such as the impact on monetary policy when the policy rate is close to the effective lower bound and the role of labor supply decisions over the life cycle. We leave a more thorough normative analysis for future research.

We undertake a comparative-statics exercise at the baseline 2015 steady state. Table 5 shows the results. The first and second columns list the policy variables and their values in the baseline calibration. Column 3 shows the response of r^* (in basis points) to a 1 pp

Table 5: Response of r^* and Earnings Taxes to Change in Policy Parameters

Policy Parameter/Instrument	Baseline Value	$\frac{dr^*}{dx}$	$\frac{dTinc}{dx}$
Consumption Tax (%)	5.0	-2.75	-0.38
Profit Tax (%)	25.0	0.58	-0.14
Capital Income Tax (%)	30.0	1.85	0.04
Estate Tax (%)	10.0	-0.63	-0.06
Progressivity of Earnings Tax	0.18	5.24	-0.05
Basic Income (% of avg. Income)	0.0	6.36	0.78
Public Debt-to-GDP (%)	72.4	1.00	-0.03
Government Consumption (% of GDP)	14.9	9.44	0.77
OOP Subsidy (%)	0.0	0.51	0.004
Retirement Replacement Rate	1.43	1.17	0.02

Notes: Responses at baseline 2015 steady state. Column 1 lists the policy parameters; column 2 lists their value at the baseline 2015 steady state. Values in column 3 denote response of r^* in basis points to change in parameter by 1 pp. Column 4 shows response of labor income tax receipts relative to output in percentage points to change in parameter by 1 pp. OOP refers to out-of-pocket expenditures.

change in the policy parameter. Column 4 shows the impact on labor income tax revenue. Throughout, we let the level parameter of the income tax, λ_0, adjust to ensure that the budget constraint holds at the new steady state.

The table shows that a change in the consumption tax rate τ_c of 1 pp from the baseline value of 5% to 6% leads to a decline in r^* by 2.75 basis points (bp). The response of r^* turns out to be fairly linear for moderate changes.[26] Consequently, given that the steady-state value of τ_c is 5%, eliminating the consumption tax would increase r^* by roughly 14 bp.

Taxes on Return on Wealth. The general picture emerging from this exercise is that policy can have a substantial effect on r^*. For example, a revenue-neutral increase in the capital income tax of 10 pp would lead to a rise in the natural rate of about 18.5 bp, while a similar increase in the profit tax would lead to an increase of 5.8 bp. In both cases, a reduction in the return on savings drives the results. The same revenue-neutral increase in the estate tax, which would double it, however, would bring a decline of 6.3 bp. The reason for this (relatively modest) decline is that most of the bequests are bequeathed fairly late in life, and thus their reduction leads to an increase in the desire to save.

[26] Appendix Table F11 replicates Table 5 but with 10 pp increases in each instrument.

Progressivity of Earnings Tax. Parameter λ_1 affects tax progressivity. As shown in Heathcote, Storesletten and Violante (2017), the kind of tax specification we assume is a device that provides a good fit with the current US earnings tax schedule. However, changes in the parameter itself do not have a real-world counterpart and thus can be difficult to interpret, especially because λ_0 is also impacted in equilibrium. Despite this difficulty, we can map changes in λ_1 (and accompanying λ_0) onto changes in marginal tax rates in different parts of the earnings distribution.

Consider the 90th percentile of earnings. Here, an increase of 1 pp in λ_1 raises the marginal tax rate from 40.2% to 41.1% and the average tax rate from 27.0% to 27.2%.[27] At the same time, for the 50th percentile of earnings the marginal tax rate rises from 28.1% to 28.6% while the average tax rate declines from 12.4% to 11.8 %. Thus, this policy is highly redistributive. As a result, it achieves an increase in r^* of 5.2 bp.

Alternatively, raising λ_1 by 10 pp changes the marginal and average tax rates at the 90th percentile respectively to 49.5% and 29.7%, with corresponding numbers at the 50th percentile being 32.9% and 6.7%. This policy change has a very strong impact on r^*, increasing it by 0.56 pp. Even though this might be considered a large increase in the progressivity of taxes, a parameter value of $\lambda_1 = 0.281$ is far from inconceivable. In fact, considering total income, Kaas et al. (2020) estimate this parameter in Germany to be around 0.35 for most age ranges. Their minimum value is 0.29, for ages 55–64. In sum, we find that moving the US tax system toward a level of progressivity more in line with European countries can have a very strong impact on the natural rate of interest.[28]

Universal Basic Income for Workers. A basic income program of 1% of GDP raises the natural rate by 6 bp. This policy is highly redistributive, reducing the average tax rate (inclusive of transfers) at the 50th percentile of earnings from 12.4% to 11.5% while raising it at the 90th percentile from 27.0% to 27.4%. The impact of this policy is large, at 6.4 bp, while the transfer corresponds to roughly $1,250 per worker per year.[29] Correspondingly, a monthly transfer of $1,000 per worker, similar to what has recently been proposed by Andrew Yang, a former Democratic primary candidate for president, would raise the natural rate by 64 bp, virtually reversing the impact of inequality from 1965 to 2015.

Public Debt. Figure 11 shows the response of r^* to changes in the steady-state ratio of public debt to output plotted on the left axis as well as the effect on labor income tax

[27] Appendix C.4 discusses the marginal and average tax rates implied by our model and plots those rates.

[28] The top 10% post-tax earnings share falls 2.9 pp (from 54.5% to 51.6%) and the top 10% wealth share by 3.2 pp (from 28.2% to 25.0%) in response to the policy.

[29] Source: Penn World Tables, measured in 2017 dollars.

revenues and output, shown on the right axis. We see that the response of r^* to changes in $\frac{B}{Y}$ is modest, although convex: it would take a value of roughly 260% to raise r^* to a level of 3%.

Figure 11: Response of r^* and Earnings Taxes to Changes in Debt-to-Output Ratio

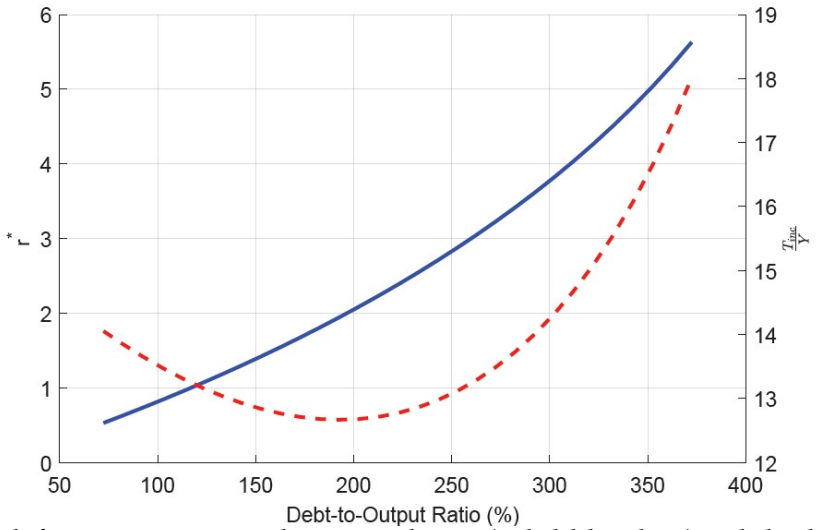

Notes: The left axis represents the natural rate (solid blue line), while the right axis represents total labor income tax receipts relative to output (dashed red line). Values for r^* and $\frac{T_{inc}}{Y}$ are in percentages. Labor income tax level parameter λ_0 adjusts to clear the government's budget constraint. All other parameters are held at the baseline 2015 steady-state value.

Locally around baseline values, our results suggest that a 1 pp increase in the debt-to-GDP ratio leads to a 1 bp increase in r^*. The number almost doubles once debt rises to 200% of GDP. Interestingly, increasing the steady-state ratio of public debt to output is a way for the government to *raise* resources:[30] the labor income tax receipts required to clear the government's budget constraint in the steady state fall as the debt ratio increases. In our baseline calibration, r^* (0.53%) is below the growth rate of the economy, consisting of roughly the sum of population and total factor productivity growth. In other words, $r^* - g < 0$. As a result, an expansion in government debt acts to reduce the total tax burden. In fact, this is minimized at the debt-to-GDP ratio of roughly 200%. This result suggests that debt-to-GDP ratios much higher than what is observed today (and in line with Congressional Budget Office projections) are sustainable—if not comfortably so. We provide a detailed analysis regarding the burden of government debt in Section G in the appendix, where we display the predicted evolution of $r^* - g$—that is, the interest rate net of the growth rate of the economy.

[30]This point has been made before—for example, in Mehrotra (2018).

Public Coverage of OOP Medical Expenditures. Our results indicate that a 1 pp increase in public coverage of OOP medical expenditures raises the natural rate by 0.5 bp, with virtually no impact on total labor income taxes (zero, up to the second decimal). That the impact is relatively small is consistent with the limited impact of the rise of the relative price of medical expenditures. However, results also indicate that covering 50% of OOP expenditures would lead to a 25 bp increase in r^*, a value we deem to be far from negligible.

Social Security. In our economy, individuals begin to accrue retirement income at 66, and the overall generosity of the system is determined by the parameter ϕ_{SS}. Figure 12 shows how the natural rate reacts to changes in the parameters of the social security system.[31] The base x is still computed based on income at age 65.

We find that a 1% reduction in the replacement-rate parameter ϕ decreases the natural rate by 1.17 bp, while a one-year increase in the age at which benefits begin to accrue lowers r^* by 5 bp. In both cases, the social security system becomes less generous, driving individuals to save more during their working life. Both effects are fairly linear for moderate changes. In addition, the bottom panel of Figure 12 highlights the dramatic effects of moving from a pay-as-you-go scheme to a fully funded one—that is, one with $\phi = 0$. In this case, the natural rate declines to -1%. Taken together, these results highlight the importance of retirement savings in determining the total demand for savings in the economy: increasing the necessity to save for retirement has a strong impact on the demand for savings, with a consequent very large impact on the natural rate.

The Power of Redistributive Policy. To illustrate how powerful public policy can be in affecting r^*, consider the following simultaneous changes, consisting of a set of redistributive policies: an increase of 5 pp in capital income and profit taxes, a UBI policy corresponding to a \$2,500 yearly transfer, a tax credit (subsidy) of 25% of OOP medical expenditures, and an increase in the parameter λ_1 by 2 bp.[32] Taken together, this suite of policies increases the natural rate of interest by 0.5 pp, all the way to 1.03%. For comparison, this corresponds to 70% of the total impact of inequality from 1965 to 2015, highlighting the powerful impact that redistributive policy can have in reversing the impact of inequality trends on r^*.

[31] In a previously circulated version of this paper, featuring endogenous labor supply and with individuals not only accruing retirement benefits at 66 but also effectively retiring at that age, we found that an increase of 1% in the parameter ϕ raises r^* by 1.66 bp, while delaying retirement by one year (including working for one additional year) increases r^* by less than 1 bp. We leave a more thorough analysis of the impact of retirement policies on the natural rate for future research.

[32] The resulting associated increases in the marginal and average tax rates at the 90th percentile are 3.5 pp and 0.7 pp. At the 50th percentile, these numbers are 3.0 pp and -3.6 pp.

Figure 12: Response of r^* to Changes in Social Security Parameters

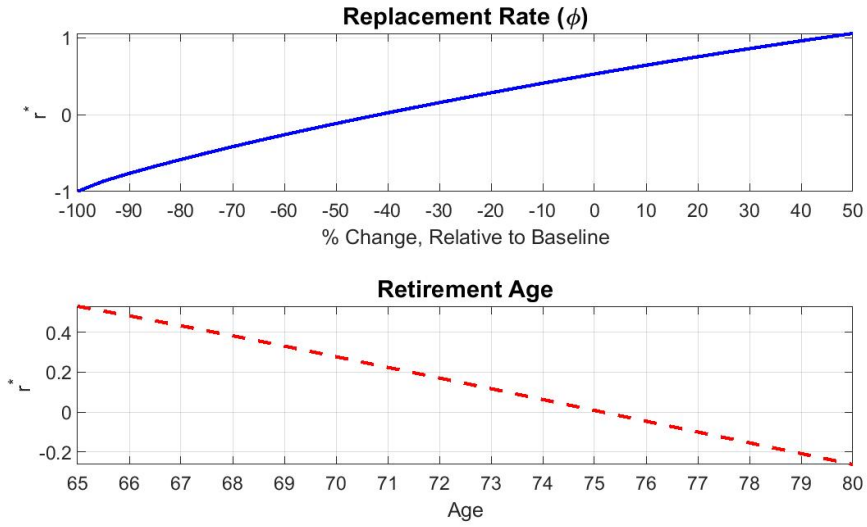

Notes: Response of r^* to a change in replacement-rate parameter ϕ (upper panel) and retirement year (bottom panel). Labor income tax level parameter λ_0 adjusts to clear the government's budget constraint. All other values are at the baseline 2015 steady-state value. Values for r^* are presented in percentages.

8 Sensitivity

Table 6 analyzes the sensitivity of the results in Section 5 to changes in the semi-elasticity of savings supply to r^*, which is a target in our calibration.[33] Appendix Tables H14 and H15 show the full sets of results. The total change in r^* ranges from 2.78 pp to 3.67 pp. Of the individual drivers, the impact of inequality is the most affected. Yet it remains important throughout the specifications. Appendix Figures H37 and H38 show that our results concerning the future outlook for r^* are insensitive to this parameter change.

9 Conclusion

In this paper we study the main drivers behind the decline of the natural rate of interest in the United States in recent decades, analyse how it will evolve, and gauge the impact of various policy instruments. To those ends, we developed a heterogeneous-agent

[33]We choose the targets as the midpoint between our baseline value and the boundaries of the interval of estimates in the literature as reported in Auclert et al. (2021). The calibration for the low target, an elasticity of 9.7, gives $\bar{\sigma} = 9.5$. This is above any of the estimates of the intertemporal elasticity of substitution surveyed in Eggertsson, Mehrotra and Robbins (2019). Therefore, we impose $\bar{\sigma} \leq 6$, the upper bound of estimates in the literature as reported in Eggertsson, Mehrotra and Robbins (2019). Appendix Tables H12 and H13 show the results for the low elasticity target of 9.7.

Table 6: Sensitivity Analysis - Change from 1965 to 2015 Steady State.

$\bar{\sigma}$	6	3.1	2.0
Elast. of Savings w.r.t. r^*	11.8	18.3	26.8
Demographics	0.65	0.72	0.58
Inequality	1.69	0.96	0.57
Productivity	1.21	1.44	1.62
Labor Share	0.48	0.60	0.69
Medical Exp.	0.33	0.21	0.14
Government Debt	-0.28	-0.29	-0.28
Government Consumption	0.21	0.16	0.13
NFA	0.45	0.36	0.28
Residual	0.10	0.10	0.10
Total	4.84	4.26	3.81

Notes: Baseline decomposition of changes (reproducing Table D8, column 3) with different assumptions regarding the elasticity of savings with respect to r^*. In each exercise, a (re)calibration is performed. NFA refers to the net foreign asset position

OLG model with non-homothetic preferences featuring the most important drivers of r^* suggested in the literature: demographics, income inequality, productivity growth, the labor share, government debt, government consumption, OOP medical expenditures, and international capital flows. The model accounts for a 4.3 pp decline in the natural rate between 1965 and 2015, within the range of existing empirical estimates.

In summary, we offer three takeaways. First, the decline in labor productivity, rising income inequality, changing demographics, and the fall in the labor share jointly account for the majority of the past observed downward trend. Second, looking ahead, these drivers are not likely to restore r^* to past values, such as those seen in the 2000s, even when we consider uncertainty regarding their future evolution. That said, although demographics will keep pulling the natural rate down, we predict a reversion in this decade, driven by public debt dynamics. Finally, we find that tax and social security policies can have a considerable impact on r^*, and a rather plausible package of redistributive policies can increase it substantially.

The likely evolution of r^* is still a matter of intense debate, and we believe that our paper provides substantial clarity about it. Yet we did not consider endogenous interactions between drivers—for instance, how demographic change affects productivity growth, which in turn impacts r^*—and the importance of potential new drivers, such as climate change and geopolitical factors. We hope future research, in which we plan to engage, will shed light on these considerations.

References

Alvaredo, Facundo, Bertrand Garbinti, and Thomas Piketty. 2017. "On the Share of Inheritance in Aggregate Wealth: Europe and the USA, 1900–2010." *Economica*, 84(334): 239–260.

Auclert, Adrien, and Matthew Rognlie. 2018. "Inequality and Aggregate Demand." National Bureau of Economic Research, Inc NBER Working Papers 24280.

Auclert, Adrien, Hannes Malmberg, Frederic Martenet, and Matthew Rognlie. 2021. "Demographics, Wealth, and Global Imbalances in the Twenty-First Century." National Bureau of Economic Research Working Paper 29161.

Barany, Zsofia, Nicolas Coeurdacier, and Stéphane Guibaud. 2018. "Capital Flows in an Aging World." Sciences Po Sciences Po publications 13180.

Barkai, Simcha. 2020. "Declining Labor and Capital Shares." *The Journal of Finance*, 75(5): 2421–2463.

Bauer, Michael D., and Glenn D. Rudebusch. 2020. "Interest Rates under Falling Stars." *American Economic Review*, 110(5): 1316–54.

BEA, U.S. Bureau of Economic Analysis. 2020. "National Income and Product Accounts." Online Edition.

Blanchard, Olivier. 2023. *Fiscal policy under low interest rates.* MIT press.

BLS, Bureau of Labor Statistics, U.S. Department of Labor. 2020. "Consumer Price Index (CU) Tables." Online Edition.

Caballero, Ricardo J., and Emmanuel Farhi. 2014. "The Safety Trap." National Bureau of Economic Research, Inc NBER Working Papers 19927.

Carroll, Christopher D. 2000. "Why Do the Rich Save So Much?" In *Does Atlas Shrug?: The Economic Consequences of Taxing the Rich.* , ed. J. Slemrod and Russell Sage Foundation, 465–84. Russell Sage Foundation.

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

Cesa-Bianchi, Ambrogio, Richard Harrison, and Rana Sajedi. 2022. "Decomposing the drivers of Global R*." Bank of England Staff Working Paper 990. Published on 12 July 2022.

Coeurdacier, Nicolas, Stéphane Guibaud, and Keyu Jin. 2015. "Credit Constraints and Growth in a Global Economy." *American Economic Review*, 105(9): 2838–81.

Congressional Budget Office. 2019. "The 2019 Long-Term Budget Outlook." Congressional Budget Office Reports 55331.

Congressional Budget Office. 2020. "An Update to the Economic Outlook: 2020 to 2030." Congressional Budget Office Reports 56517.

Congressional Budget Office. 2022. "The Budget and Economic Outlook: 2022 to 2032."

De Nardi, Mariacristina. 2004. "Wealth Inequality and Intergenerational Links." *Review of Economic Studies*, 71(3): 743–768.

De Nardi, Mariacristina, and Fang Yang. 2014. "Bequests and heterogeneity in retirement wealth." *European Economic Review*, 72(C): 182–196.

De Nardi, Mariacristina, Eric French, and John B. Jones. 2010. "Why Do the Elderly Save? The Role of Medical Expenses." *Journal of Political Economy*, 118(1): 39–75.

Dynan, Karen E., Jonathan Skinner, and Stephen P. Zeldes. 2004. "Do the Rich Save More?" *Journal of Political Economy*, 112(2): 397–444.

Eggertsson, Gauti B, Jacob A Robbins, and Ella Getz Wold. 2018. "Kaldor and Piketty's Facts: The Rise of Monopoly Power in the United States." National Bureau of Economic Research Working Paper 24287.

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

Feenstra, Robert C., Robert Inklaar, and Marcel P. Timmer. 2015. "The Next Generation of the Penn World Table." *American Economic Review*, 105(10): 3150–3182.

Fernald, John G. 2014. "A quarterly, utilization-adjusted series on total factor productivity." Federal Reserve Bank of San Francisco Working Paper Series 2012-19.

Gagnon, Etienne, Benjamin K. Johannsen, and David López-Salido. 2021. "Understanding the New Normal: The Role of Demographics." *IMF Economic Review*, 69(2): 357–390.

Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song. 2021. "What do data on millions of US workers reveal about lifecycle earnings dynamics?" *Econometrica*, 89(5): 2303–2339.

Guvenen, Fatih, Greg Kaplan, and Jae Song. 2014. "The Glass Ceiling and The Paper Floor: Gender Differences among Top Earners, 1981-2012." National Bureau of Economic Research Working Paper 20560.

Guvenen, Fatih, Gueorgui Kambourov, Burhanettin Kuruscu, Sergio Ocampo-Diaz, and Daphne Chen. 2019. "Use It or Lose It: Efficiency Gains from Wealth Taxation." National Bureau of Economic Research Working Paper 26284.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante. 2017. "Optimal Tax Progressivity: An Analytical Framework." *The Quarterly Journal of Economics*, 132(4): 1693–1754.

Heider, Florian, and Agnese Leonello. 2021. "Monetary policy in a low interest rate environment: Reversal rate and risk-taking."

Hendricks, Lutz. 2001. "Bequests and Retirement Wealth in the United States." mimeo.

Holston, Kathryn, Thomas Laubach, and John C. Williams. 2017. "Measuring the natural rate of interest: International trends and determinants." *Journal of International Economics*, 108: S59–S75. 39th Annual NBER International Seminar on Macroeconomics.

Houthakker, H. S. 1960. "Additive Preferences." *Econometrica*, 28(2): 244–257.

Huggett, Mark, and Gustavo Ventura. 2000. "Understanding why high income households save more than low income households." *Journal of Monetary Economics*, 45(2): 361–397.

Hurd, Michael D, and James P Smith. 2002. "Expected bequests and their distribution."

Jones, Callum. 2018. "Aging, Secular Stagnation and the Business Cycle." International Monetary Fund IMF Working Papers 2018/067.

Kaas, Leo, Georgi Kocharkov, Edgar Preugschat, and Nawid Siassi. 2020. "Low Homeownership in Germany—a Quantitative Exploration." *Journal of the European Economic Association*, 19(1): 128–164.

Kitao, Sagiri. 2014. "A life-cycle model of unemployment and disability insurance." *Journal of Monetary Economics*, 68: 1–18.

Kopecky, Joseph, and Alan M. Taylor. 2020. "The Murder-Suicide of the Rentier: Population Aging and the Risk Premium." Trinity College Dublin, Department of Economics Trinity Economics Papers tep1220.

Kopecky, Karen A., and Tatyana Koreshkova. 2014. "The Impact of Medical and Nursing Home Expenses on Savings." *American Economic Journal: Macroeconomics*, 6(3): 29–72.

Laubach, Thomas, and John C. Williams. 2003. "Measuring the Natural Rate of Interest." *The Review of Economics and Statistics*, 85(4): 1063–1070.

Lisack, Noëmie, Rana Sajedi, and Gregory Thwaites. 2019. "Population Ageing and the Macroeconomy." Banque de France Working papers 745.

Liu, Ernest, Atif Mian, and Amir Sufi. 2022. "Low Interest Rates, Market Power, and Productivity Growth." *Econometrica*, 90(1): 193–221.

Mehrotra, Neil. 2018. "Implications of Low Productivity Growth for Debt Sustainability." mimeo.

Mian, Atif, Ludwig Straub, and Amir Sufi. 2021. "The Saving Glut of the Rich." Replication kit available here.

Piketty, Thomas, and Emmanuel Saez. 2003. "Income Inequality in the United States, 1913–1998." *The Quarterly Journal of Economics*, 118(1): 1–41.

Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman. 2018. "Distributional National Accounts: Methods and Estimates for the United States." *The Quarterly Journal of Economics*, 133(2): 553–609.

Platzer, Josef, Robin Tietz, and Jesper Linde. 2022. "Natural versus Neutral Rate of Interest: Parsing Disagreement about Future Short-Term Interest Rates."

Porcellacchia, Davide. 2023. "The Tipping Point: Interest Rates and Financial Stability."

Rachel, Lukasz, and Lawrence H. Summers. 2019. "On Secular Stagnation in the Industrialized World Keywords: stagnation, fiscal policy, economic recovery, industrial economy. market indicators, interest rates, monetary policy, regulatory policy." *Brookings Papers on Economic Activity*, 50(1 (Spring): 1–76.

Straub, Ludwig. 2019. "Consumption, Savings, and the Distribution of Permanent Income." Revise and resubmit at Econometrica.

Tauchen, George. 1986. "Finite state markov-chain approximations to univariate and vector autoregressions." *Economics Letters*, 20(2): 177–181.

United Nations. 2022*a*. "World Population Prospects 2022: Methodology of the United Nations population estimates and projections." Department of Economic and Social Affairs, Population Division UN DESA/POP/ 2022/TR/NO. 4.

United Nations. 2022*b*. "World Population Prospects 2022, Online Edition." Department of Economic and Social Affairs, Population Division.

A Additional Data Details

A.1 Interest Rates in the United States

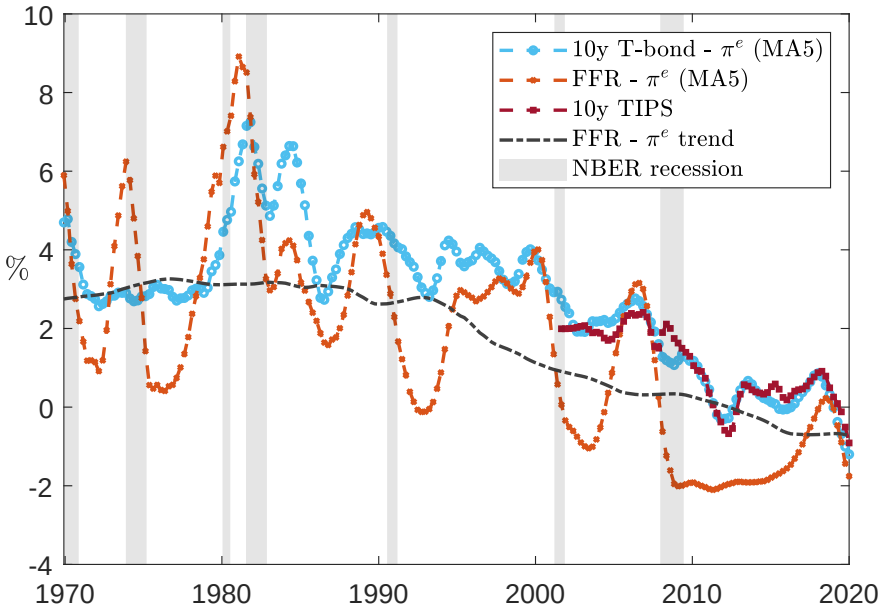

Figure A13: Various interest rates and trend of short-term interest rate. Details see text.

Figure A13 shows the yield on a 10 year U.S. Treasury Bond minus inflation expectations, the U.S. Federal Funds Rate (FFR) minus inflation expectations, and the yield on a 10 year Treasury-Inflation Protected Security (TIPS). All series are smoothed using a five period moving average, MA(5). In addition, "FFR - π^e trend" shows U.S. Federal Funds Rate (FFR) minus inflation expectations smoothed using MA(120). For quarterly data this corresponds to a moving average over 30 years.[34] The figure also shows NBER recessions.

A.2 Additional Details on Data Sources

Sources - Public Policy variables. We use data that supplements Congressional Budget Office (2020). Debt is "Debt Held by the Public". G is constructed as "Total Outlays" - "Net Interest" - "Social Security" - "Federal Civilian and Military Retirement". The later three items come from the model, so we exclude them from exogenous government spending G. In the data "Social Security" + "Federal Civilian and Military Retirement" amount to 5.0% and 5.8% in 1975 and 2015, respectively.

[34]The data are taken from Bauer and Rudebusch (2020). Inflation expectations are based on surveys.

Sources - Top Earnings Shares. Earnings shares data is taken from Piketty and Saez (2003). We use wage income shares, table B2 from the updated file February 2020, from Emmanuel Saez' webpage. See https://eml.berkeley.edu/∼ saez/, accessed 8/5/2020.

Productivity Growth. For the TFP growth rate γ we rely on estimates from Fernald (2014), covering the full postwar period.[35]

Labor share. We use labor share data from Penn World Table, version 10.0, variable label "labsh" (Feenstra, Inklaar and Timmer, 2015).

Relative Price of OOP Medical Expenditures. We use data from BLS (2020). We use "All items in U.S. city average, all urban consumers, seasonally adjusted", code "CUSR0000SA0" for the general price index and "Medical care in U.S. city average, all urban consumers, seasonally adjusted", code "CUSR0000SAM" for price index of medical goods and services. We use the amount that medical care increased relative to the general index as measure for increase in OOP price.

Intermediation wedge. The value of the intermediation wedge is $\nu = 0.02$. This value is between the average spreads of U.S. 10 year government bond yields and average investment grade corporate bond yields, from 2010 to 2019. We use series "Moody's Seasoned Aaa Corporate Bond Yield [DAAA]," "Moody's Seasoned Baa Corporate Bond Yield [DBAA]," and "Market Yield on U.S. Treasury Securities at 10-Year Constant Maturity, Quoted on an Investment Basis [DGS10]." All series are retrieved from FRED, Federal Reserve Bank of St. Louis.

Capital-to-Output Ratio. The calibration target for the capital-to-output ratio is $\frac{K}{Y} = 3.7$. This is about the mid-point of a more narrow and a broader measure of capital: the first uses NIPA tables from the U.S. Bureau of Economic Analysis (BEA), "Table 1.1 Current-Cost Net Stock of Fixed Assets and Consumer Durable Goods," line 2 ("Fixed assets") and "Table 1.1.5. Gross Domestic Product," line 1 ("GDP"), which gives a 2010 to 2019 average of $\frac{K}{Y} = 3.05$. The second uses net national wealth relative to Net National Income from WID, which we can denote $\frac{A}{Y}$. Following Auclert et al. (2021), we use this figure to calculate the capital-to-output ratio as a residual: $\frac{K}{Y} = \frac{A}{Y} - \frac{NIIP}{Y} - \frac{B}{Y}$, where $\frac{NIIP}{Y}$ is the net

[35] An updated TFP series is provided at https://www.frbsf.org/economic-research/indicators-data/total-factor-productivity-tfp/

international investment position as a share of output, and $\frac{B}{Y}$ is our measure of public debt. With this approach we get $\frac{K}{Y} = 4.29$.

A.3 Changes in Idiosyncratic Risk

Our baseline calibration for the parameter σ_ϵ, representing the standard deviation of the idiosyncratic earnings shock, is $\sigma_\epsilon = 0.04$. We compute counterfactual changes in this parameter by taking into account changes in the unexplained component of earnings, using a standard wage dispersion regression.

We use data from the Panel Study of Income Dynamics (PSID) from 1970 to 2017. At each year of the sample, we run a cross-sectional regression of total yearly earnings on individual fixed effects, a quadratic polynomial on age, dummies for education years, race, sex, state, and interactions between dummies for sex and age, sex and years of education, and education and race. We then compute the variance of the regression residual and compute its relative changes over time. We then adjust σ_ϵ so that changes in the cross-sectional variance of earnings (conditional on the life cycle profile) in the model matches the one in the data for each period of our analysis.

Figure A14 below shows the evolution of the cross-sectional variance of the regression residual. The profile shows a marked increase in the variance of the unexplained component of earnings, which we find to have a limited impact on r^*.

Figure A14: Cross-Sectional Variance of the Regression Residual

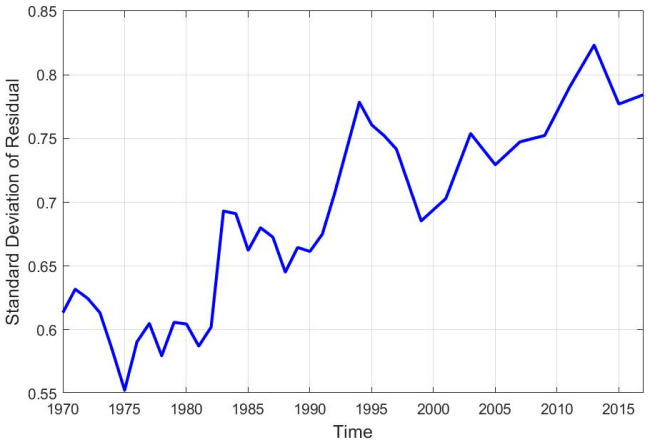

Notes: The figure shows the variance of the residual of a cross-sectional regression, for each year in our sample, of total yearly earnings on individual fixed effects, a quadratic polynomial on age, dummies for education years, race, sex, state, and interactions between dummies for sex and age, sex and years of education, and education and race.

B Additional Model Details

B.1 Demographics - Details

Let N_{gt} be the *size* of generation g at time t, M_{gt} be the *size* of the net migration flow at time t. Accordingly, $\phi_{gt} \equiv \frac{N_{gt}}{N_t}$, with $\sum_g N_{gt} = N_t$. The flow of population aged g is determined by:

$$N_{gt} = p_{g-1,t-1} N_{g-1,t-1} + M_{g-1,t-1}$$

This can be stated in terms of shares. Dividing through by N_t:

$$\begin{aligned} \phi_{gt} &= p_{g-1,t-1} \frac{N_{g-1,t-1}}{N_t} + \frac{M_{g-1,t-1}}{N_t} \\ &= \frac{1}{1+n_t} (p_{g-1,t-1} \phi_{g-1,t-1} + m_{g-1,t-1}), \end{aligned}$$

where $m_{g-1,t-1} \equiv \frac{M_{g-1,t-1}}{N_{t-1}}$. The expression above can be inverted to compute migration flows, given (data on) population shares, survival rates, and population growth.

Throughout or analyses, we assume that, conditional on a particular age g, the wealth and earnings distributions among migrants is the same as that among domestic citizens.

B.2 Production Side - Intermediate Goods

The aggregation of intermediate goods follows a CES function. Here i denotes an infinitesimal sized firm:

$$Y_t = \left[\int_0^1 y_{it}^{\frac{\theta_t - 1}{\theta_t}} \right]^{\frac{\theta_t}{\theta_t - 1}} \tag{15}$$

with retailer elasticity of substitution θ_t. The markup is then given by $\mu_t = \frac{\theta_t}{\theta_t - 1}$

The aggregate capital stock K_t evolves according to the standard law of motion

$$K_{t+1} = (1 - \delta) K_t + I_t \tag{16}$$

where δ is the depreciation rate and I_t denotes aggregate investment.

B.3 Means-tested benefit

The means-tested benefit ensures a consumption floor $\underline{c}$. The transfer is only available to households that run down all of their assets. Since the social security system usually provides enough benefits to raise consumption above the consumption floor, total amount of transfers usually turns out to be small. In the 2015 steady state, no household turns out to be eligible for the means-tested transfer.

Omitting time subscripts, The transfer $T^M \equiv T^M(z,h,a,g)$ is given as:

$$
\begin{aligned}
&T^M = 0,\ m = M \text{ if: } \ (1+\tau_c)\underline{c} + m \leq (1+ret)a + (1-\tau^{beq})beq + \xi + T^M + \lambda_0(wz)^{1-\lambda_1} \\
&T^M = 0,\ m = \max\left(\underline{M}, \min\left(M, (1+ret)a + (1-\tau^{beq})beq + \xi + T^M + \lambda_0(wz)^{1-\lambda_1} - (1+\tau_c)\underline{c}\right)\right) \\
&\quad \text{if: } \ (1+\tau_c)\underline{c} + \underline{M} < (1+ret)a + (1-\tau^{beq})beq + \xi + T^M + \lambda_0(wz)^{1-\lambda_1} \leq (1+\tau_c)\underline{c} + M \\
&T^M = \max\left(0, (1+\tau_c)\underline{c} + \underline{M} - (1+ret)a - (1-\tau^{beq})beq + \xi + T^M + \lambda_0(wz)^{1-\lambda_1}\right),\ m = \underline{M} \\
&\quad \text{if: } \ (1+ret)a + (1-\tau^{beq})beq + \xi + T^M + \lambda_0(wz)^{1-\lambda_1} \leq (1+\tau_c)\underline{c} + \underline{M}
\end{aligned}
$$

where $\underline{M}$ is a basic OOP medical expense. What these conditions state is that the transfer T^M is means-tested and only available when initial wealth $((1+ret)a + (1-\tau^{beq})beq)$ is null and consumption floor $\underline{c}$ and basic medical expenses $\underline{M}$ surpass disposable income when old. We will typically make the following assumptions on medical expenses: for low permanent types $\underline{M}$ coincides with the full medical expense incurred. High permanent types typically have higher out-of-pocket medical expenses M. Once their consumption becomes constrained, they will first cut consumption until they reach $\underline{c}$, then reduce medical spending until $m = \underline{M}$. Once these conditions are fulfilled, they are also eligible for means-tested benefit T^M. We set $\underline{M}$ of high permanent types equal to the OOP medical expenses of low permanent types.

B.4 First-order Conditions

Individual's first-order conditions:

$$
\begin{aligned}
\frac{1}{1+\tau_c}u_c(c,g) \geq{}& \beta p_{g,t+1}\mathbb{E}_{(z',h')|(z,h)}\left[V'_a(z',h',a',g+1)\right] \\
&+ \beta(1-p_{g,t+1})b_0(1-\tau'_{beq})\frac{1}{o'}\left(k_b + (1-\tau'_{beq})\frac{a'}{o'}\right)^{-b_{1,g}}, \qquad (17)
\end{aligned}
$$

where $V_a(z,h,a,g) = \frac{1}{1+\tau_c}(1+ret)u_c(c,g)$. The first-order conditions will hold with equality so long as $a' > 0$.

B.5 Equilibrium Conditions

In the equilibrium, households maximize their problem, firms maximize profits, and the government's debt evolves according to (10), taking the (path of) real interest rates and real wages as given. Let $\lambda(z, a, g)$ be a measure of states, which includes age-specific population shares.This measure is normalized to one in every period.

For convenience, we repeat the market clearing condition for assets:

$$B_{t+1} + K_{t+1} = \sum_{g} \sum_{z,h,a} g_t^a(z, h, a, g)\lambda_t(z, h, a, g) - NFA_t \tag{18}$$

Goods market clearing is given by:

$$Y_t = \sum_{g} \sum_{z,h,a} \left[g_t^c(z, a, g) + p_m m_g(z, h)\right] \lambda_t(z, h, a; g) + K_{t+1} - (1 - \delta - \nu)K_t + G_t + NX_t, \tag{19}$$

where $g_c(\cdot)$ is the policy function for consumption and NX_t corresponds to net exports:

$$NX_t = -(NFA_t - (1 + ret_t)NFA_{t-1})$$

The market clearing condition for effective labor is:

$$L_t^D = \sum_{z,h,a,g} z\lambda_t(z, h, a, g), \tag{20}$$

with L^D aggregate labor demand from the supply side. Firms' labor demand has to be met by labor supply by households taking into account their individual productivity.

Total bequests are given by:

$$beq_t = (1 + ret_t) \sum_{g=1}^{G} (1 - p_{g-1,t-1})\lambda_{t-1}(z, h, a, g)g_{t-1}^a(z, h, a, g) \tag{21}$$

Total government primary deficit, $\Xi_t - \Lambda_t$, is given by:

$$\Xi_t - \Lambda_t = G_t + \sum_{z,h,a,g} \lambda_t(z,h,a,g)\xi(z,g) + \sum_{z,h,a,g} \lambda_t(z,h,a,g)T^M(z,h,a,g)$$
$$- \sum_{z,h,a,g} \lambda_t(z,h,a,g)\left(T_y(wz) + \tau^c g_t^c(z,h,a,g)\right) - \tau^{beq} beq_t - \tau^{corp} d_t - \tau^k(r_t^k K_t + r_t^* B_t)$$

The law of motion for the measure $\lambda(\cdot)$ is

$$\lambda_{t+1}(z',h',a',g) = \left(\frac{\phi_{g,t+1}}{\phi_{g-1,t}}\right) \sum_{\substack{z,h,a \\ a'=g_t^a(z,h,a,g-1)}} P(h'|h)P(z'|z)\lambda_t(z,h,a,g-1) \text{ if } g > 1 \quad (22)$$

In the equation above, the term $\left(\frac{\phi_{g,t+1}}{\phi_{g-1,t}}\right)$ represents the (gross) yearly growth rate of a particular cohort aged $g-1$ at year $t-1$, including migration.

For the cohorts entering the model, their non-bequest wealth is null:

$$\lambda_{t+1}(z,0,1) = \frac{n_t}{1+n_t} P(z),$$

with the absence of h reflecting the fact that health shocks are absent for individuals aged 65 and younger. Finally, note that $\lambda_{t+1}(z,h,1) = 0$ if $a \neq 0$. In the definition above, $P(z)$ corresponds to the ergodic distribution of labor productivities (including permanent, transitory, and star-earner components).

B.6 Definition of Competitive Equilibrium

A competitive equilibrium in the baseline model of section 3 is a set of aggregate allocations $\{Y_t, K_t, L_t\}_{t=0}^{\infty}$, price processes $\{r_t, w_t, d_t, ret_t\}$, distribution $\{\lambda_t\}_{t=0}^{\infty}$, paths for policy functions $\{g_t^a(z,h,a,g)\}_{t=0}^{\infty}$ and $\{g_t^c(z,h,a,g),\}_{t=0}^{\infty}$ that jointly satisfy:

1. Household optimality conditions are satisfied.
2. Firm first-order conditions yield competitive input prices ((4) and (5))
3. Government budget constraint (10)
4. Asset, goods, labor markets clear ((8), (19), and (20))

5. Law of motion for the measure $\lambda(\cdot)$ (22)
6. Total dividends are given by (6), and returns given by (9)
7. The aggregate bequests received are consistent with those implied by choices by (soon-to-be) deceased individuals (21).

As mentioned before, in the steady-state equilibrium $\tilde{Y}_t$, $\tilde{K}_t$, $\tilde{L}_t$, and $\tilde{B}_t$ (among others) are constant over time, where $\tilde{X}_t = \frac{X_t}{N_t A_t}$. In that case, output, capital, (effective) labor, and total government debt growth at at a rate $(1+n)(1+g)-1$.

C Additional Details - Functional Forms and Calibration

C.1 Non-Homothetic Preferences

Figure C15 plots the profile of σ_g under the baseline calibration.

Figure C15: Profile of σ_g under baseline calibration

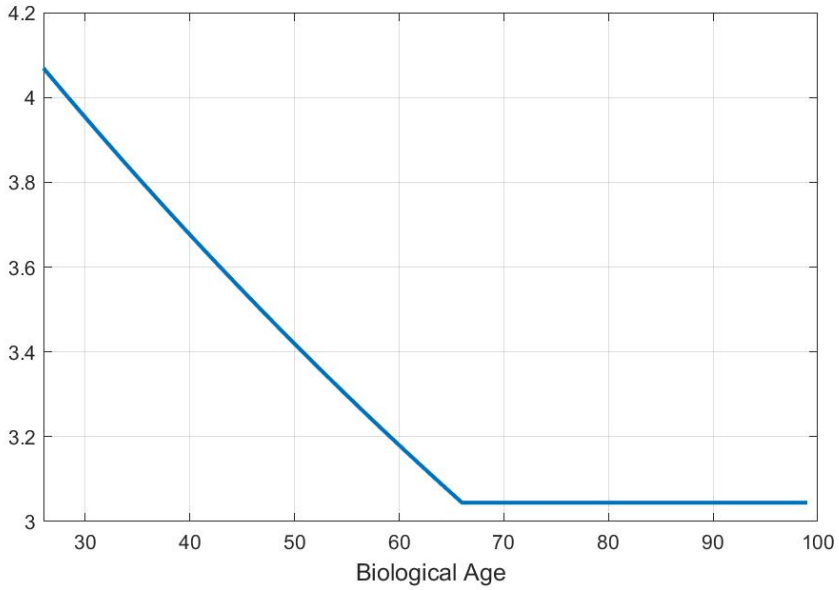

Notes: The figure shows the different values σ assumes at different ages in out baseline calibration.

In order to build further intuition for addilog preferences, we discuss a simple two-goods example: assume a consumer with addilog preferences as in equation 11, where g refers to goods instead of age. Denote the consumption of two goods as c_L and c_H, respectively, and their corresponding sigmas σ_L and σ_H, with $\sigma_L < \sigma_H$. For households with a low level

of income, almost all consumption is of the good c_H – that is, the one with the larger sigma, σ_H. The expenditure share of good c_H is high for such a household. For a rich household, almost all consumption is of good c_L. As we increase the income of the poor household, the expenditure on c_L rises, as the income elasticity of good c_L is larger than 1. While the expenditure on c_H can also rise, in relative terms it decreases; that is, the expenditure share on c_H declines in income. With the parameter o we can manage how fast this transition happens, or equivalently, what we mean by "low" and "high" income in this example. For instance, the level of o pins down the level of income at which expenditure shares on both goods are equal. Clearly, setting o very low or very high would not make sense, since then all agents would almost exclusively consume either good c_H or c_L, depending on the case.

Going back to our model economy, we assume that σ_g are monotonically decreasing in g. In the OLG setting, this implies that high-income households want to consume a larger share of their lifetime income when they are old. Consequently, savings rates are higher for high-income workers so that they can afford the desired consumption when they are old. We do not want this effect to hold true for low-income households, as empirically it is the case that they accumulate almost no savings over their lifetime.

The motivation for this preference structure relates to the empirical regularity that high-income individuals have higher savings rates, or, equivalently, that the marginal propensity to consume out of *permanent* income decreases with income.[36] One way to measure this is to look at the elasticity of consumption out of permanent income, ϕ_{PI}. Straub (2019) estimates ϕ_{PI} using data from the Panel Study of Income Dynamics (PSID) and finds a value of 0.7 – significantly below 1. Permanent income is used here to abstract from idiosyncratic earnings shocks that eventually reverse and imply a different consumption response. $\phi_{PI} < 1$ implies that high-income households have higher savings rates during their working lives.

An elasticity of consumption out of permanent income lower than 1 can also result for reasons other than nonhomothetic consumption preferences. Within our model, a nonhomothetic bequest motive or a redistributive tax-and-transfer system can have the same effect. However, as Straub (2019) shows and we confirm, the calibration procedure cannot match $\sigma_{PI} = 0.7$ together with the calibration targets for other parameters. Thus, addilog preferences are a tractable way to bring the model closer to an aspect of the data that is highly relevant for the impact of inequality on aggregate savings.

[36] See Carroll (2000) and Dynan, Skinner and Zeldes (2004) for evidence on savings rates by income group.

C.2 Life-Cycle Earnings

To compute the life cycle component of earnings, we proceed as follows. For each level of permanent income, we compute a *pseudo* life cycle earnings profile using a quadratic polynomial. In particular, we set the maximum increase in life-cycle earnings of 60% for the low permanent type and 4.8-fold for the high permanent type, with corresponding peak ages being 49 and 54, according to Guvenen et al. (2021), figure 11(a).[37] For the low permanent type we take a 60% maximum income growth over the life-cycle ($e_L^{lc} = 0.47$), which is equal to the increase of lifetime earnings between age 25 and 55 of the median worker. The high type increase we set to ($e_H^{lc} = 1.57$), which corresponds to the lifetime earnings growth between age 25 and 55 at the 95th percentile. [38]

We then use this *pseudo* profile to compute the relative growth in earnings across permanent income levels, keeping the average profile fixed, taken from LIS. In other words, we ensure that different permanent income levels experience different earnings growth profiles over the life cycle, but the average profile matches the data in LIS. Following, Guvenen et al. (2019), the resulting polynomials ensure that the proposes life-cycle average earnings profile is implemented:

$$h_g^z(e^H) = \frac{56 \times (g-1) - (g-1)^2}{28^2} \times (1.57) \; g \in [1, 60]$$
$$h_g^z(e^L) = \frac{46 \times (g-1) - (g-1)^2}{23^2} \times (0.47) \; g \in [1, 60]$$

In the equations above, note that 1.57 and 0.47 correspond to the maximum increase in life-cycle earnings in log points ($\log 4.8 = 1.57$ and $\log 1.60 = 0.47$). Also, due to the lack of LIS data, we assume that earnings are null after individuals reach the biological age of 86. In any case note that earnings at these ages are very low.

Figure C16 below displays the resulting life-cycle profiles for each worker, as well as the average one (obtained directly from LIS).

[37] The 4.8-fold increase is based on workers at the 95th percentile of the income distribution

[38] For lack of exact data, we set the peak increase to age 54, which corresponds to the maximum of the estimated life-cycle profile of average log earnings in Guvenen et al. (2021), online appendix figure C.36. Online appendix figure C.37 suggests that the peak for the median worker is reached before age 55, but above age 35. The exact peak age cannot be determined from this figure. We decide for a peak of biological age of 49 since average income growth between age 45 and 55 for the median worker is slightly negative.

Figure C16: Earnings Profile - Average and Permanent Types

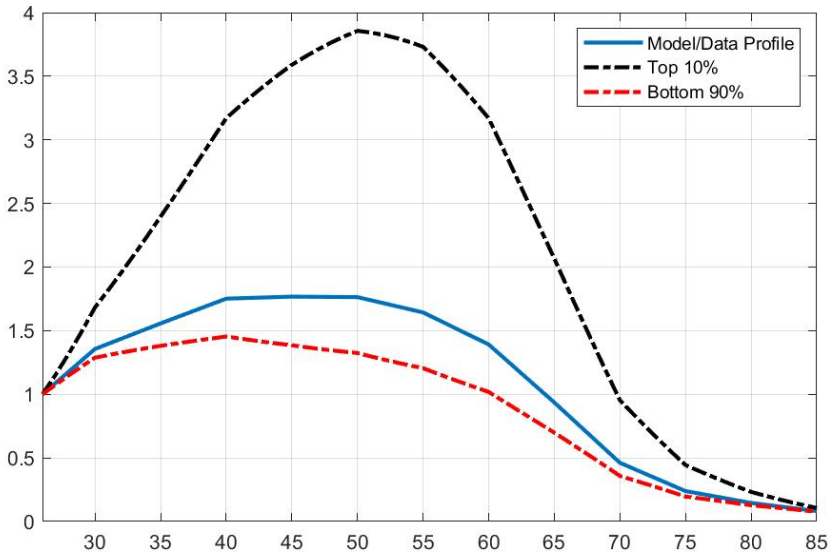

Notes: The figure shows how average income for e^h (black), e^L (red), and for all (blue) individuals in the economy. In each case, earnings are relative to their value at age 26.

C.3 OOP Expenses Process

We take the deterministic expense profile from Kopecky and Koreshkova (2014). They use a fixed-effects estimator on medical expenses excluding nursing home stay from the HRS. They include permanent income dummies and find considerable differences of OOP health expenses by permanent income. In our model, households can face one of three profiles of $h_m^{z,g}$ (details below). Households with high permanent productivity type e^H have the highest profile. Households of e^L type and that enter retirement with below median ζ_i realizations have a low $h_m^{z,g}$ profile, the remaining e^L households an intermediate profile. Figure C17 shows $\exp(h_m^{z,g})$ if in bad health state, by productivity state z_i.

Transition probabilities and size of the shock are taken from Kopecky and Koreshkova (2014) as well. We transform the 4 state Markov chain to a two state chain. Being in bad health raises OOP medical expenditure by about 13.5 times. About a third of retired households finds themselves in bad health condition.

$h_m^{z,g}$ profiles: For the e^H type we take quintile 5 estimate of the deterministic profile in Kopecky and Koreshkova (2014). For the e^L type, an average of their quintile 1 to quintile 4 lead to about 25% of retirees at the consumption floor. This is much higher than

Figure C17: Profile of out-of-pocket (OOP) Health Expenditure under the Bad Health State.

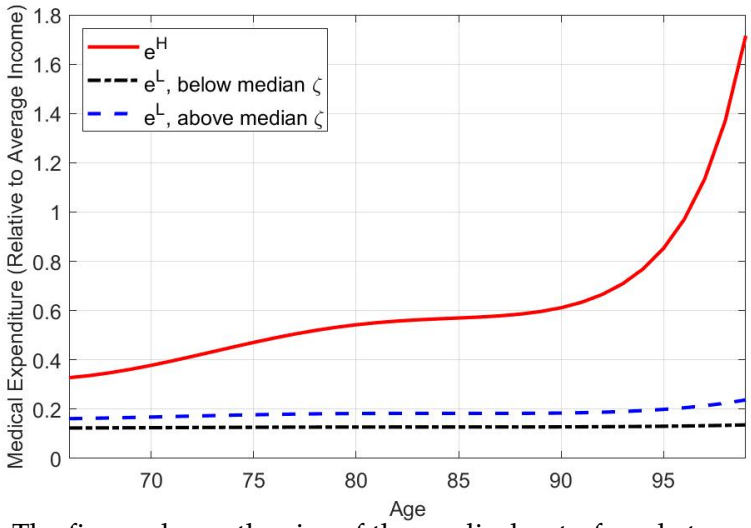

Notes: The figure shows the size of the medical out-of-pocket expenditure "bad health" shock, depending on their permanent income level and (for low permanent income individuals) their realization of the earnings risk shock at the age of 65.

empirical estimates of old-age poverty levels. Thus, we use quintiles 1 and 2 for below and above median ζ_i realizations, respectively. Kopecky and Koreshkova (2014) mention that quintile 1 and quintile 2 are qualitatively different than the other quintiles, likely because individuals in these groups rely more on government-provided services. We can then think of e^L types in our model mostly relying on government-provided health goods and services.

Shock process: The transition probabilities are taken from Kopecky and Koreshkova (2014), appendix Computation B. We transform the four-state Markov chain into a two-state chain. We do this simply by calculating the conditional transition probabilities for the two lowest and the two highest states. The Markov transition matrix Λ_{hh} is given by expression (23), where entry in row one, column one, denotes the probability to stay in good health if one is already in good health. For the state values, we take their grid points, transforming them by taking a weighted average at the ergodic distribution. Being in bad health implies 13.48 times higher health expenditures. To preserve the deterministic profile of our model, we do not use their distribution Γ_h as initial distribution, but the ergodic distribution. Note that the initial distribution converges to the ergodic distribution fast: after 5 periods we are basically at the ergodic distribution. The ergodic distribution for the health shock in our model is $[0.63, 0.37]$, that is, 37% of households in retirement are in bad health.

$$\Lambda_{hh} = \begin{bmatrix} 0.7918 & 0.2082 \\ 0.3610 & 0.6390 \end{bmatrix} \tag{23}$$

C.4 Tax Function

Labor income tax is given by equation (24) below. This is equation (13) restated, with y denoting labor income (adjusted for productivity growth). Figure C18 shows some characteristics of that tax schedule. The top panel plots disposable against total earnings. The function is concave - high-earnings individuals contribute more in taxes, also as proportion of their income.

In fact, our formulation admits an analytical representation of marginal and average tax rates. Recall that:

$$T_y(y) = y - \lambda_0 y^{1-\lambda_1} \tag{24}$$

Marginal tax rates (MTR) are given by:

$$T_y^{'}(y) = 1 - \lambda_0(1-\lambda_1)y^{-\lambda_1}$$

Average tax rates (ATR) are given by:

$$\frac{T_y}{y} = 1 - \lambda_0 y^{-\lambda_1}$$

The bottom panel of figure C18 shows that indeed average tax rates are higher as earnings increase. The same is true for marginal tax rates. At the 50th percentile, the ATR and the MTR are respectively 12.4% and 28.1%. For the 90th percentile, those numbers are 27.0% and 40.2%.

C.5 The Distribution of Bequests

We compute the bequest incidence function, $beq(e, g)$ using the 2016 wave of the Survey of Consumer Finances. This survey asks respondents about three particular instances in which they received their largest (eventual) inheritances, together with the year in which they received them. There is also a fourth variable, corresponding to "other bequests received" that does not feature an accompanying year. We drop this fourth variable. It

Figure C18: Marginal and Average Tax rates by Earnings

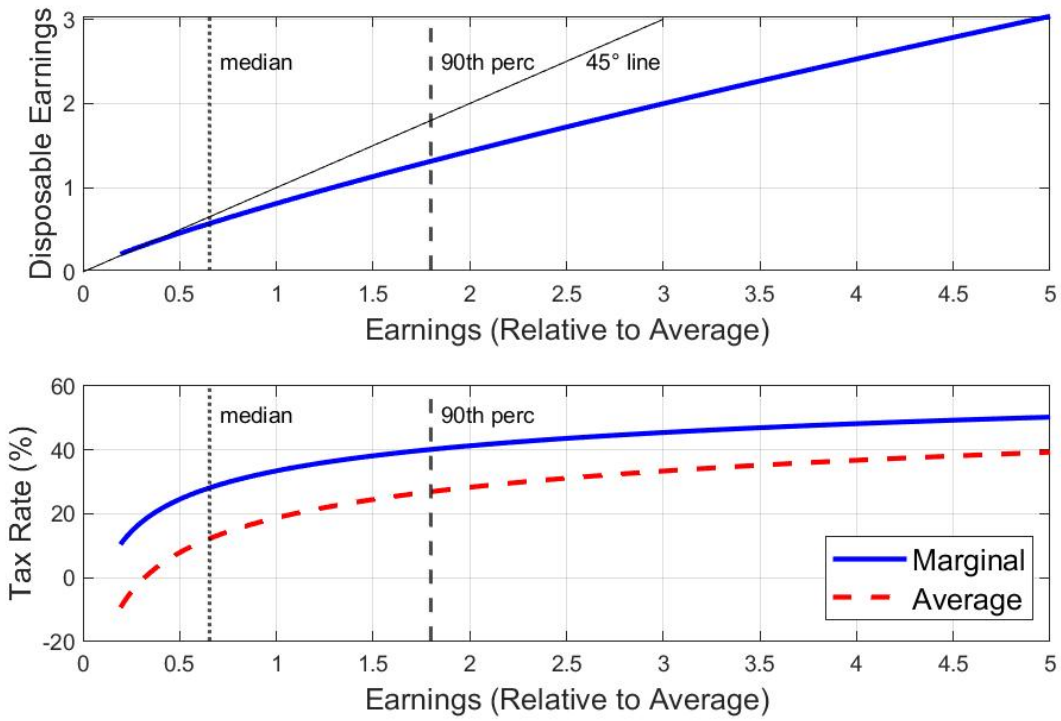

Notes: The top panel displays disposable earnings as a function of pre-tax earnings. The bottom panels displays the implied marginal (solid blue line) and average (dashed red line) tax rates. Pre-tax earnings, in the x-axis, is relative to average pre-tax earnings in the baseline calibration. "90th perc" refers to the 90th percentile.

only has positive observations for 0.6% of people, while the first variable features more than a fourth of observations receiving some inheritance.

We then construct a panel where the unit of observation is a person and the time dimension corresponds to age. For each unit, we have bequests received at each age (deflated by the corresponding CPI). Note that for each individual there are *at most* 3 observations. We then separate the sample into "top 10%" and bottom "90%" conditional on wage income, dropping individuals without any wage income in the reference year.

Our incidence function corresponds to the "average bequest received by an individual for each age", conditional on permanent income group. We assume that within permanent groups, bequests are equal. In practice, we compute averages for 10-year age groups, and then interpolate the resulting values. We then construct another incidence function, that equals the original but with age shifted by 26 periods. This is aimed at capturing spousal

(or within household unit) bequests.[39]. Finally, because around 50% of bequests are spousal (Hurd and Smith (2002)) we adjust the final incidence by constructing a linear combination with equal weights between the two incidences described above. The resulting distribution is shown in figure C19.

Figure C19: Distribution of Bequests by Age and Permanent Income

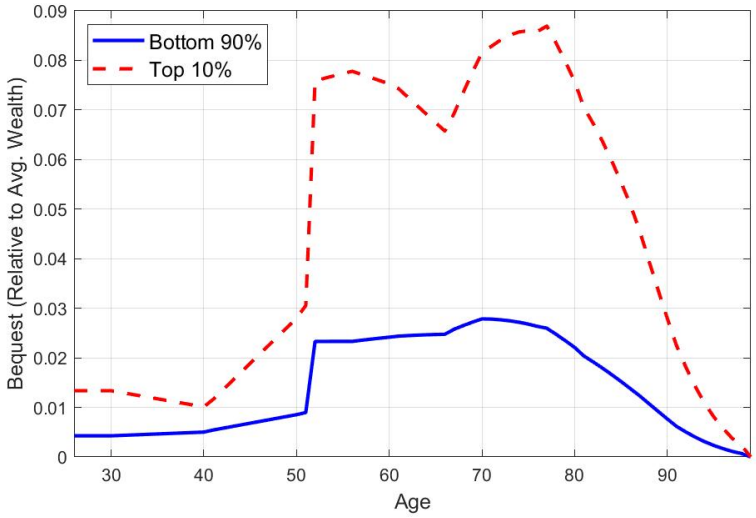

Notes: The figure shows the distribution of bequests received over the life time for different biological ages in our baseline calibration and by permanent income group. See text for further explanation.

C.6 Estimation of ϕ_{PI} on Model-Simulated Data

One of the targets in the calibration routine of section 4.7 is the elasticity of consumption out of permanent income, ϕ_{PI}. To calculate this elasticity, we closely follow the approach outlined in Straub (2019), appendix E.2.

First, we simulate productivity profiles using Monte-Carlo simulations.[40] We use our model, in particular policy functions for household choices, to derive post-tax income and consumption choices. All calculations are done at the respective steady state, usually the baseline 2015 steady state. Then, the income data are multiplied by a measurement error term $\exp(\nu)$, where ν is drawn from a normal distribution with variance 0.02. Then we employ a 2SLS estimation approach. We regress consumption $c_{i,g}$ on post-tax labor income

[39]For each age, we also re-weight this incidence by the probability of survival up until that age, to account for the probability of surviving long enough to receive that bequest.

[40]We simulate data for 7,500 model households and find the estimate of ϕ_{PI} to be reasonably robust to increasing the number of simulations.

adjusted for measurement error, $\hat{y}_{i,g}^{post}$, and age dummies:

$$\ln c_{i,g} = \alpha + \phi_{PI} \ln \hat{y}_{i,g}^{post} + d_{age} + \epsilon_{i,g}, \quad \text{for } g \in \{5, 40\}$$

We instrument for $\hat{y}_{i,g}^{post}$ using income leads, $\ln z_{i,g} = \ln \hat{y}_{i,g+j+1}^{post} - \rho \ln \hat{y}_{i,g+j}^{post}$, $j > 0$. ρ is the persistence parameter from the AR1 process of the persistent earnings shock ζ. We take into account the biennial nature of the PSID which was used by Straub (2019) to estimate ϕ_{PI} in the data, so only consider every second observation of a household's simulated data series in the 2SLS estimation.

Note that this procedure is used inside the calibration routine of section 4.7.

C.7 Social Security Function

Figure C20: Social Security Function and Replacement Rate

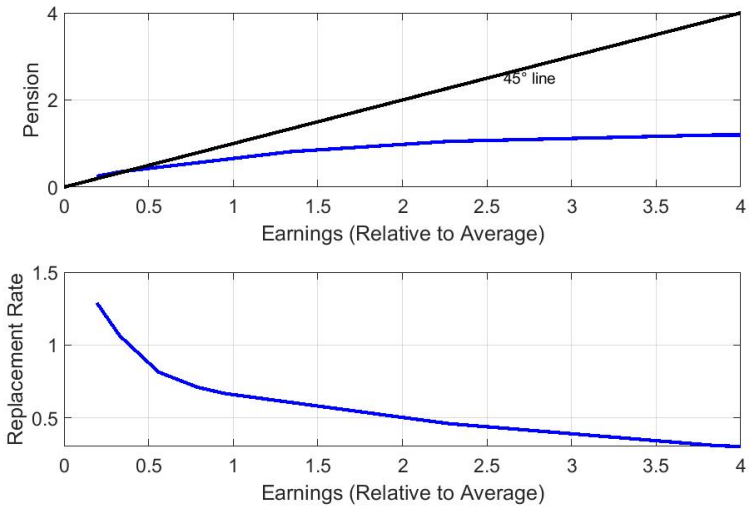

Notes: The top panel shows the pension ξ as a function of earnings at 65 years old (relative to average). The bottom panel shows the implied replacement rate.

Figure C20 shows the social security benefit payment as function of income under the baseline calibration, as well as the implied replacement rate.

The target for social security spending relative to output is set to 5.8 and taken from data supplement to Congressional Budget Office (2020), table "10. Mandatory Outlays Since 1962, as a Percentage of GDP", the sum of columns "Social Security" and "Federal Civilian and Military Retirement", 2010-2019 average.

High OOP medical expenses late in life can potentially exhaust a retiree's resources. We assume that the government provides a means-tested transfer T^M to ensure a consumption floor $\underline{c}$. In most of our exercises, it turns out that the social security system provides enough benefits and total amount of transfer T^M is tiny. Further discussion see below.

D Additional Results and Details - Section 5

D.1 The Compositional Effect of Demographics

This section briefly explains the concept behind the composition effect of aging on total savings, and consequently on the natural rate. A comprehensive study about the impact of this effect is done by Auclert et al. (2021), on which the following summary is based.

The composition effect can be described as follows. Suppose that we compute the average wealth for individuals at different ages - which is what is shown in figure 4. Now suppose that this profile is constant across different years. Then, if there is an increase in the share of older people in the population, and those people hold a relatively large amount of savings, *all else equal*, total savings in the economy will rise. This will exercise downward pressure on the interest rate.

Auclert et al. (2021) show that the effect of this change in the population composition in r^* depends crucially on not only the change in aggregate savings, but in fact in the change in the **wealth-to-earnings ratio**.The total impact of composition changes on r^* also depends on the shape of the demand for savings, which in turn depends on preference and production function parameters. For a detailed analysis, see Auclert et al. (2021), Proposition 2. Here, however, we focus our discussion on the change in wealth to income ratios implied by our model, and its match with the data.

To be clear, the *all else equal* statement does not hold in our experiments - prices, returns, and taxes change as other drivers change over time and, most importantly, the survival probabilities directly impact households' savings behavior, potentially altering the wealth-age profile depicted in 4.

Consider the following expression for aggregate wealth-to-earnings ratio, where tilde variables are normalized by effective units of labor:

$$\frac{\widetilde{\mathcal{W}}_t}{\tilde{\mathcal{Y}}_t} = \frac{\sum_g \phi_{gt} \mathbb{I}\mathbb{E} \tilde{g}^a_{gt}}{\sum_g \phi_{gt} \mathbb{I}\mathbb{E} \tilde{\mathcal{Y}}_{tg}},$$

where $\mathbb{E}\tilde{\mathcal{Y}}_{tg} \equiv \mathbb{E}_{z,h,a}\tilde{\mathcal{Y}}_{tg} = \sum_{z,h,a} \widetilde{\mathcal{W}}_t \lambda(z,h,a,g) z_{gt}$ represents the average gross earnings at each age g and $\mathbb{E}\tilde{g}^a_{gt} \equiv \mathbb{E}_{z,h,a}\tilde{g}^a_{gt} = \sum_{z,h,a} \lambda(z,h,a,g) g^a_{gt}$ represents the average end-of-period wealth holdings at the end of every period, again in effective labor terms. Note that we can deal away with the tilde variables in the ratio above:

$$\frac{\mathcal{W}_t}{\mathcal{Y}_t} = \frac{\sum_g \phi_{gt} \mathbb{E} g^a_{gt}}{\sum_g \phi_{gt} \mathbb{E}\mathcal{Y}_{tg}},$$

where, to be clear, $\mathcal{Y}_{tg} = w_t \mathbb{E}_{z,h,a} z_{gt}$. Over time, we can compute the induced change in the wealth to income ratio, conditional on savings behavior and on the life cycle earnings profile:

$$\Delta^{comp}_{t,\tau} = \frac{\frac{\mathcal{W}_t}{\mathcal{Y}_t} - \frac{\mathcal{W}_{t,\tau}}{\mathcal{Y}_{t,\tau}}}{\frac{\mathcal{W}_t}{\mathcal{Y}_t}} \tag{25}$$

for any year τ.[41] In the expression above, $\mathcal{W}_{t,\tau} \equiv \sum_g \phi_{g\tau} \mathbb{E}\tilde{g}^a_{gt}$ and $\mathcal{Y}_{t,\tau} \equiv \sum_g \phi_{g\tau} \mathbb{E}\tilde{\mathcal{Y}}_{tg}$ respectively determine the changes in aggregate savings and earnings due to the change in the composition of the population, represented by the vector of ϕ's.

It turns out that we can measure Δ^{comp} in the model and in the data, and evaluate the quality of the match. For the population composition, we use data from the United Nations (2022*b*); for earnings, we use the Luxembourg Income Studies - that is, the average earnings at each age in the model and in the data are the same by construction; the empirical wealth profile is obtained from the Survey of Consumer Finances (2016) - figure 4 compares it with the model-obtained one.

We then set the base year $t = 2015$ and plot $\Delta^{comp}_{t,\tau}$ for different years (τ). Figure D21 plots the comparison. The top panel plots a small variation of equation (25), keeping the average earnings term constant at 2015 levels. That is, the top panel plots $\frac{\frac{\mathcal{W}_t}{\mathcal{Y}_t} - \frac{\mathcal{W}_{t,\tau}}{\mathcal{Y}_{t,t}}}{\frac{\mathcal{W}_t}{\mathcal{Y}_t}}$, isolating the impact of changes in aggregate savings only. The bottom panel displays the total change across years.

The model captures the changes in aggregate wealth to income ratio remarkably well *throughout all years in the analysis*. This shows that the model is able to capture the total impact of those changes throughout our analysis period, not only to the past, but also towards the future. The quality of the match depends crucially on the model fit of the wealth-age profile depicted in figure 4.

[41] Auclert et al. (2021) approximate the formula with logs.

Figure D21: Compositional Effect of Demographics - Model and Data

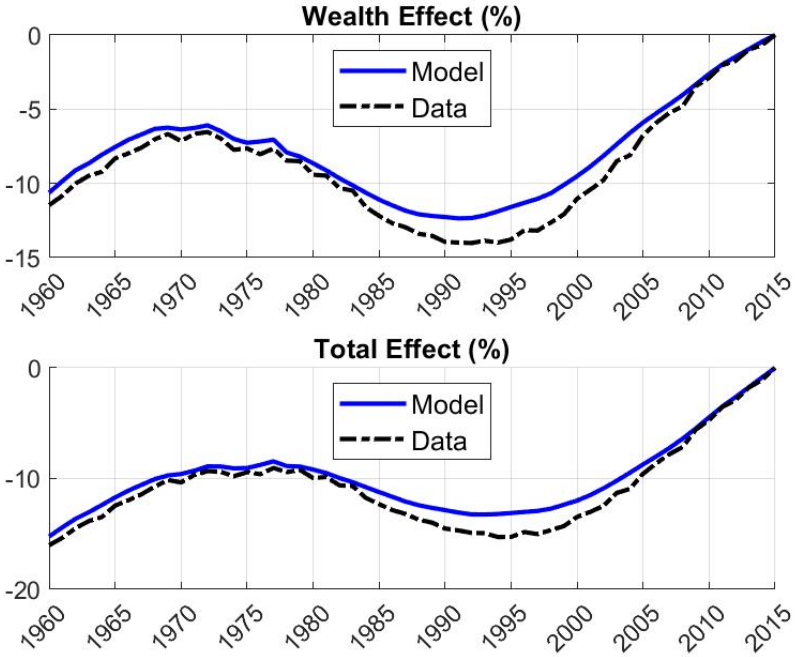

Notes: Sources (data): LIS, SCF 2016, and United Nations (2022*b*); and (model) baseline calibration. The top panel displays a variation of equation (25), keeping the average earnings term constant at 2015 levels. The bottom panel displays equation (25) evaluated from $\tau = 1960$ to $\tau = 2015$.

D.2 Details - Drivers

Table D7 displays the values the drivers assume at each decade. To represent those drivers that consist of a vector of parameters – population shares and survival probabilities – we include the probability of a 26-year old living up to 100 years, as well as the dependency ratio.

D.3 Decomposition Results

Table D8 displays the full decomposition discussed in section 5.

Table D7: Drivers at Different Decades

Driver	1965	1975	1985	1995	2005	2015
Population Growth (n_t) (%)	1.29	1.08	1.07	1.30	0.98	0.77
$\mathbb{P}$ 26 to 99 (%)	0.3	0.6	1.0	1.2	1.3	2.2
(Model) Dependency Ratio	0.18	0.19	0.20	0.21	0.21	0.23
Top 10% (%)	25.4	26.5	29.7	32.8	34.4	34.8
Top 1% (%)	5.2	5.6	7.5	9.7	11.2	11.0
Variance of Earnings Shock	0.025	0.024	0.028	0.034	0.037	0.040
$100 \times \gamma$	2.47	1.74	1.02	1.61	1.51	0.70
Debt-to-GDP (%)	37.3	25.9	33.7	44.7	36.6	72.4
G-to-GDP (%)	16.4	16.9	16.0	15.1	15.2	14.9
NFA-to-GDP (%)	4.3	5.8	4.1	-5.5	-17.9	-37.4
Labor Share (%)	0.63	0.63	0.61	0.61	0.62	0.59
OOP Relative Price	0.43	0.48	0.56	0.75	0.87	1.00

Note: $\mathbb{P}$ 26 to 99 (%) refers to the probability of surviving from 26 to 99 years old, corresponding to the ages represented in the model. The model-based dependency ratio consists of the ratio of individual aged 66 or older, relative to individuals aged 26 to 65. Top 10% and top 1% refer to pre-tax earnings shares.

Table D8: Baseline Steady-State Decomposition

Driver	1965	1975	1985	1995	2005
Demographics	0.72	0.50	0.51	0.55	0.30
Inequality	0.96	0.87	0.54	0.20	0.05
Productivity	1.44	0.75	0.13	0.61	0.54
Labor Share	0.60	0.65	0.24	0.29	0.30
Medical Exp.	0.21	0.19	0.16	0.10	0.05
Government Debt	-0.29	-0.49	-0.39	-0.24	-0.34
Government Consumption	0.16	0.23	0.12	0.02	0.03
NFA	0.36	0.37	0.36	0.27	0.15
Interactions	0.10	0.07	-0.06	-0.03	0.01
Total	4.26	3.14	1.61	1.76	1.10

Notes: Each row-column pair shows the marginal contribution of the row driver in the column year, in percentage points, relative to 2015.

D.4 An Alternative Way to Perform the Steady-State Decomposition

We now consider a slightly different decomposition exercise than the one shown in Section 5. The approach in that section is as follows. First, to compute the total change from 2015 to 2005, we alter the values of all underlying drivers to the respective level in 2005, one by one, and solve the model. Then, to compute the change from 2015 to 1995, we proceed in two steps. First, we change the values of each driver, one by one, to their 1995 level, but keep all other drivers at their 2005 level. The result is a decomposition of the change from 2005 to 1995. The second step is then to add the resulting contributions to the ones computed from 2015 to 2005. We repeat this process for other decades into the past.

Consider now a different approach: keep all drivers at their 2015 levels, and then change only one driver to its assumed decade, repeating the process for each decade. Importantly, always keep all other drivers at their 2015 level.

Figure D22 displays the results. By construction, the total change in r^* is the same. Comparing the contributions across drivers, the takeaways are unchanged. The only notable difference is the role of residuals, i.e. interactions across variables. This indicates that the drivers' marginal contributions are state-dependent: in Figure D22, when we keep all other drivers at their 2015, the total sum of marginal contributions is smaller than the total contribution, as opposed to Figure 5. This indicates the importance of having a unified framework to capture changes in r^* over time, in which the evolution of one driver can affect the marginal impact of the others.

D.5 Decomposition of Demographics - With Population Growth

Figure D23 shows the decomposition of demographics, including population growth n. Note that, in this exercise, population growth has no impact whatsoever on population shares ϕ_g. Its impacts are restricted to the evolution of debt per effective unit of labor ($\tilde{B}$) and associated taxes, as well as affecting the evolution of capital per effective unit of labor $\tilde{K}$, which explains its irrelevance. Note that this result does not in any way mean that population growth is not an important driver r^*. Instead, it just means that in practice its effects are entirely due to its impact on the population shares ϕ_g.

Figure D22: Alternative Steady-State Decomposition

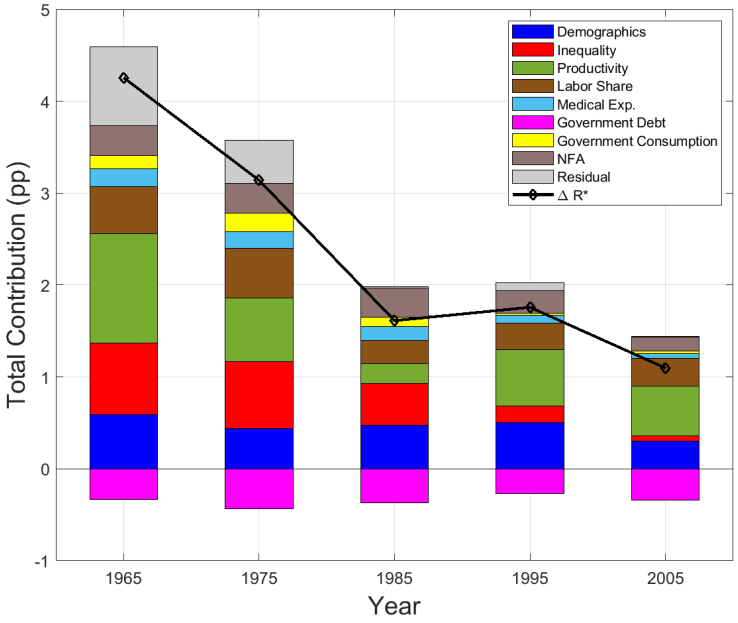

Notes: The figure shows the decomposition of change, in percentage points, into various drivers for decades since 1965. The decomposition method is described in Appendix Section D.4. black diamonds show total change in r^* at respective years relative to 2015. The stacked bars show the contribution to the change relative to 2015 from various drivers.

Figure D23: Demographics Decomposition Including Population Growth

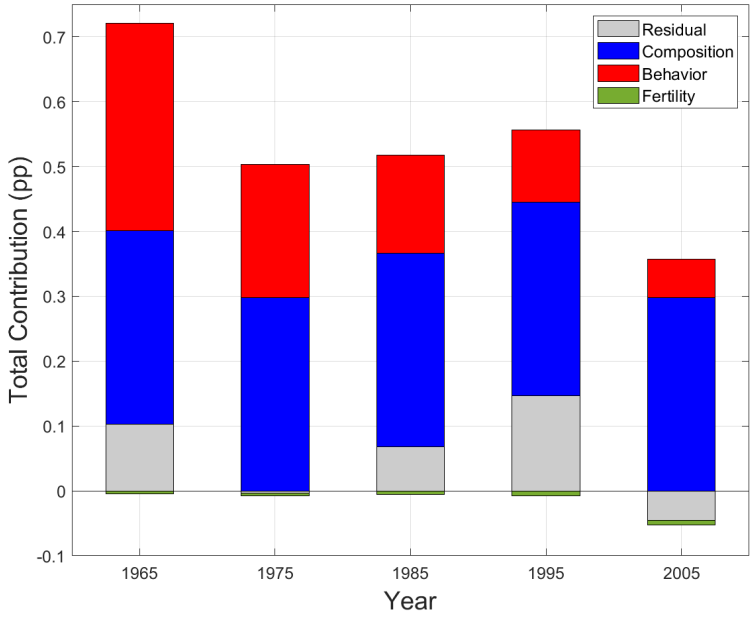

Notes: The figure shows the decomposition of change, in percentage points, into the sub-components of demographics, including population growth (conditional on age-specific population shares) for decades since 1965. The stacked bars show the contribution to the change relative to 2015.

D.6 Additional Results - Homothetic Economy

Figure D24 and table D9 reproduce the baseline results with the homothetic economy, as explained in section 5. Figures D25 and D26 reproduce the decompositions for demographics and inequality respectively.

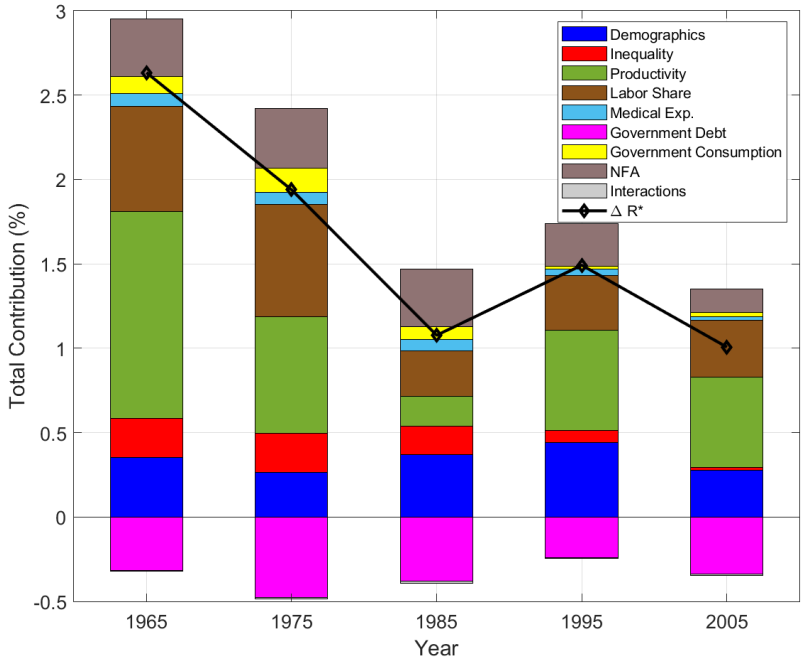

Figure D24: Change in r^* and decomposition of change into various drivers for decades since 1965, homothetic utility. The black diamonds show total change in r^* at respective years relative to 2015. The stacked bars show the contribution to the change relative to 2015 from various drivers.

Table D9: Decomposition of Results - Homothetic economy.

Driver	1965	1975	1985	1995	2005
Demographics	0.36	0.27	0.37	0.44	0.28
Inequality	0.23	0.23	0.17	0.07	0.02
Productivity	1.22	0.69	0.18	0.60	0.53
Labor Share	0.62	0.67	0.27	0.32	0.34
Medical Exp.	0.08	0.07	0.07	0.04	0.02
Government Debt	-0.32	-0.47	-0.38	-0.24	-0.34
Government Consumption	0.10	0.14	0.08	0.01	0.02
NFA	0.34	0.36	0.34	0.25	0.14
Interactions	-0.00	-0.01	-0.01	-0.00	-0.01
Total	2.63	1.94	1.08	1.49	1.01

Notes: Each row-column pair shows the marginal contribution of the row driver in the column year, in percentage points, relative to 2015, for the homothetic economy.

Figure D25: Decomposition of Changes in r^* - Demographics (Homothetic Economy)

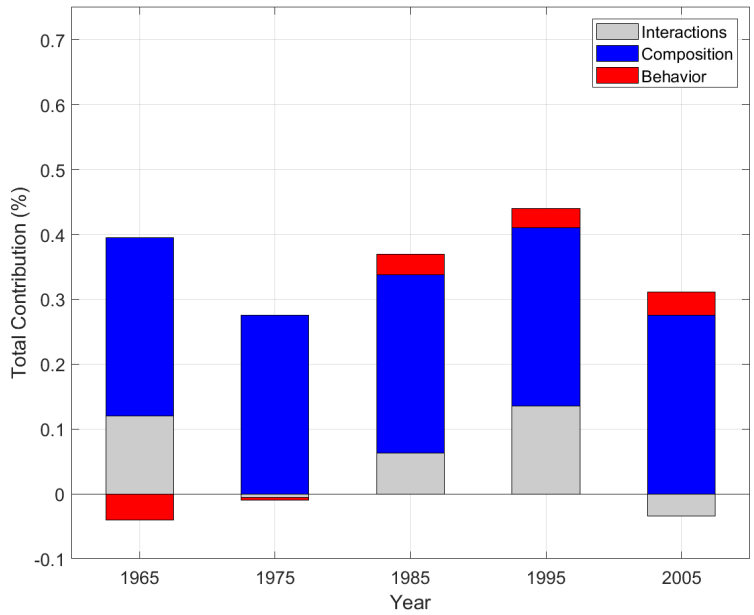

Notes: The stacked bars show the contribution, in percentage points, to the change relative to 2015 from sub-components of demographic change in the homothetic economy. The stacked bars show the contribution to the change relative to 2015.

Figure D26: Decomposition of Changes in r^* - Inequality (Homothetic Economy)

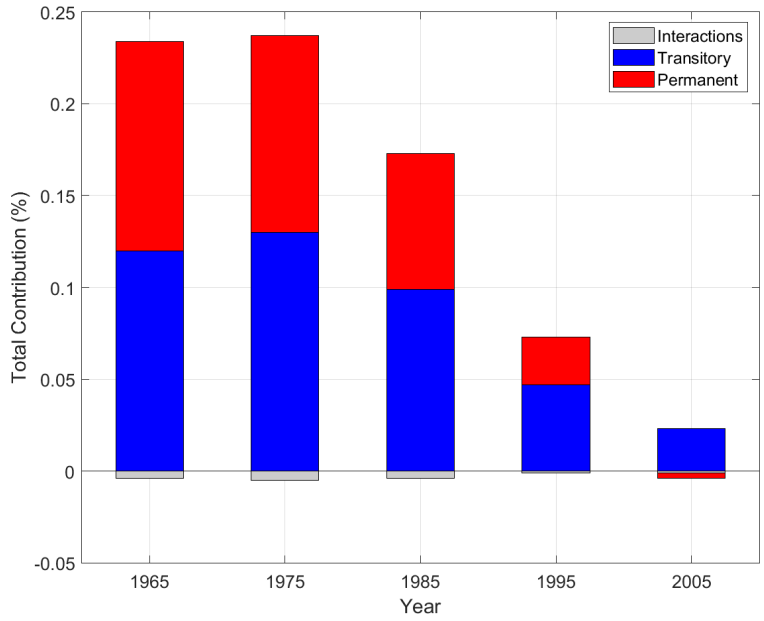

Notes: The figure shows the decomposition of change, in percentage points, into the sub-components of inequality for decades since 1965 in the homothetic economy. The stacked bars show the contribution to the change relative to 2015

E Additional Details and Results - Section 6

E.1 Transition Path - Drivers

In all our exercises in section 6, with the exception of government debt and demographic variables, we assume that drivers converge to their terminal values in forty years. Specifically, we proceed as follows: for each driver, we assume that there is a linear path, beginning in the year where the latest data we used is available and ending forty years later.[42]. For government debt, we try to remain close to the CBO estimate and assume that it rises linearly from 2019 to 2080 and stabilizes thereafter.

For demographic drivers (population growth, age shares, and survival probabilities) we use realized values up until 2022 and UN projections thereafter, as explained in section 6. For each driver, including each component of the vectors (e.g. $\phi_{g,t}$), after joining the past and future (projected series), we take 10-year moving averages.[43]

Figure E27: Time Series for Drivers - Baseline Scenario Analysis

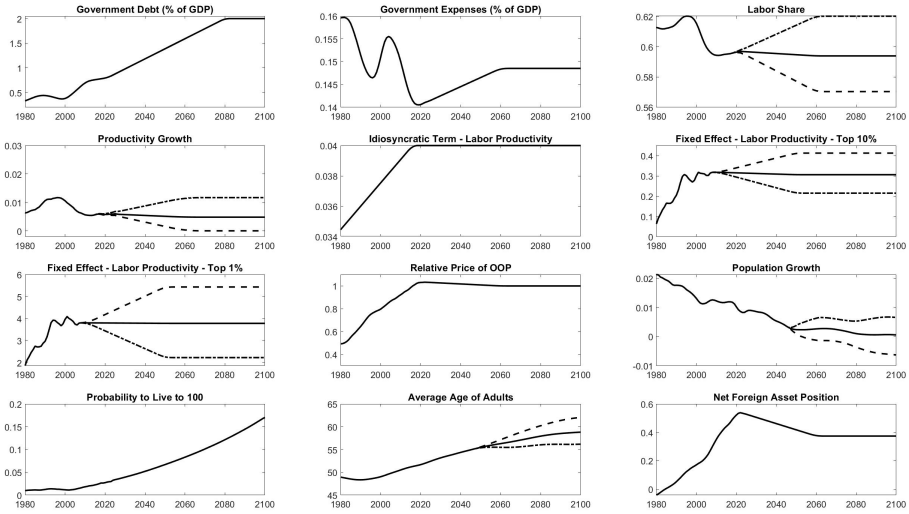

Notes: The figure shows the evolution of distinct drivers from 1980 to 2100 used in our baseline scenario analysis. All drivers stabilize at their terminal values after 2100, but simulations are truncated in 2200. The dashed line corresponds to the "high" scenario, while the dash-dotted line represents the "low" scenario.

Table E10 summarizes the assumptions regarding the value of the standard deviation of each driver in our baseline probabilistic analysis.

[42]With the exception of top earnings shares, whose latest available year in Piketty and Saez (2003) 2011. the latest year is 2019.

[43]For TFP, by far our most volatile series, we use a 20-year moving average.

Table E10: Probabilistic Analysis - Standard Deviations

Driver	Standard Deviation
Demographics	UN Projections
TFP	0.0029
Labor Share	0.0125
Top 1% Share	0.0123
Top 10% Share	0.0172

Notes:Values correspond to half the change in each driver from 1995 to 2015. Draws for terminal values of each driver in our probabilistic analysis are based on the displayed standard deviation, with the latter two drivers being perfectly correlated.

E.2 Statistics - Stochastic Simulation with Demographics Only

Figure E28 shows the distribution of r^* at different years as well as the distribution of the trough year for the stochastic simulation with only Demographics.

Figure E28: Selected Statistics - Distribution of r^* - Simulation with only Demographics

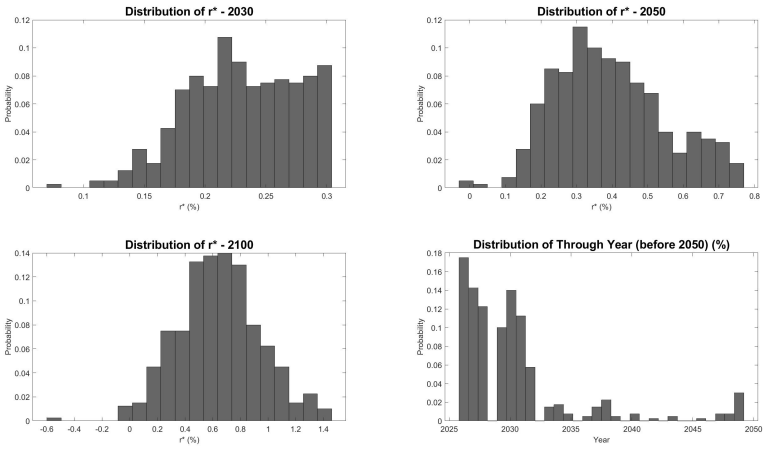

Notes: The two top and the bottom-left panels represent the distribution of r^* from the Monte Carlo simulations. The bottom-right panel displays the distribution of the trough year, provided that it occurs before 2050. Only draws from future Demographics distributions are considered. Government Debt-to-GDP converges to 200% in 2082.

E.3 Statistics - Stochastic Simulation with Four Main Drivers

Figure E29 shows the distribution of r^* at different years as well as the distribution of the trough year for the stochastic simulation with the four main drivers.

Figure E29: Selected Statistics - Distribution of r^* - Simulation with All Drivers

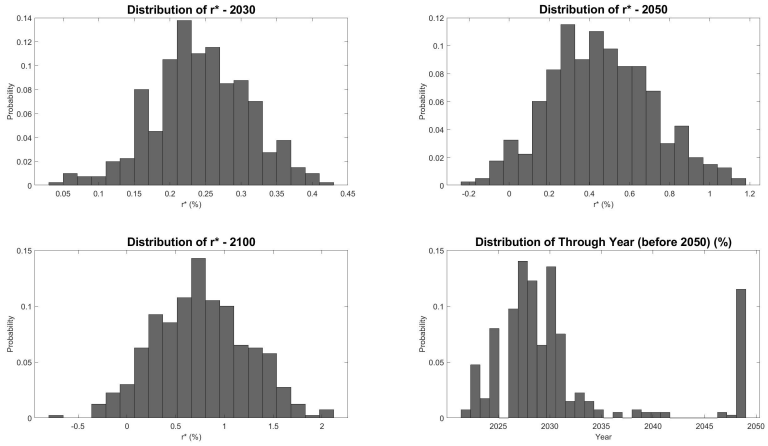

Notes: The two top and the bottom-left panels represent the distribution of r^* from the Monte Carlo simulations. The bottom-right panel displays the distribution of the trough year, provided that it occurs before 2050. Government Debt-to-GDP converges to 200% in 2082.

E.4 Decomposition of the Future of r^* - Demographics

Figure E30 below decomposes the future impact of demographics. Specifically, the dashed line shows the implied path of r^* if survival probabilities p_g where kept constant from 2022 onward, while in the dash-dotted the population shares ϕ_g are kept fixed at their 2022 levels. Note that, under perfect foresight, the path before 2022 can also be affected.

Figure E30: Decomposition of r^* Path - Demographics Components

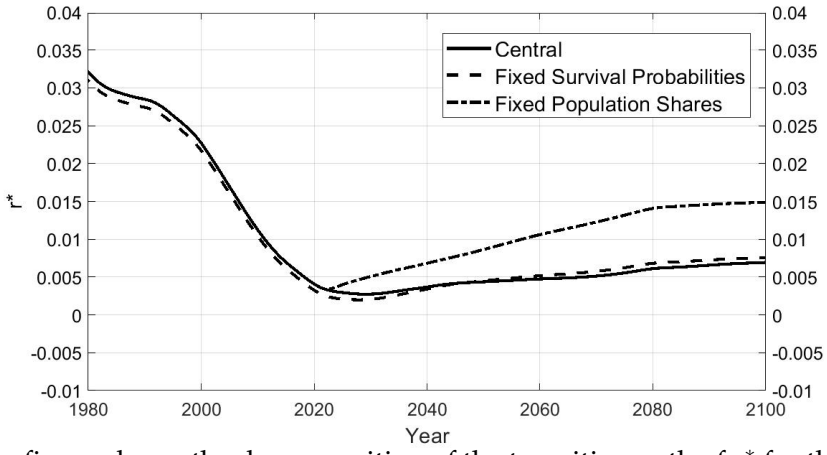

Notes: The figure shows the decomposition of the transition path of r^* for the central scenario, as well as the central scenario but fixing population shares (dash-dotted line) and fixing survival probabilities (dashed line).

E.5 The Future of r^* with Debt-to-Output Converging to 100%

Figures E32, E33, and E34 reproduce our results in section 6 but with Debt-to-Output converging to 100% (instead of 200%). Figure E31 shows the evolution of the drivers in the scenario analysis.

Figure E31: Time Series for Drivers - Scenario Analysis with $\frac{B}{Y} = 100\%$.

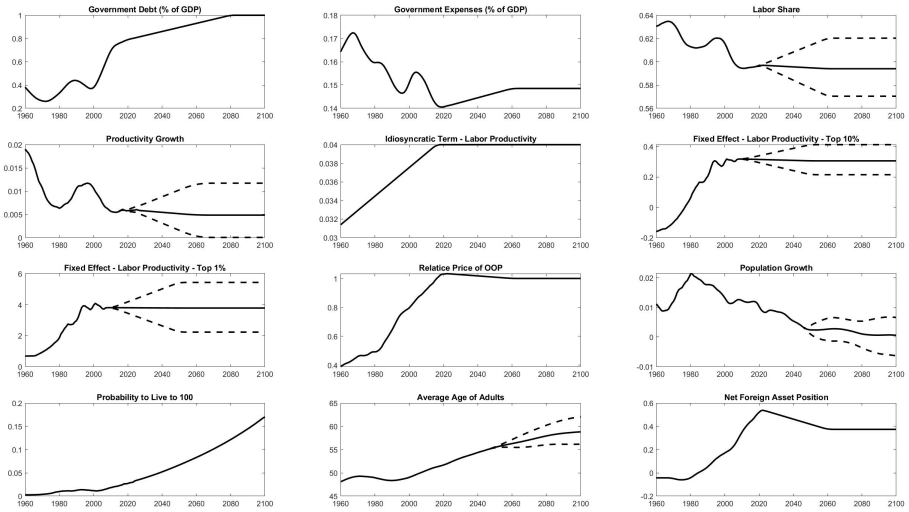

Notes: The figure shows the evolution of distinct drivers from 1960 to 2100 used in the scenario analysis with Debt-to-GDP stabilizing at 100%. All drivers stabilize at their terminal values after 2100, but simulations are truncated in 2200. The dashed line corresponds to the "high" scenario, while the dash-dotted line represents the "low" scenario.

Figure E32: Scenario Analysis with Debt-to-GDP Stabilizing at 100%

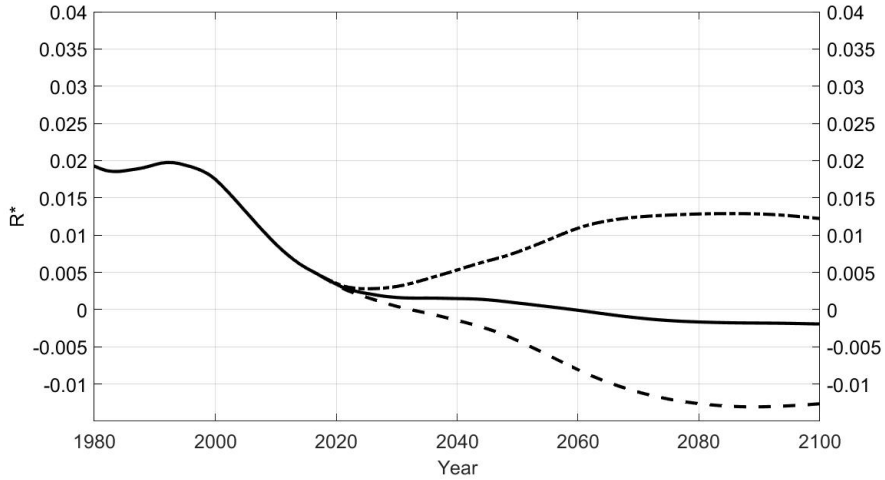

Notes: The figure shows the evolution of r^* into the future based on different scenarios for the evolution of the drivers. The solid, dashed, and dot-dashed line represent respectively the central, low, and high scenarios. The terminal value for Debt-to-GDP is 100%.

Figure E33: Probabilistic Analysis with Demographics Only and Debt-to-GDP Stabilizing at 100%

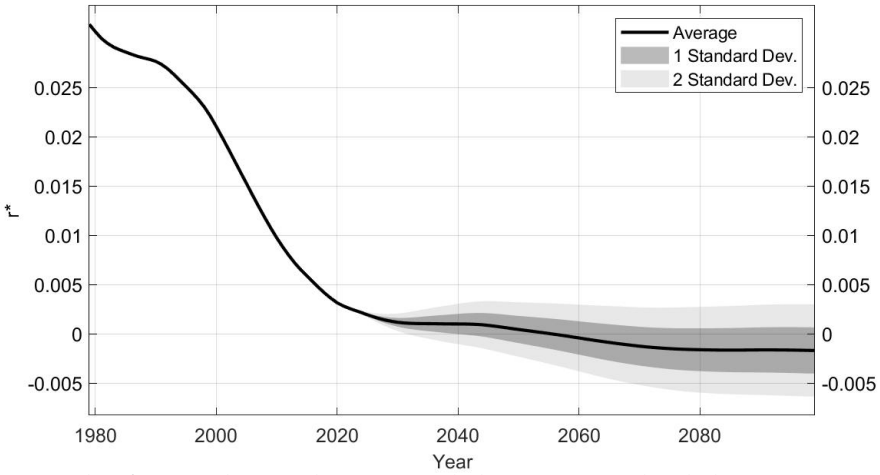

Notes: The figure shows the one and the two standard deviation confidence intervals for the Monte Carlo simulations of the path of r^* only due to Demographics. Government Debt-to-GDP converges to 100% in 2080.

Figure E34: Selected Statistics - Distribution of r^* - Simulation with only Demographics and Debt-to-GDP Stabilizing at 100%

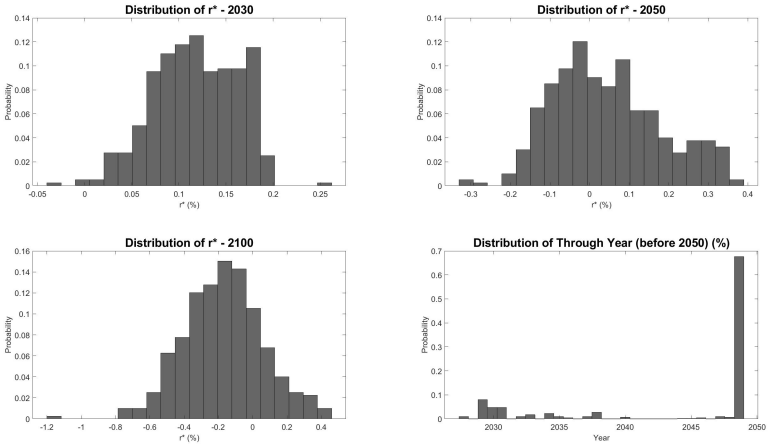

Notes: Notes: The two top and the bottom-left panels represent the distribution of r^* from the Monte Carlo simulations. The bottom-right panel displays the distribution of the trough year, provided that it occurs before 2050. Only draws from future Demographics distributions are considered. Government Debt-to-GDP converges to 100% in 2082

F Further Results on Policy

Table F11 shows the impact of a 10 percentage point change in each of the policy instruments considered,

Table F11: Response of r^* and Earnings Taxes to 10 pp Change in Policy Parameters

Policy Parameter/Instrument	Baseline Value	Total Change (pp)	$\frac{\Delta r^*}{\Delta x}$	$\frac{\Delta Tinc}{\Delta x}$
Consumption Tax (%)	5.0	-0.26	-2.56	-0.38
Profit Tax (%)	25.0	0.06	0.61	-0.14
Capital Income Tax (%)	30.0	0.19	1.88	0.03
Estate Tax (%)	10.0	-0.06	-0.61	-0.06
Progressivity of Earnings Tax	0.18	0.56	5.56	-0.04
Basic Income (% of avg. Income)	0.0	0.65	6.53	0.76
Public Debt-to-GDP (%)	72.4	0.10	1.01	-0.03
G (% of GDP)	14.9	1.08	10.85	0.72
OOP Subsidy (%)	0.0	0.05	0.51	0.00
Replacement Rate Parameter	1.43	0.11	1.44	0.02

Notes: Responses at baseline 2015 steady state. Column 1 lists the policy parameters; column 2 lists their value at the baseline 2015 steady state. Values in column 3 denote response of r^* in basis points to change in parameter by 1 percentage point (pp). Column 4 shows response of labor income tax receipts relative to output in percentage points to change in parameter by 1 pp.

G The Burden of Debt

In this section, we show the results concerning the prediction of the evolution of the natural rate, net of productivity and population growth. This is commonly referred to "r minus g", that is the cost of rolling over the debt net of the growth rate of the economy. This is a key indicator of government debt sustainability. In our economy, debt is non-defautable, and thus our analysis is necessarily incomplete. This is why we relegate these results to the appendix. Yet we think that our results provide some guidance as to whether the current and future projected values of government debt are sustainable. As is shown below in figure G35, our results indicate that the cost of rolling over the debt will remain close to the growth rate of the economy. In this figure, g represents the growth rate of the economy, given by $\frac{Y_t}{Y_{t-1}} - 1$.

In either scenario, the difference remains negative until circa 2065, indicating that the increasing level of debt - which, recall, goes up to 200% of GDP in the terminal steady state - might not be an issue for the cost of debt servicing. Looking across scenarios, the forces

Figure G35: $r^* - g$ - Scenario Analysis

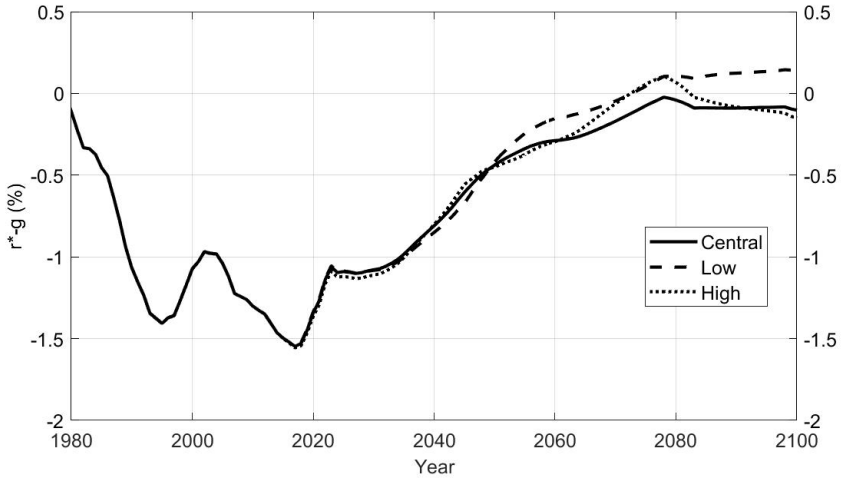

Notes: The lines show the evolution of r^* minus g in each scenario of our baseline analysis in section 6. g is defined as the growth rate of of output, i.e. $g_t = \frac{Y_t}{Y_{t-1}} - 1$

seem to even each other out. That is, in the pessimistic scenario the decline in n and γ is accompanied by a similar decline in r^*, and the opposite is true for the optimistic case. The reason is that, in the bad scenario, low population and productivity growth pull the natural rate downwards, and so the net effect on $r^* - g$ is muted. The reverse movements happen in the good scenario, with the net effect also muted.

However, when we consider our full stochastic simulation (with demographics, inequality, productivity, and the labor share), shown in figure G36, the picture is somewhat different. The reason is that, for this simulation, population and productivity growth do not necessarily co-move as they do in the scenario analysis. The simulation suggests that there is some uncertainty in $r^* - g$ towards the future, especially over the very long run. There is, however, a very high probability that it stays in negative territory for the short- and medium-run. For instance, the simulations attribute less than 5% probability of $r^* - g$ being greater than 0 in 2040.

In all, our results suggest that, at the range considered, the cost of sustaining the government debt is low. In addition, for the short- and medium-run, it is very unlikely that $r^* - g$ will remain substantially above 0.

Figure G36: $r^* - g$ - Probabilistic Analysis

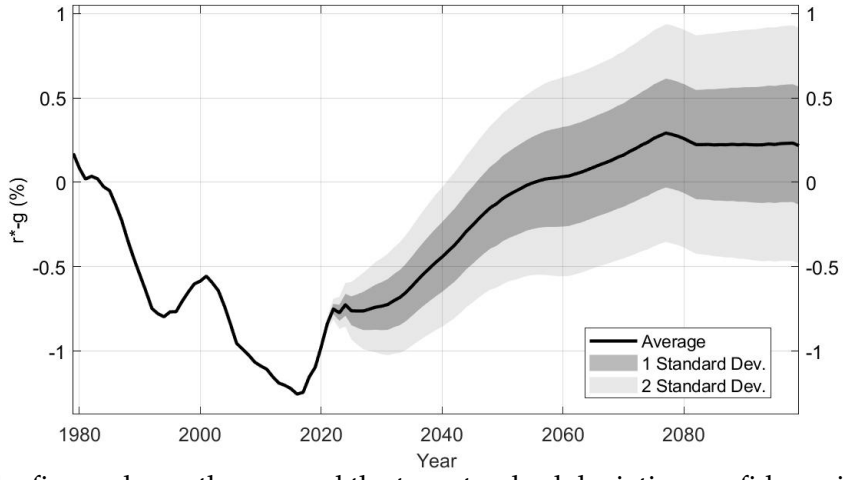

Notes: The figure shows the one and the two standard deviation confidence intervals for the Monte Carlo simulations of the path of $r^* - g$ due to changes Demographics, Productivity, Permanent Inequality, and the Labor Share. Government Debt-to-GDP converges to 200% in 2080.

H Details - Sensitivity

Table H12 below displays the full set of sensitivity results

Table H15: Decomposition (decline relative to 2015), high elasticity target. (elasticity = 26.8)

Driver	1965	1975	1985	1995	2005
Demographics	0.58	0.43	0.43	0.45	0.24
Inequality	0.57	0.52	0.35	0.14	0.04
Productivity	1.62	0.87	0.19	0.73	0.65
Labor Share	0.69	0.75	0.28	0.34	0.36
Medical Exp.	0.14	0.12	0.11	0.06	0.03
Government Debt	-0.28	-0.44	-0.35	-0.22	-0.31
Government Consumption	0.13	0.18	0.09	0.02	0.03
NFA	0.28	0.29	0.28	0.21	0.12
Residual	0.10	0.07	-0.04	-0.02	0.01
Total	3.81	2.78	1.34	1.70	1.16

Notes: Each row-column pair shows the marginal contribution of the row driver in the column year, in percentage points, relative to 2015.

Table H12: Sensitivity steady state - Different Calibrations

$\bar{\sigma}$	9.5	6	3.1	2.0
elast savings wrt r^*	9.7	11.8	18.3	26.8
Demographics	0.55	0.65	0.72	0.58
Inequality	2.24	1.69	0.96	0.57
Productivity	1.09	1.21	1.44	1.62
Labor Share	0.44	0.48	0.60	0.69
Medical Exp.	0.42	0.33	0.21	0.14
Government Debt	-0.24	-0.28	-0.29	-0.28
Government Consumption	0.22	0.21	0.16	0.13
NFA	0.49	0.45	0.36	0.28
Residual	0.05	0.10	0.10	0.10
Total	5.26	4.84	4.26	3.81

Notes: Baseline decomposition of changes (reproducing table D8, column 4) with different assumptions regarding the elasticity of savings with respect to r^*. In each exercise, a (re-)calibration is performed.

Figure H37 below display the scenario exercise, considering the high-elasticity calibration (sensitivity). In this case, there is in practice no noticeable change from our baseline predictions.

Figure H37: Scenario Analysis - Calibration with elasticity = 26.8.

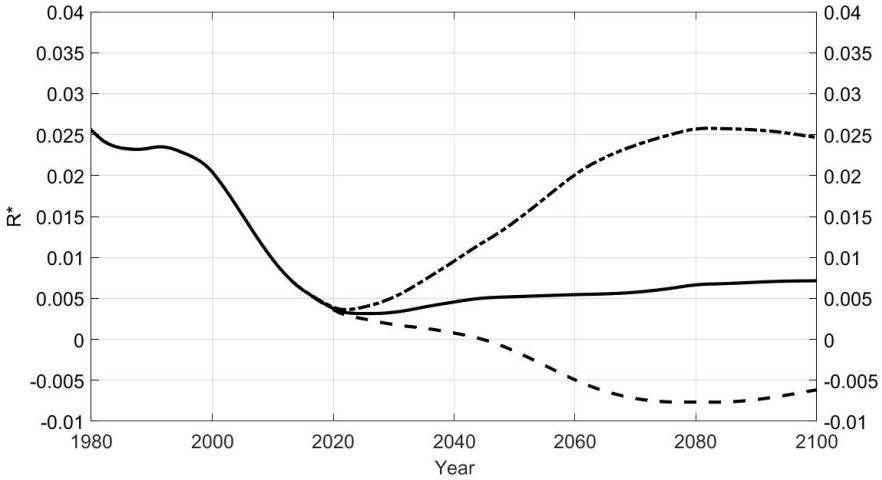

Notes: The figure shows the evolution of r^* into the future based on different scenarios for the evolution of the drivers, with the calibration selecting the elasticity of savings with respect to returns to equal 26.8. The solid, dashed, and dot-dashed line represent respectively the central, low, and high scenarios. See text for further explanations.

Table H13: Decomposition (decline relative to 2015), very low elasticity target (elasticity = 9.7)

Driver	1965	1975	1985	1995	2005
Demographics	0.55	0.34	0.49	0.61	0.34
Inequality	2.24	2.10	1.16	0.38	0.11
Productivity	1.09	0.55	0.01	0.40	0.35
Labor Share	0.44	0.47	0.15	0.19	0.20
Medical Exp.	0.42	0.36	0.31	0.18	0.09
Government Debt	-0.24	-0.60	-0.45	-0.27	-0.39
Government Consumption	0.22	0.34	0.17	0.03	0.04
NFA	0.49	0.53	0.50	0.37	0.21
Residual	0.05	0.02	-0.08	-0.02	0.00
Total	5.26	4.12	2.27	1.87	0.96

Notes: Each row-column pair shows the marginal contribution of the row driver in the column year, in percentage points, relative to 2015.

Table H14: Decomposition (decline relative to 2015), low elasticity target (elasticity = 11.8)

Driver	1965	1975	1985	1995	2005
Demographics	0.65	0.41	0.51	0.60	0.34
Inequality	1.69	1.56	0.89	0.31	0.08
Productivity	1.21	0.62	0.05	0.47	0.42
Labor Share	0.48	0.52	0.17	0.21	0.23
Medical Exp.	0.33	0.28	0.25	0.15	0.07
Government Debt	-0.28	-0.56	-0.43	-0.26	-0.37
Government Consumption	0.21	0.31	0.16	0.03	0.04
NFA	0.45	0.48	0.46	0.34	0.19
Residual	0.10	0.06	-0.08	-0.03	0.01
Total	4.84	3.67	1.99	1.81	1.00

Notes: Each row-column pair shows the marginal contribution of the row driver in the column year, in percentage points, relative to 2015.

Figure H38 below display the scenario exercise, considering the high-elasticity calibration (sensitivity). In this case, although the upper and lower scenario paths are very similar to the baseline, the central prediction projects a slightly higher natural rate over time. This is due to the increased sensitivity of r^* to government debt implied by the new calibration target. Yet the takeaway of our baseline scenario analysis is robust (quantitatively) to different alternatives regarding the elasticity to savings with respect to returns.

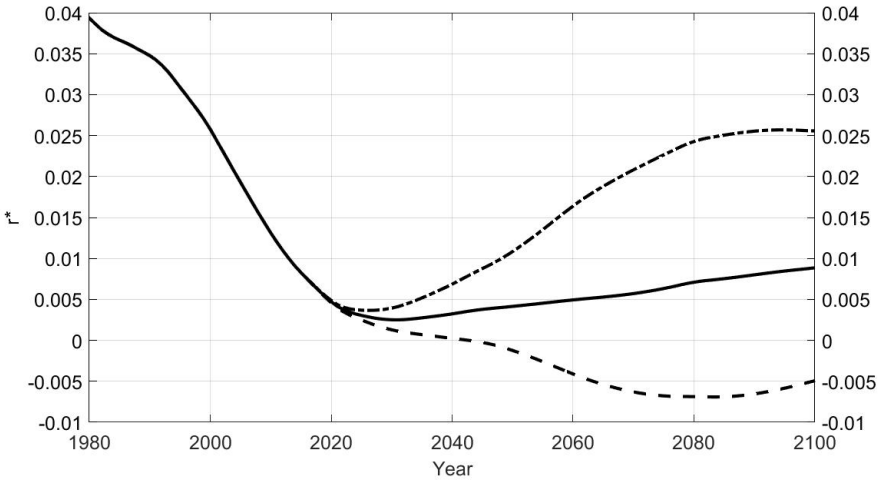

Figure H38: Scenario Analysis - Calibration with elasticity $= 11.8$.

Notes: The figure shows the evolution of r^* into the future based on different scenarios for the evolution of the drivers, with the calibration selecting the elasticity of savings with respect to returns to equal 11.8. The solid, dashed, and dot-dashed line represent respectively the central, low, and high scenarios. See text for further explanations.