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

Josef Platzer and Marcel Peruffo

WP/22/30

IMF Working Papers describe research in progress by the author(s) and are published to elicit comments and to encourage debate.

The views expressed in IMF Working Papers are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

**2022
FEB**

IMF Working Paper
Research Department

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

Prepared by Josef Platzer and Marcel Peruffo*

Authorized for distribution by Prachi Mishra
February 2022

IMF Working Papers describe research in progress by the author(s) and are published to elicit comments and to encourage debate. The views expressed in IMF Working Papers are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

ABSTRACT: We develop a heterogeneous agent, overlapping generations model with nonhomothetic preferences that nests several explanations for the decline in the natural rate of interest (r^*) suggested in the literature: demographic change, a slowdown in productivity growth, a rise in income inequality, and public policy. The model can account for a 2.2 percentage point (pp) decline in r^* between 1975 and 2015, which is within the range of empirical estimates. Rising income inequality is an important driver (-0.70 pp), and together with demographic change (-0.71 pp) and the slowdown in productivity growth (-1.0 pp) explains most of the decline. Growing public debt is the major counteracting force (+0.31 pp). Permanent income inequality is of greater importance than inequality due to uninsurable income risk, and matching the degree of nonhomotheticity in consumption and savings behavior to empirical estimates is essential for this result. We predict that r^* will reach a low of 0.38% by 2030, after which a slow reversal will begin. The natural rate will stabilize at 1% in the long run, a low level when compared with the postwar path of r^* implied by the model. This remains true even if we take into account soaring public debt levels due to the COVID-19 pandemic. Policy can have considerable impact on the level of r^* through the tax and transfer system.

JEL Classification Numbers:	D31, E21, E43, E60
Keywords:	Demographic Change, Inequality, Life-cycle, Natural Rate of Interest, Nonhomothetic Preferences, Secular Stagnation
Author's E-Mail Address:	jplatzer@imf.org ; marcel_peruffo@brown.edu

---

*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, Balazs Zelity, as well as conference and seminar participants at Brown University, CEA, CEF, CES, FRB Boston, Federal Reserve Board, IFABS, IMF, SEA, National University of Singapore, Bank of Norway, Peking University School of Economics, and Wharton Macro Finance workshop for helpful comments and discussions. Josef Platzer kindly acknowledges the generous financial support of the James M. and Cathleen D. Stone Wealth and Income Inequality Project.

Contents

1	Introduction	7
2	Literature Review	10
3	Model	12
3.1	Demographics	13
3.2	Households	13
3.3	The Consumer's Problem	14
3.3.1	The Worker	14
3.3.2	The Retiree	14
3.4	Productivity Process	15
3.5	Production	16
3.6	Government	16
3.7	Market Clearing and Equilibrium	17
4	Functional-Form Assumptions and Calibration	17
4.1	Utility Functions	18
4.1.1	Consumption and Leisure Utility	18
4.1.2	Bequest Utility	20
4.2	Demographics	20
4.3	Productivity Process	21
4.4	Production	21
4.5	Medical Expenses	22
4.6	Government	22
4.7	Numerical Calibration Procedure	23
5	Numerical Analysis	27
5.1	Steady-State Analysis	27
5.1.1	Baseline Experiment	27
5.1.2	Homothetic-Economy Results	29
5.1.3	Inequality Due to Permanent Productivity versus Income Risk	31

5.2	Transition-Path Analysis	32
5.2.1	Scenario Analysis	34
5.3	Policy Analysis	36
5.4	Sensitivity	40
6	Conclusion	41
A	Appendix	42
A.1	Data on Interest Rates in U.S.	42
A.2	Appendix to Sections 3 and 4	42
A.3	Market Clearing Conditions	54
A.4	Definition of Competitive Equilibrium	55
A.5	Additional Material for Section 4.7	55
A.6	Homothetic Economy	56
A.7	First-order Conditions	57
A.8	Extension with Exogenous International Capital Flows	57
A.8.1	Foreign Capital Flows: Numerical Analysis	59
A.9	Decomposing Demographic Change	59
A.10	The Role of Out-of-Pocket Medical Expenses	61
A.11	Productivity Growth and Nonhomothetic Consumption Preferences	63
A.12	Additional Tables and Figures	66
A.12.1	Steady State	66
A.12.2	Transition Path	69
A.12.3	Scenario Analysis	73
A.12.4	Policy Analysis	78
A.13	Further Results on Policy	80
A.14	Overview Related Literature	81
A.15	Algorithm	82

Figures

1	Profile of σ_g under baseline calibration	19
2	Transition path of r^* and top 10% labor-income share	33
3	Transition path r^* implied by model under baseline calibration and under “ALL low” and “ALL high” scenarios	36
4	Response of r^* to change in ratio of public debt-to-output $\frac{B}{Y}$ at the steady state	38
5	Response of r^* to a change in replacement-rate parameter ϕ	39
6	Various interest rates and trend of short-term interest rate	42
7	Population growth rate and probability to survive until age 99 conditional on age 26 along the transition path	43
8	Deterministic productivity profile $h_z^g(e_i)$ by permanent productivity type	45
9	Transition path of e_H and σ_e	46
10	Profile of out-of-pocket (OOP) health expenditure under bad health state	48
11	Income tax system and social security benefit payment as function of income	49
12	Public debt along the transition path, in percent of output	51
13	Effective replacement rate implied by social security system	52
14	Share of profits going to each age group	53
15	Transition path for shock variables under baseline transition-path experiment	53
16	U.S. Net International Investment Position (NIIP) from 1976 to 2020, in percent of GDP	58
17	Response of r^* to a change in means-tested transfer T^M	63
18	Lorenz curves from model in baseline 2015 steady state and data.	66
19	Life-cycle profiles in baseline 2015 steady state, by permanent productivity type	67
20	Various life-cycle asset profiles in baseline 2015 steady state, model and data	68
21	Transition path r^* model, and average and range of estimates of r^* from literature	69
22	Transition path r^* model and trend of short-term interest rate	69
23	Transition path for selected variables under baseline	70
24	Transition path for selected variables under baseline	71
25	Transition path for selected variables under baseline	72
26	Transition path for population growth rate under three different scenarios	73

27	Transition path for public debt under three different scenarios	74
28	Transition path for earnings process parameters e_H (left) and σ_e (right) under three different scenarios: baseline	74
29	Transition path for TFP growth rate γ under three different scenarios	75
30	Implied top 10% labor income share under different scenarios: baseline	75
31	Implied path for exogenous government spending G to output (left) and total tax revenues to output (right) under different scenarios	76
32	Implied path for r^* under different scenarios for path of population growth rate after year 2020: baseline, low, and high	76
33	Implied path for r^* under different scenarios for path of public debt after year 2020: baseline, low, and high	77
34	Implied path for r^* under different scenarios for path of earnings process parameters e_H and σ_e after year 2020	77
35	Implied path for r^* under different scenarios for path of TFP growth rate γ after year 2020	78
36	Response of labor income tax receipts T_{inc}/Y and output in response to change in coverage of OOP medical expenses	79
37	Response of r^* to change in mandatory retirement age $\bar{g}$ at the steady state	79
38	Marginal tax rate (MTR) and average tax rate (ATR) as a function of pre-tax labor income wzl at baseline 2015 steady state	81

Tables

1	Calibration results for Type III parameters Note: OOP stands for out-of-pocket	25
2	Parameters under baseline calibration (Type I and Type II)	26
3	Selected data for variables in 1975 and 2015	27
4	Change in r^* going from 1975 steady state to 2015 steady state, and decomposition in baseline model	28
5	Change in r^* going from 1975 to 2015 steady state, and decomposition in three versions of the model	30
6	Selected variables going from 1975 to 2015 steady state in three versions of the model	30
7	Change in r^* and selected statistics going from 1975 to 2015 steady state, and decomposition by inequality source in baseline model	31
8	Parameters in terminal steady state of transition path for three different scenarios: “baseline,” “low,” and “high”	34
9	r^* in terminal steady state under various scenarios for driving variables	35

10	Response of selected variables to change in policy parameter at baseline 2015 steady state	37
11	Change in r^* , and decomposition under various parametrizations of the model	40
12	Change in r^* going from 1975 to 2015 steady state and decomposition in model with NIIP	59
13	Change in r^* going from 1975 to 2015 steady state due to demographic change and sub-components	60
14	Change in r^* going from 1975 to 2015 steady state due to demographic change and sub-components	61
15	Change in r^* going from 1975 to 2015 steady state, and decomposition in baseline model and model without out-of-pocket (OOP) medical expenses	61
16	Change in r^* going from 1975 to 2015 steady state, and selected decomposition in baseline versions of the model and under assumption of constant o_t term	64
17	Elasticity of r^* in response to change in policy parameter at baseline 2015 steady state . .	80
18	Overview of change in r^* across several papers in literature	82

1 Introduction

There is a growing consensus among economists that the natural rate of interest (r^*) has declined in recent decades in the United States and other advanced economies. This development is of great importance for monetary and fiscal policy and for explaining economic trends such as movements in asset prices. The recent COVID-19-induced downturn highlights the challenges for countercyclical policy in a low-rate environment and the enduring relevance of the topic. A good understanding of the drivers of the decline in r^* and of whether it is of temporary or permanent nature is key for devising an appropriate policy response.

The quantitative literature on the decline in r^* initially focused on channels such as demographic change, a slowdown in productivity growth, and public policy variables – particularly, transfers and public debt levels (Gagnon, Johannsen and Lopez-Salido, 2016; Eggertsson, Mehrotra and Robbins, 2019). The analysis was conducted in an overlapping generations (OLG) setting, without intragenerational heterogeneity. Over the years, the literature has been extended to account for other drivers, among them inequality. Several papers (Auclert and Rognlie, 2018; Straub, 2019) have addressed the question of inequality's contribution to the fall in r^* in models geared toward the analysis of inequality and found the contribution to be sizeable. While we have a good understanding of the most promising candidates to explain the decline in r^* , we lack a clear picture of their relative importance. We do not have a comprehensive study of the drivers of the fall in r^* that is conducted in a unified framework.

This paper tries to fill this gap. Our main contribution is to systematically explore various explanations for the decline in r^* within one unifying framework. To do this, we build a quantitative model that includes the most common drivers proposed in the literature as well as some previously unexplored ones. The model provides an environment to conduct a horse race between competing explanations for the decline in r^* . In addition we use our framework to clarify open questions in the literature, among them whether permanent productivity differences or idiosyncratic income risk drives inequality's contribution, what the role of productivity growth is in models that deviate from homotheticity, and whether seniors' out-of-pocket health expenditures are of importance for the determination of r^* . The path of the natural rate in coming decades is of high relevance for policy makers. We use our model to project r^* into the future. Finally, we ask whether policy can have a meaningful effect on the natural rate and what prominent reform proposals imply for the future trajectory of r^* .

The main exercise is split into a steady-state comparison and a transition-path analysis. In both we change the variables pertaining to the candidate explanations for the r^* decline – we frequently refer to them as drivers – and look at the implied impact on r^* . In the first exercise, we compare a steady state in which each driver assumes its observed values around 2015 with a counterfactual economy with drivers at their 1975 levels.¹ While this experiment is clear cut and computationally relatively simple, our

---

¹The model is calibrated to match key moments of the data in 2015. We then change the candidate drivers to their 1975 value to analyze the change in r^* . The 1970s are a common reference point in the r^* literature, as it is the decade when the baby boom generation entered the labor force. It is also the time when other often-discussed macroeconomic developments began, such as the

assumption that the economy is in a steady state at the respective dates is questionable. Therefore, as a second experiment, we analyze the transition path between two steady states. We start with a post-World War II economy in 1950. At this time, households with perfect foresight are made aware of a projection of the path of driving variables to 2060, beyond which it is difficult to find high-quality forecasts of the drivers we study. The transition-path analysis allows us to project the trend of r^* into the future.

The paper combines a heterogeneous-agent framework with an OLG setting. It features seventy-four generations and endogenous labor supply. The model includes means-tested government support for medical expenditures and a social security system as well as a variety of progressive and proportional taxes and transfers, mimicking the US system. Agents are exposed to idiosyncratic productivity shocks and health-expenditure shocks in old age. They also face differences in permanent income. In a seminal paper, [Straub \(2019\)](#) shows that permanent income inequality can affect r^* if the model allows for nonhomothetic consumption and savings behavior. Addilog preferences and a bequest motive with a luxury-good component are important sources of nonhomotheticity in our model.

Comparing a 1975 steady state with a 2015 steady state, the model explains a 2.2 percentage point (pp) decline in r^* . The slowdown in total factor productivity (TFP) growth is the most important driver under the baseline calibration, with a contribution of -1.0 pp, or 46.2% of the total decline. Demographic factors, namely a decline in fertility and a rise in life-expectancy, and inequality are the second and third most important drivers, each explaining about a 0.7 pp decline, or 32.5%. Our findings thus suggest that inequality should be considered a major driving force of the r^* decline, next to the slowdown in TFP growth and demographic change. The most important counteracting factor is the increase in public debt by about 50 percent of output between the steady states, which increased r^* by 0.31 pp, or 14.3%.

The transition path of r^* shows a mild decline between 1970 and 2000, after which the downward trend accelerates considerably. The natural rate stands at 0.76% in 2020 according to the model. The projection forward suggests that r^* continues to decline for another decade but hits a low of 0.38% in 2031 and starts a slow reversal thereafter. This is driven by the large baby boom generation's depleting of their savings in old age and ultimately dropping out of the model and also driven by a steady increase in the forecast level of public debt ([Congressional Budget Office, 2019](#)). However, even though public debt rises to 200% in our terminal steady state, r^* settles at 1% over the long run. This is a low level in the context of our model, which covers the post-World War II period with initial levels of r^* as high as 4%.

This projection naturally depends on assumptions about the future path of driving variables. To study sensitivity, we construct low and high scenarios for the most important drivers: demographic variables, public debt, inequality drivers, and TFP growth. A result from this analysis is that we can't rule out a permanently negative natural-rate environment. Under a scenario in which all forces that put downward pressure on r^* coincide – arguably a low-probability event but still illustrative – the natural rate goes negative around 2035 and settles at a low of -0.92% after another twenty years. On the flipside, the high-end scenario, which includes public debt levels soaring to 350% of GDP, predicts a rise of r^* to 3%.

---

productivity slowdown and rise in inequality.

Equipped with our model, we want to clarify several open issues in the r^* literature. The first of them is the role of inequality. Our model features two prominent mechanisms that can generate intragenerational inequality: permanent productivity differences and idiosyncratic income risk.² For permanent productivity differences to matter, we ensure that the model features a degree of nonhomotheticity in consumption and savings as estimated in the data. We find that the rise in income inequality is important for the r^* decline. Differences in permanent productivity explain almost all of the effect: in the steady-state comparison, more than 80% of inequality's contribution is due to larger permanent productivity differences. In our model, neither is idiosyncratic income risk able to explain much of the change in inequality, measured by the top 10% labor income share, nor does it have a large effect on r^* .

The nonhomothetic environment in our model deserves further attention. It is one of the model features regarding which existing contributions in the literature vary considerably. Homotheticity implies that the response of consumption to permanent changes in income is independent of income itself. As a result, changes in inequality driven by permanent income have essentially no effect on the aggregate supply of savings. (Permanent income is distinct from temporary income shocks, which eventually reverse.) The data, however, reject the homotheticity (see [Straub \(2019\)](#)). In our nonhomothetic model, richer households have higher savings rates, so higher inequality can lead to more aggregate savings and consequently depress the natural rate. Matching the degree of nonhomotheticity implied by estimates is key for the results. If we change the model to a fully homothetic version, the role of inequality is extremely muted; it accounts for only 0.03 pp of the total decline.

Our model features several policy variables, ranging from tax rates, direct transfers, and the social security system to the level of public debt. We find that policy can potentially have a large effect on the natural rate. This leads us to propose that r^* should be treated as a *policy choice*. While our analysis in this regard is incomplete – as we can't analyze the full costs and benefits associated with changes in the natural rate in our purely real model – we can conduct a positive analysis. We quantify the effect of policy changes on both r^* and output, which highlights a possible trade-off between raising the natural rate and reducing the level of output.

A particularly stark but illustrative example of policy's power to affect r^* is a move from the current pay-as-you-go social security system to a fully funded system. In the latter system, households would accumulate a lot of savings before retirement. The implication is that r^* would drop by more than 2 pp from its 2015 steady-state value into deeply negative territory (-1.6%). Introducing a health care system that covers most OOP health expenditures would raise the natural rate by about 40 basis points (bp). Tax policy is another way to affect r^* . A rise in the consumption tax or the profit tax would lower the natural rate, while increasing the capital income tax, estate tax, or progressivity of the labor income tax would raise the natural rate. The government could thus adjust the tax system in a revenue-neutral way

---

²[Straub \(2019\)](#) shows the effect of changes in permanent productivity differences on r^* in a model that also contains idiosyncratic earnings risk, but he does not show the effects of changing both sources of inequality together – neither the effect on inequality nor the effect on the natural rate. [Auclert and Rognlie \(2018\)](#) and [Rachel and Summers \(2019\)](#) show the effect of changing earnings risk on r^* , but neither model has both permanent productivity differences together with nonhomothetic consumption preferences.

to change r^* . Increasing public debt by 1 pp would raise the natural rate by 0.64 bp, with an associated output cost of 1.2 bp. An increase of 10% in the debt ratio, which corresponds to the long-term impact of COVID-19 on the debt ratio according to recent estimates,³ would raise the steady-state value of the natural rate by 6.4 bp, at a steady-state output cost of 12 bp.

The two papers closest to this study are [Rachel and Summers \(2019\)](#) and [Straub \(2019\)](#). In [Rachel and Summers \(2019\)](#), the decline in r^* is the central research question and many of the drivers analyzed in this study are included. However, there are two crucial differences. First, they conduct the analysis in two separate stylized models. One has no intragenerational inequality but has three overlapping generations. The other one has idiosyncratic earnings risk but no OLG structure. Second, they do not include the mechanism by which permanent productivity differences together with nonhomothetic consumption and savings behavior can lead to rising inequality and downward pressure on r^* . This mechanism drives most of the results related to inequality in our model.

[Straub \(2019\)](#) highlights the importance of permanent income differences together with nonhomothetic consumption and savings behavior that can generate the empirically observed positive relation between savings rates and permanent income. The crucial difference from this paper is our focus on the natural rate. We attempt to quantify the contribution of several drivers, not just permanent income differences, on r^* , to project r^* forward, and to analyze the impact of policy.

The structure of the paper is as follows. Section 2 gives a more detailed literature review. Section 3 develops the model, and section 4 is dedicated to functional-form assumptions and model calibration. Section 5 shows the results, and section 6 concludes.

2 Literature Review

The early papers trying to explain the decline of r^* within general-equilibrium models found demographic change and a slowdown in productivity growth to be the main driving forces. [Gagnon, Johannsen and Lopez-Salido \(2016\)](#) explores the impact of demographic dynamics in the United States using a life-cycle model with a large number of generations.⁴ [Eggertsson, Mehrotra and Robbins \(2019\)](#), henceforth EMR, develops a quantitative model to explore drivers of low interest rates and the possibility of real rates permanently below zero. Besides considering demographics and TFP growth, they consider several other possible determinants of r^* , most importantly the change in public debt. EMR discuss the possible effect of rising inequality on the natural rate in a stylized model but do not add it to the full quantitative analysis. Our paper is very much in the same spirit as these two papers. The biggest difference is that we add a within-age-group heterogeneity dimension. We find that rising inequality due to changes in within-age-group earnings is as important a force as demographic change or the slowdown in TFP growth. Another contribution is that our forecast horizon for the future trajectory of r^* is longer.

³This number is from [Congressional Budget Office \(2020\)](#).

⁴Table 18 in appendix A.14 gives an overview of estimates of the r^* decline in our model and others in the literature.

A more recent paper within the same literature is [Rachel and Summers \(2019\)](#). Their first main contribution is to show that the bloc of industrial economies can be reasonably considered as a closed economy, thereby ruling out capital flows as a possible driver for them. Second, they explore in detail the counteracting effect on r^* of public policy changes. They use two separate models to make their point: one with three generations to capture demographic dynamics but without intragenerational inequality, and another one abstracting from age but including idiosyncratic income risk. Our contribution here is to study inequality and demographics within a single model. We consider this preferable from a methodological standpoint, as it ensures that possible interaction effects are captured and allows us to have consistent testing conditions – in particular, the same model environment and parametrization.⁵ We also add an additional channel for inequality: permanent productivity differences together with nonhomothetic consumption and savings behavior. We find that this channel drives inequality's contribution. We include a broader set of policy variables, which allows us to compare the effect of a wider range of policies on r^* .

[Auclert et al. \(2020\)](#) provide important insights on the effect of demographic change on wealth-to-income ratios and r^* . They conduct a shift-share analysis of demographic change, and derive conditions under which the pure compositional effect of demographic change, i.e. a metric that only relies on data, is sufficient to derive the impact on r^* . In addition, they highlight the importance of elasticities of asset demand and supply and shed light on quantitative differences in r^* estimates in the literature.⁶ Our paper focuses on the demographic dimension, so one of our contributions is our consideration of other drivers. A distinction in our model is that we explore the role of nonhomothetic consumption preferences.

Several papers studying the macroeconomic consequences of inequality have considered the impact of its rise on the natural rate. [Straub \(2019\)](#) provides important theoretical and empirical insights on permanent income differences and nonhomothetic consumption and savings behavior. Permanent income can be thought of as a fixed effect in an earnings regression and is distinct from a temporary-shock term that generates income fluctuations. He shows that the data reject the assumption of homotheticity in consumption and develops a quantitative model to analyze the implied macroeconomic consequences. [Straub \(2019\)](#) finds that rising permanent income differences can explain at least a 1 pp decline in the natural rate. Our modeling choices regarding consumption nonhomotheticity draw heavily on this finding. The main difference is our focus on the natural rate. We are interested in quantifying the effect of various drivers, not just permanent income differences, on r^* and in forecasting the effect of policy. The differences in the models are motivated by our focus on r^* as well: we account for the substantial changes in factor-income shares in recent decades and their possible impact on r^* . Endogenous labor introduces a disincentivizing effect of higher income taxation, an interesting addition in light of our policy results. [Auclert and Rognlie \(2018\)](#) is another paper highlighting the possible implications of rising income inequality, but this time with a focus on inequality due to idiosyncratic earnings risk. The decline

⁵For example, [Rachel and Summers \(2019\)](#), p. 37, table 7 simply add the contribution of drivers from their two models when they calculate the total r^* decline. While interaction effects are small in our model, we find that magnitudes can be greatly affected moving from one model environment to another.

⁶[Ho \(2019\)](#) points out that the effect of demographic change on r^* is weakly identified in OLG models.

in r^* is not their primary research question, so their model doesn't feature some common elements in the literature, such as demographic change. They find that rising inequality over recent decades can account for a decline in r^* of at least 70 bp.

In [Mian, Straub and Sufi \(2021\)](#) the authors examine the relative importance of demographic change versus rising income inequality on saving behavior and consequently r^* . Their evidence suggests that rising income inequality is the more relevant factor explaining the decline in the natural rate. They conduct a shift-share analysis, while we rely on a more theory-driven approach, developing a quantitative model. We can isolate individual channels and study the impact on r^* dynamics, and consider additional drivers like productivity growth and public policy.

Additional drivers for the r^* decline have been suggested, and papers have studied other countries than the United States.⁷ [Carvalho, Ferrero and Nechio \(2016\)](#) and [Jones \(2018\)](#) are additional contributions focusing on demographic factors. [Papetti \(2019\)](#) and [Ikeda and Saito \(2014\)](#) look at Europe and Japan, respectively. Following the [Bernanke \(2005\)](#) global-savings-glut hypothesis, several papers explore the international dimension – particularly, the role of capital flows – and find it to be important ([Coeurdacier, Guibaud and Jin, 2015](#); [Lisack, Sajedi and Thwaites, 2019](#); [Barany, Coeurdacier and Guibaud, 2018](#)).⁸ [Caballero and Farhi \(2014\)](#) discusses the role of scarcity of safe assets. [Sajedi and Thwaites \(2016\)](#) focuses on the contribution of a decline in the relative price of capital goods. [Kopecky and Taylor \(2020\)](#) looks at changes in risk premia induced by demographic change and the implications for the natural rate.

While we do not include all of these channels in our model, we do consider our integration of inequality, nonhomothetic consumption and savings behavior, and OOP health expenses an important contribution for understanding the r^* decline. We leave further elaborations for future research.

3 Model

In this section we outline the model. Functional-form assumptions and calibration are discussed in the next section. Our modeling choices are meant to capture the most important aspects of equilibrium determination in savings markets in a closed economy. On the savings-supply side, our model features households in an OLG structure that are subject to uninsurable idiosyncratic risk and permanent productivity differences. Utility functions allow for nonhomothetic consumption and savings behavior. Households' savings decision is driven by life-cycle, precautionary, and bequest considerations. On the demand side, there is a representative producer that requires savings in order to invest in capital, and there is a central government that issues debt and taxes its citizens in order to fund its activities, which include government consumption and a social security system. In what follows, we describe each of these features. Further details are relegated to appendix [A.2](#).

---

⁷[Abel \(2003\)](#), [Geanakoplos, Magill and Quinzii \(2004\)](#), and [Krueger and Ludwig \(2007\)](#) point out the importance of demographic factors for interest rate trends and were written before the Great Recession.

⁸We discuss an extension of our baseline model with exogenous international capital flows in appendix [A.8](#).

3.1 Demographics

We model a closed economy in which time is discrete ($t = 0, 1, 2, \dots$) and there are G^d overlapping generations at any time. Households face uncertainty about their length of life. A household 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 $1 - p_{g,t}$. There is a predetermined oldest generation ($g = G^d$) for which the probability of dying is 1 – that is, $p_{G^d,t} = 0$.

In every period, individuals are born, and $\tilde{n}_t$ denotes the fertility rate. The growth rate of the population N_t is given by n_t such that $N_{t+1} = (1 + n_t)N_t$. Note that given survival probabilities $\{p_{g,t}\}$, a given $\tilde{n}_t$ implies a certain n_t and vice versa. Demographic change can be caused by an increase or decline in $\tilde{n}$, a change in gross survival probabilities $\{p_{g,t}\}$, or both, where the latter can account for migration flows.⁹

3.2 Households

An individual household is indexed by i , but we often suppress this notation for the sake of readability. We interpret an individual (or agent) i as a household. Mortality rates and other household-specific variables are not gender specific. Households have offspring, but we do not track the ancestry across generations, and all agents in the model are adults.¹⁰

Individual preferences are represented by the following time-separable utility function:¹¹

$$\mathbb{E} \sum_{g=1}^G \beta^{g-1} s_{g-1} \left(p_{g-1} u_g(c, l) + (1 - p_{g-1}) v_g(a) \right) \quad (1)$$

Here β is the discount factor and s_g the probability of surviving from birth until age g .¹² Agents derive instantaneous utility from consumption c and disutility from supplying labor l according to function $u(\cdot)$. Upon death, assets a are transferred to surviving generations. There is a warm-glow bequest motive represented by $v(\cdot)$ that gives an incentive to leave a bequest. The instantaneous-utility functions may depend on age g for reasons explained below.

Agents in this economy move through two stages of life: work and retirement. Working individuals are of age $g < \bar{g}$ and endogenously supply their labor to firms in exchange for wages. Retired individuals do not work, and they obtain earnings from their savings and from the social security system. Individuals deterministically enter retirement at age $g = \bar{g}$.

Each agent i is endowed with an individual productivity state $z_{i,g,t}$ and a health state h . We cover both in more detail in the following subsections.

⁹The law of large numbers implies that $p_{g,t}$ is the share of households of age group g at t that survive into the next period. When we allow for migration, this share has to be adjusted for migration flows, and we denote the migration-adjusted shares by $\bar{p}_{g,t}$. See appendix A.2 for further details on our treatment of migration flows.

¹⁰There are OLG models that explicitly account for consumption of dependents – for example, in Gagnon, Johannsen and Lopez-Salido (2016). Their results, though, are insensitive to this assumption.

¹¹Time subscripts suppressed for readability.

¹²It holds that $s_g = \prod_{i=1}^{g-1} p_i$.

There are two assets in this economy. Households can save by holding government bonds b and capital k . We assume assets are held in proportion to their aggregate supply. Over any perfect-foresight path, an arbitrage condition ensures the returns on these two assets are equalized, such that households must be indifferent between holding bonds and capital.

The individuals' state variables include their current productivity realization (z), current health status (h), current asset holdings (a), and current age (g).

3.3 The Consumer's Problem

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 $t + 1$ gross their return.

We describe the consumer's problem separately for workers and retirees. Companies' pricing decisions are frictionless, so all quantities are real.

3.3.1 The Worker

The worker's problem in recursive form is:¹³

$$V_g(z, h, a) = \max_{a', l} u_g(c, l) + \beta p_g \mathbb{E}_{z', h'} V'_{g+1}(z', h', a') + \beta(1 - p_g) v_{g+1}(a') \quad (2)$$

The worker's budget constraint is:

$$(1 + \tau_c)c + a' = (1 + (1 - \tau_k)r) [a + beq] + wzl - T_y(wzl) + T^W + (1 - \tau_d)d \quad (3)$$

Here $a = b + k$ denotes assets (sum of individual holdings of bonds b and capital k), r is the real interest rate, w is the wage, beq is bequests, z is individual productivity, τ_i denotes various tax rates, T_y is the labor income tax, T^W is an unconditional transfer to workers (or universal basic income), and d is dividends. In this purely real model, the natural rate of interest coincides with the real interest rate ($r \equiv r^*$).

Workers face a borrowing constraint, so $a \geq 0$ at all times.¹⁴

3.3.2 The Retiree

If $g \geq \bar{g}$, the household's problem is:

$$V_g(z, h, a) = \max_{a'} u_g(c) + \beta p_g \mathbb{E}_{h'} V'_{g+1}(z', h', a') + \beta(1 - p_g) v_{g+1}(a') \quad (4)$$

¹³Primed variables denote time $t + 1$.

¹⁴The absence of borrowing is a usual assumption in an OLG model with incomplete markets. It is a simple way to deal with the possibility of agents dying with negative asset holdings. We explored the impact of a borrowing constraint below 0 and found the relevance for the r^* decline to be slight. This result is in line with Eggertsson, Mehrotra and Robbins (2019).

subject to

$$(1 + \tau_c)c + p_m m_g + a' = (1 + (1 - \tau_k)r) [a + beq] + \xi + T^M + (1 - \tau_d)d \quad (5)$$

Here ξ is the retiree's social security income, m and p_m are OOP medical goods and services quantity and their relative price, respectively, and T^M is a means-tested transfer.

Retirees face OOP medical expenses m . 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^* . OOP medical expenses may be relevant for at least two reasons. First, they lead to a less steep decline of asset holdings in late stages of life, thereby potentially affecting the impact of demographic change on r^* . Second, an upward trend in OOP medical expenses can directly affect r^* by increasing the savings motive, especially at later stages in life.

Retirees face a borrowing constraint of $a \geq 0$ for all t . Further details on the solution to the worker's and retiree's problems can be found in [appendix A.2](#).

3.4 Productivity Process

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

$$\log z_{i,g,t} = e_i + h_z^g(e_i) + \zeta_{i,t} \quad (6)$$

Productivity is the combination of three terms. First is a permanent productivity term, e_i . Households draw e_i at birth and stay with this value 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 may depend on the permanent-productivity draw e_i . Third, $\zeta_{i,t}$ is a persistent productivity shock, which evolves according to a log-normal AR(1) process:

$$\zeta_{i,t} = \rho \zeta_{i,t-1} + \epsilon_{i,t}^\zeta$$

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

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 $h_z^g(e_i)$ is chosen to create a hump-shaped labor-income profile as found in the data. The persistent-shock term $\zeta_{i,t}$ introduces idiosyncratic risk to earnings.

The productivity term $z_{i,g,t}$ is also the source of within-cohort labor-income inequality and an important source of wealth inequality. There are two sources of within-cohort inequality: the difference in permanent productivity and the idiosyncratic shock. Below we often refer to the former as permanent inequality

and the latter as inequality due to idiosyncratic risk (or inequality due to earnings risk). We study the different impact of the two shocks on inequality and r^* and their potential to explain trends in inequality and r^* in recent decades.

3.5 Production

A continuum of companies use capital and labor in their production technologies. Each firm produces a differentiated good y_{it} . The production function of intermediate goods producers is of the Cobb-Douglas form:

$$y_{it} = k_{it}^\alpha (A_t l_{it})^{1-\alpha}, \quad (7)$$

A_t is productivity, and $A_{t+1}/A_t = (1 + \gamma_t)$, with γ_t being the productivity growth rate. Intermediate goods are costlessly aggregated into the final good following a constant elasticity of substitution (CES) function. Imperfect substitutability gives rise to a markup, which we denote by μ_t .

3.6 Government

The government has four roles in our model economy. First, its debt sets the supply of bonds, thus directly impacting the equilibrium interest rate. Second, it establishes a social security system. Third, it spends G , which includes all other 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:

$$B_{t+1} = (1 + r_t)B_t + (\Xi_t - \Lambda_t) \quad (8)$$

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 , payroll-tax allowance τ_{pr} , and labor income tax $T_y(\cdot)$.¹⁵

Government expenditures are represented by the sum of exogenous government spending G , means-tested transfers T^M , unconditional transfers to workers T^W , social security spending ξ , and debt-servicing costs.¹⁶ Here we assume that the government and the social security system are consolidated. In reality, the US social security agency is its own entity. We expect that the government would bail out social security in the event of solvency concerns, and so we prefer to take the consolidated approach.

¹⁵The payroll tax is explained further in the text adjacent to equation (26) in appendix A.2.

¹⁶Exogenous government spending G enters the model as a pure resource cost.

3.7 Market Clearing and Equilibrium

We reformulate the model in terms of effective labor. For any variable X , it holds that $\hat{X}_t \equiv \frac{X_t}{N_t}$, $\check{X} \equiv \frac{X_t}{A_t}$, and $\tilde{X} \equiv \frac{X_t}{A_t N_t}$. In the steady state, variables expressed in tilde notation satisfy $\tilde{X}_{t+1} = \tilde{X}_t = X^{SS}$. The market-clearing conditions and the definition of a competitive equilibrium can be found in appendices [A.3](#) and [A.4](#), respectively.

4 Functional-Form Assumptions and Calibration

Our baseline economy is aimed at capturing features of the US economy in 2015. A model period corresponds to one year. We categorize into three groups the parameters to be calibrated. Those of Type I can be calibrated independently without a numerical procedure. In this category belong parameters for which there is a one-to-one correspondence with objects in the data. Here we simply take the values from the data for model periods in the past and, if available, projections for the future. Examples are fertility and mortality rates, tax rates, government spending, and public debt. In addition, we include parameters for the earnings and medical-expense processes. Parameters of Type II are deep structural parameters that we can't directly take from the data and for which we rely on estimates from the literature. Examples are the Frisch elasticity $\frac{1}{\psi}$ and the depreciation rate δ . Finally, Type III parameters are structurally estimated. In estimating them, we minimize the distance between appropriate moments from the data and corresponding quantities in the model. Section [4.7](#) provides further details.

We run two main experiments in the paper. The first is a comparison between hypothetical steady states. Here we compare a 2015 steady state with a 1975 steady state. The main calibration is for 2015. The structural-estimation routine for Type III parameters targets values in the 2015 steady state. When simulating the 1975 steady state, we only change the parameter values that correspond to our candidate drivers of the r^* decline. All of them can be changed directly, except for inequality. For inequality changes, we change the permanent-productivity realization e_i and idiosyncratic-shock variance σ_e to match the change in the top 10% labor-income share using a numerical optimization routine.

The second experiment is a transition-path analysis. The starting date is 1950, and the terminal date is set to ensure that the model has fully converged. All parameters except those related to r^* drivers are taken from the calibration at the 2015 steady state. The parameters related to the candidate drivers of the r^* decline change along the transition path. The start and end dates of these parameter transitions vary and depend on data availability.

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 the 1975 steady state are discussed and summarized in table [3](#) of section [5.1.1](#). Details on the path of driving variables along the transition path can be found in appendix [A.2](#). Lorenz curves and life-cycle profiles for selected variables in the 2015 steady state can be found in appendix [A.12.1](#).

4.1 Utility Functions

4.1.1 Consumption and Leisure Utility

Instantaneous utility from consumption c and labor l for an individual of age g is given by the following:

$$u_g(c, l) = \frac{(c/o)^{1-\sigma_g}}{1-\sigma_g} - \chi \frac{l^{1+\psi}}{1+\psi} \quad (9)$$

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. The Frisch elasticity is given by $\frac{1}{\psi}$, and χ is a scaling term for the disutility of working.

Assuming $\sigma_g = \sigma \forall g$ implies the familiar case of homothetic consumption preferences. Let us interpret consumption at different ages as separate goods. In a well-behaved setting – for example, one with no credit constraints – expenditure shares of each good are independent of income. Income elasticities of consumption are equalized among all goods and in particular are equal to 1. In this case the o term is irrelevant for allocations and just rescales the household's utility values.

Heterogeneous σ_g imply nonhomothetic preferences, in which case the conditions of the previous paragraph no longer hold. Expenditure shares and income elasticities now depend on income. For goods with a small σ_g , expenditure shares are low initially – that is, low for small incomes – but increase in income. There is an inverse relationship between σ_g and a good's income elasticity: relatively low- σ_g goods have a higher income elasticity. The parameter o allows us to control how fast expenditure shares and income elasticities converge from their starting values (that is, at low income) to their terminal values (income $\rightarrow \infty$).¹⁷

Let us use a simple two-goods example to make these properties clear. Denote the consumption of the 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.

---

¹⁷One way to see this is to set all prices equal to one and $c = o$ for all goods. At this implied level of income, all expenditure shares are equal. Increasing income from this point will lead to higher expenditure shares for low- σ_g goods and lower expenditure shares for high- σ_g goods.

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. σ is set accordingly.

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.¹⁸ 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 empirical estimate of $\sigma_{PI} = 0.7$ cannot even be approached with these elements alone. 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.

Figure 1: Profile of σ_g under baseline calibration

To parametrize the consumption preferences, we follow [Straub \(2019\)](#). The path of $\{\sigma_g\}$ is determined by two parameters: first, intercept parameter $\bar{\sigma}$, where the median σ_g is set equal to $\bar{\sigma}$; second, slope

¹⁸See [Carroll \(2000\)](#) and [Dynan, Skinner and Zeldes \(2004\)](#) for evidence on savings rates by income group.

parameter σ_{slope} , which determines the decay of σ_g . We assume

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

That implies exponential decay until retirement age $\bar{g}$ and constant σ_g thereafter.

In models with homogeneous σ , the curvature parameter is usually set in the range of 1.5 to 2.5. Since $\bar{\sigma}$ is the median σ_g , we want to use this range as a benchmark. We choose $\bar{\sigma} = 1.5$ in our baseline calibration and conduct sensitivity analysis with respect to this parameter. σ_{slope} is one of the parameters included in the numerical calibration routine. In our model, setting $\sigma_{slope} < 1$ is essential in order to be able to hit $\phi_{PI} = 0.7$ – that is, a low-enough elasticity of consumption out of permanent income. Figure 1 plots the profile of σ_g under the baseline calibration.

The discount factor β is included in the numerical calibration routine. 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 TFP. One interpretation of this assumption is that the level of desired consumption when young relative to the level desired when old depends on average income in the economy.^{19 20}

4.1.2 Bequest Utility

The utility function for bequests is given by the following:²¹

$$v_g(a) = b_0 \frac{1}{1 - b_{1,g}} \left[k_b + (1 - \tau_{beq}) \frac{a}{o} \right]^{1 - b_{1,g}} \quad (10)$$

Here b_0 governs the overall desire to leave a bequest, $b_{1,g}$ determines the curvature of the bequest utility function, and k_b can be used to induce a luxury-bequest motive – that is, to make rich households leave more of a bequest. τ_{beq} is an estate tax. o is time dependent and ensures that a balanced growth path exists in the case of $k_b > 0$. We set $b_{1,g} = \sigma_g$ throughout our experiments. Assets of all the deceased within a period are pooled and distributed equally among households of ages thirty to sixty-five.

4.2 Demographics

The initial generation in the model, $g = 1$, corresponds to biological age twenty-six. Below this age, labor force participation is lower because of schooling, and since the model abstracts from educational choice we exclude ages below twenty-six. There is a final generation G^d at which mortality is 100%. This

¹⁹This assumption has the flavor of status models, in which an agent's utility depends on that of other agents. In fact, [Straub \(2019\)](#) introduces a status model with *relative* wealth in the utility function that is able to match the estimate of ϕ_{PI} as well as the model with addilog preferences does.

²⁰Even with homogeneous σ_g , a utility function of the form (9) with endogenous labor does not admit a balanced growth path, unless $\sigma = 1$ (log consumption preferences). Introducing a time-dependent o term can again solve this issue. See, for example, [Boppart and Krusell \(2020\)](#) for a more extensive discussion.

²¹A functional form like this has been used as early as [De Nardi \(2004\)](#).

happens at model age $G^d = 74$, corresponding to biological age ninety-nine. We use population data and projections from the Census Intercensal Population Estimates provided by the US Census Bureau. The projections start at year 2017 and range to the last available year, 2060 (U.S. Census Bureau, 2016, 2017). An important fact to keep in mind is that given our choices regarding model age, particularly the fact that individuals enter the model at biological age twenty-six, the baby boom appears in the model in around 1975.

4.3 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, given the age-productivity profile and the idiosyncratic shock. Since we normalize the population-weighted sum of individual productivity states to 1, only the value of e_H has to be set. We add e_H to the numerical calibration routine. It is an important parameter to match the top 10% pre-tax labor-income shares, which we take from Piketty and Saez (2003).

The deterministic age profile $h_z^g(e_i)$ follows a quadratic polynomial as in Guvenen et al. (2019). We set the maximum increase in life-cycle earnings to 47% for the low permanent type and 110% for the high permanent type, based on estimates in Guvenen et al. (2015).²²

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 Rouwenhorst method and four productivity states ζ in the baseline experiment.

Inequality in the model is endogenously determined, so in order to match a specific inequality target in the 1975 steady state, we have to recalibrate some productivity-process parameters. To proceed, we first change to the 1975 values all other parameters that change between steady states. Next, we recalibrate parameters of the income process to hit the 1975 targets of our preferred inequality measure, the top 10% labor-income share. In our analysis we focus on two drivers of changes in inequality: the high permanent-productivity state e_H and the variance of the innovation of the persistent shock σ_ϵ . In principle we can change only e_H , change only σ_ϵ , or change both. In the baseline experiment, we change the variance of e_i and σ_ϵ by equal amounts until we hit the top 10% labor-income-share target. As a sensitivity analysis, we redo the analysis by changing only one of the parameters e_H and σ_ϵ in isolation.

4.4 Production

Parameters α and μ are set to match a labor share of 60.2% and a capital-to-output ratio $\frac{K}{Y} = 3$ in the 2015 steady state. 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 . This changes the profit share and hence the amount of dividends distributed. A falling labor share and a corresponding increase in the profit share are documented in

²²Figure 8 in appendix A.2 plots the profiles of $h_z^g(e_i)$.

Barkai (2020) and Eggertsson, Robbins and Wold (2018). The labor-share series on the transition path is shown in figure 15 in appendix A.2. Profits are distributed according to the empirically observed distribution by age and income reported in Piketty, Saez and Zucman (2018).

4.5 Medical Expenses

Medical expenses at age g for a household with productivity draw z and health status h , $m_g(z, h)$, follow the process

$$\log m_g(z, h) = h_m^{z,g} + h \quad (11)$$

Here $h_m^{z,g}$ is a deterministic component and h is a health, or equivalently 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. h introduces idiosyncratic risk for retirees, and only retired households face OOP-medical-expense risk. We calibrate the relative price of medical expenses p_m in the 2015 steady state to hit a ratio of OOP medical expense to output of 1.5%. This target is taken from Kopecky and Koreshkova (2014) and includes OOP nursing home expenses.

We do not have data on the long-term developments of OOP medical expenses. However, price indices for medical goods and services are available as far back as 1970 from the Bureau of Labor Statistics. 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. The increase of price of medical goods and services from 1975 to 2015 was twice the increase of the general price level.

4.6 Government

We choose labor income tax $T_y(\cdot)$ as in Heathcote, Storesletten and Violante (2017), implying a post-tax labor income given by (23):

$$wzl - T_y(wzl) = A_t \lambda_0 \left(\frac{wzl}{A} \right)^{1-\lambda_1} \quad (12)$$

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 income tax is indexed to productivity growth, ensuring a balanced growth path. Figure 11a in appendix A.2 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 US Social Security (Huggett and Ventura, 2000; De Nardi and Yang, 2014). Pre-tax benefits $T^{SS}(\cdot)$ are calculated using a measure of average lifetime income as the benefit base. The system is regressive: a higher base implies a lower relative benefit (that is, a lower replacement rate). Benefits

are taxed according to the same system as labor income, with an allowance for payroll tax τ_{pr} . The government's budget constraint, equation (8), has to hold in every period. This 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. In our setting, we can choose an exogenous path for two of the following three variables that affect the government's budget constraint: debt $\{B_t\}$, exogenous government spending $\{G_t\}$, and income tax parameter $\{\lambda_{0,t}\}$. The third variable then adjusts accordingly to clear the government's budget. In the steady-state experiments of our analysis, we take values for B and G from the data and let λ_0 adjust. To avoid counterfactual fluctuations in marginal tax rates at annual frequency and for computational reasons, we feed in a path for $\lambda_{0,t}$ together with B_t on the transition path and let G_t adjust.

4.7 Numerical Calibration Procedure

Type III parameters are jointly determined to minimize the distance between the data and model moments. We use a weighted quadratic distance:

$$\min_{\theta_{III} \in \Theta} (m_{data} - m_{model}(\theta_{III}))' W (m_{data} - m_{model}(\theta_{III})) \quad (13)$$

Here m_{data} and m_{model} are the data and model moments, respectively; θ_{III} denotes the vector of Type III parameters; Θ is the feasible parameter space; and W is a weighting matrix.

We include seven Type III parameters in the numerical procedure: first, discount factor β ; second, permanent-productivity gap between low and high type e_H ; third, the slope of the profile of age-dependent risk-aversion parameters σ_{slope} ; fourth, the relative price of medical expenses p_m ;²³ fifth, work distaste χ ; sixth, a shifter, k_b , that determines to what extent the bequest motive is stronger as income increases; and seventh, the overall strength of the bequest motive, b_0 .

The data targets to be hit are the real interest rate of 0.53%, the top 10% labor-income share of 34.87%, elasticity of consumption out of permanent income $\phi_{PI} = 0.7$, $\frac{OOP}{Y} = 1.5\%$, output (in effective-labor terms) $\tilde{y} = 1$, a share of households with bequests smaller than 18.5% of average output of 30% (of all deceased households), and the bequest-to-output ratio of 5%. We target $r^* = 0.53\%$ since it corresponds to the average of eight estimates of the level of the natural rate in 2015.²⁴

Even though all seven parameters are jointly estimated with the simulated method of moments, each of them is more strongly associated with one particular moment. Discount factor β matches r^* . Permanent productivity e_H is essential for matching the top 10% labor-income share. Likewise, $\sigma_{slope} < 1$ is necessary for the target ϕ_{PI} . We use the relative price of medical expenses p_m to hit the empirically observed ratio of OOP medical expenses to output. The appropriate choice of labor disutility χ sets output (in

²³The value of p_m in the 2015 steady state does not have meaning, as we can always change the quantity of goods and services provided per unit of cost. This is because medical goods and services do not enter the utility function in our model.

²⁴Seven 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 thirty years. See appendix A.12.2 for further details. The yield on a ten-year Treasury Inflation-Protected Security averaged 0.45% in 2015, which can be seen as additional support for our choice.

effective-labor terms) $\tilde{y} = 1$. The bequest-to-output ratio is associated closest with bequest-motive parameter b_0 , and the distributional moment of bequests is closest to k_b .

All of the targets except ϕ_{PI} have a straightforward counterpart in the model. In order to derive ϕ_{PI} , we follow [Straub \(2019\)](#). We use the model and a Monte Carlo procedure to simulate the data. We then apply the same estimation approach that has been used to estimate ϕ_{PI} with real-world data. The estimation approach for ϕ_{PI} is 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 [A.5](#).

In our numerical analysis we experiment with different versions of the baseline model – for example, a homothetic version of the model. Each time, we recalibrate the model to match the targets described above. The baseline calibration hits model targets in a fictitious 2015 steady state. We choose this year because for ϕ_{PI} we only have an estimate around this time.²⁵ Results can be seen in table [1](#), panel A.²⁶

Two of the Type III parameters are uniquely pinned down by targets and equilibrium conditions in the model. We can solve for input elasticity α and markup μ once we have targets for r^* , the labor share, and the capital-to-output ratio. We do not include these parameters in the numerical calibration procedure. The values can be seen in table [1](#), panel B.

Part of the analysis consists of comparing the 2015 steady state to a 1975 steady state. All parameters are unchanged except the ones pertaining to the candidate drivers of r^* . The same applies to the transition-path analysis.

---

²⁵As mentioned in [Straub \(2019\)](#), limited data availability prevents us from estimating ϕ_{PI} for an earlier period, including the 1970s.

²⁶The weighting matrix W is chosen using a guess-and-verify approach to minimize squared errors. Further details can be found in appendix [A.5](#).

PANEL A: Numerical Procedure			
Parameters	Sign	Value	
Discount factor	β	1.0025	
Permanent productivity differential	e_H	1.53	
Slope of σ_g profile	σ_{slope}	0.9625	
Relative price of OOP medical expenses	p_m	7.143	
Work distaste	χ	0.47	
Shifter curvature bequest motive	k_b	10	
Bequest motive	b_0	5.3	
Targets	Data	Model	Source
r^*	0.53%	0.552%	Bauer and Rudebusch (2020) and own calculations, see text
Top 10% labor-income share	34.87%	34.81%	Piketty and Saez (2003) (updated)
$\frac{\phi_{PI}}{Y}$	0.70	0.698	Straub (2019)
$\frac{OOP}{Y}$	1.5%	1.49%	Kopecky and Koreshkova (2014)
$\tilde{y}$		1.019	target $\tilde{y} = 1$
Share beq $\leq 18.5\%$ of $\tilde{y}$	30%	33.20%	De Nardi, French and Jones (2010)
$\frac{beq}{Y}$	5.0%	5.02%	Hendricks (2001), Alvaredo, Garbinti and Piketty (2017)
PANEL B: Implied by Equilibrium Conditions			
Parameters	Sign	Value	
Output elasticity wrt K	α	0.298	
Markup	μ	1.167	
Targets	Data	Model	Source
Labor share	60.17%	60.17%	BEA (2020)
K/Y	3.0	2.98	BEA (2020)

Table 1: Calibration results for Type III parameters
Note: OOP stands for out-of-pocket

Parameter	Sign	Value	Source
PANEL A: Type I Parameters			
<i>Demographics:</i>			
No. generations	ng	74	
Retirement age	$\bar{g} + 25$	66	
Survival probabilities	$p_{g,t}, \bar{p}_{g,t}$		(U.S. Census Bureau, 2016, 2017)
Population growth rate	n	1.15%	(U.S. Census Bureau, 2016, 2017)
<i>Production:</i>			
TFP growth rate	γ	0.7%	Fernald (2014) (updated)
<i>Fiscal:</i>			
Public debt/Y	B	74.2	Congressional Budget Office (2019)
Exogenous gov spending	G	13.5	Congressional Budget Office (2019)
Consumption tax	τ_c	0.05	Kitao (2014)
Capital income tax	τ_k	0.4	Straub (2019)
Profit tax	τ_{corp}	0.25	Straub (2019)
Bequest tax	τ_{beq}	0.1	Straub (2019)
Curvature income tax function	λ_1	0.181	Heathcote, Storesletten and Violante (2017)
<i>Earnings Process:</i>			
Persistence earnings shock	ρ_w	0.9	Guvenen et al. (2019)
Var earnings shock	σ_ϵ^2	0.04	Guvenen et al. (2019)
Age profile max increase	e^{hc}	[0.47, 1.1]	Guvenen et al. (2015)
<i>Health:</i>			
Deterministic med expenditure	$h_m^{e,g}$		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)
PANEL B: Type II Parameters			
Median coefficient of RRA	$\bar{\sigma}$	1.5	Gourinchas and Parker (2003)
Depreciation rate	δ	0.08	Nadiri and Prucha (1996)
1/Frisch elasticity	ψ	0.5	Whalen and Reichling (2017)
Scaling parameter consumption utility	ϕ	0.3 $\tilde{y}$	Straub (2019)

Table 2: Parameters under baseline calibration (Type I and Type II)

5 Numerical Analysis

5.1 Steady-State Analysis

5.1.1 Baseline Experiment

The first set of results looks at the change in the natural rate between two steady states. We compare a 1975 steady state to a 2015 steady state and decompose the total decline into the contributions of individual drivers. We choose 1975 as the reference year because it is around this time when the baby boom generation enters our model and triggers a substantial change in the demographic structure.

The natural rate is calibrated to hit 0.53% in 2015, and the model comes close under the baseline parametrization, with a value of $r^* = 0.55$. As described in the previous section, this target is chosen according to the average of eight estimates of the natural rate. If we apply the same procedure to 1975, we get an estimate of the natural rate of 1.91%.²⁷ This implies a decline of 1.38% between 1975 and 2015. The size of the decline in r^* over this period ranges from 0.45% to 3.2% in the eight reference series.²⁸

We want to highlight that the model-implied change in r^* between the two points is *not* targeted. This holds not only for the comparison of steady states between 1975 and 2015 in this section, but also for initial and terminal steady states on the transition path in section 5.2. Only the level of r^* in the 2015 steady state enters the calibration routine.

	1975	2015
TFP growth rate(%)	1.5	0.7
Population growth rate (%)	1.92	1.15
IP survival from age 26 to 99 (%)	5.1	8.8
Top 10% share (%)	26.56	34.87
Public debt (%)	25.4	74.2
OOP price (rel. to 1975)	1.0	2.1
Labor share (%)	64.3	60.2
Exog gov spending (%)	13.2	13.5

Table 3: Selected data for variables in 1975 and 2015. Regarding demographics drivers, both the population growth rate n and survival probabilities $\bar{p}_g$ are changing between steady states. For the top 10% labor-income share, we change e_H and σ_e^2 as described in section 4.3. “Public debt” is the ratio of government debt to GDP. Price for out-of-pocket (OOP) medical expenses is relative to price in 1975. For “Labor share” we change the markup μ . “Exog gov spending” refers to G and excludes social security and old-age medical spending.

The only parameter values to be changed going from the 2015 steady state to the 1975 steady state are the ones pertaining to the drivers, summarized in table 3. All other parameters are left unchanged. To isolate the effect of each driver, we change the parameter(s) connected to each driver alone while leaving all other parameters at the 2015 steady-state value. Following this procedure, table 4 displays the individual effect of each driver. The effect of changing all drivers simultaneously is shown in the second row.

²⁷While the 1970s featured major swings in interest rates, we find the r^* estimates to be quite stable. If we take the average of the eight r^* estimates over the full decade, we get a value of 1.98%.

²⁸Figures 21 and 22 in appendix A.12.2 plot the range of estimates from the literature and the trend of r^* constructed as a moving average, respectively.

	Δr^*	%
Total r^* change	-2.161	100.0
TFP growth	-1.00	46.2
Demographics	-0.71	32.7
Inequality	-0.70	32.4
Public debt	0.31	-14.3
OOP	-0.14	6.6
Labor share	0.11	-5.0
Exog gov spending	0.03	-1.4
Interactions	-0.06	3.0

Table 4: Change in r^* going from 1975 steady state to 2015 steady state, and decomposition in baseline model. Rows 2-9 change one driver of r^* at a time. Column 2 shows the change in r^* in percentage points while column 3 shows the relative contribution of each row to the total. “Exog gov spending” refers to exogenous public spending (G) and does not include social security. Rows 2-9 sum to row 1. “Inequality” changes both permanent productivity and persistent shock variance.

Our model accounts for a decline in r^* of 2.16 pp when moving from the 1975 steady state to the 2015 steady state. The largest contribution comes from the slowdown in productivity growth, with -1 pp, or 46.1% of the total decline. The second-largest part of the decline is explained by demographics: r^* declines by 0.71 pp when we change demographic variables alone. Inequality comes very close and is the third most important driver, with -0.7%, or 32.4% of the total. Most of the decline is due to these three forces. The increase in public debt is a notable counteracting force, increasing r^* by 0.31 pp, or 14.3%.

We find that OOP medical expenses can only account for a 0.14 pp decline, or 6.6%. We elaborate on this result in appendix A.10. The labor share’s contribution is of similar magnitude but opposite sign. The contribution of a declining labor share is positive because it implies a larger share of profits to be distributed. Profits disproportionately go to the elderly in the form of dividends. Since the elderly receive more dividends, they do not have to save as much of their labor income, thus reducing the aggregate supply of savings. The change in public spending G from the 1975 steady state to the 2015 steady state is simply too small to have much of an impact.²⁹

Surprisingly, interactions are small, even though we work in a rich model with many moving parts. This result is robust to recalibration. However, it would be wrong to conclude that this implies that we can easily compare estimates across models and across papers. As we highlight in sections 5.1.2 and 5.4, some of the results are sensitive to shutting down elements of the model and changing the parametrization.

Appendix A.8 argues that international capital flows are likely not as quantitatively important in explaining the decline in r^* as the top three explanations found in table 4. Appendix A.9 shows that changes in life expectancy drive most of the impact of demographic change on r^* in our analysis. In particular, what matters are individual households’ responses in optimal consumption-savings decisions, while the effect of demographic change on the composition of the population is of lesser importance.

The results of table 4 are consistent with a general picture that emerges after conducting sensitivity analysis: the model can account for a decline in r^* of 2 pp. Most of this movement is due to demographics,

²⁹Changes along the transition path will be higher by a factor of more than ten, so this driver will have more relevance.

growth, and inequality, with each explaining a decline of about equal magnitude. We consider inequality's large role an important finding. Related studies have highlighted that inequality's contribution can be large. We are able to show that this is indeed true when directly comparing inequality with other important drivers within the same model.

5.1.2 Homothetic-Economy Results

Table 5 highlights the importance of different assumptions regarding homotheticity in the model.³⁰ Here we again look at the change in r^* between 1975 and 2015 steady states and make a decomposition by drivers. We compare three economies. Baseline "Fully Nonhomothetic," shown in columns 2 and 3, replicates the results from the previous section and is the reference point. "Homo C," columns 4 and 5, stands for homothetic consumption preferences and assumes $\sigma_g = \bar{\sigma} \forall g$. It highlights the impact of introducing age-dependent coefficients of risk aversion for our analysis. "Fully Homothetic," columns 6 and 7, is useful for illustrating the concept of homotheticity.

The Fully Homothetic economy is one in which changes in permanent productivity e_i have no effect on aggregate prices such as r^* and w and no effect on aggregate quantities if the population-weighted sum of productivity stays unchanged. It is an extreme case in which changes in permanent inequality by definition do not affect r^* . The intuition for this result is that changes in permanent income inequality by definition entail redistribution of income across households in a permanent fashion. If winners and losers of this redistribution have the same savings rate, then the overall savings supply in the economy is unchanged and so is r^* . We make the following changes relative to the baseline: (1) homothetic consumption preferences, such that $\sigma_g = \bar{\sigma} \forall g$, (2) a linear social security system (see equation (40) in appendix A.6), (3) no bequest motive and 100% taxation of accidental bequest, such that $b_0 = 0$ and $\tau_{beq} = 1$, (4) linear labor income tax, such that $\lambda_1 = 0$, (5) exogenous labor supply, (6) full taxation of profits, such that $\tau_d = 1$, (7) homogeneous age-productivity profile $h_z^g(e_L) = h_z^g(e_L) \forall g$, and (8) OOP medical expenses are proportional to permanent productivity type e_i . These assumptions ensure that the utility function is homothetic, that tax- and transfer-schedules are linear, and that there is no redistribution through bequest or otherwise.³¹ Appendix A.6 provides further details on the parametrization of this economy. For each economy we recalibrate to hit the targets of table 1.³²

Results: The total drop in r^* is fairly stable, always close to 2 pp. The most interesting result is seen when looking at the contribution of inequality. In the baseline model, it explains about a third of the decline, or 0.7 pp. Introducing the assumption of homothetic consumption preferences cuts this number by more than half to 0.32 pp, or 16.2%. In the fully homothetic economy, the contribution of inequality is negligible at 0.03 pp, or 1.2%. Interestingly, the lower contribution of inequality is almost entirely

³⁰ Appendix A.11 explores an additional aspect of working with nonhomothetic consumption preferences: the implications for including productivity growth in the model. There we partially relax the assumption that the o term in the utility function grows at the same rate as productivity.

³¹ This economy therefore meets assumptions 1-4 of propositions 1-3 in Straub (2019). The propositions prove linearity in permanent income.

³² Economy Homo C can hit all targets but ϕ_{PI} . Fully Homothetic can hit all targets but ϕ_{PI} and the targets related to bequest.

	Baseline		Homo C		Fully Homothetic	
	$\Delta r^* \text{pp}$	%Total	$\Delta r^* \text{pp}$	%Total	$\Delta r^* \text{pp}$	%Total
Total	-2.161	100.0	-1.953	100.0	-2.194	100.0
TFP growth	-1.00	46.2	-1.00	51.3	-0.87	39.5
Demographics	-0.71	32.7	-0.80	41.0	-1.19	54.1
Inequality	-0.70	32.4	-0.32	16.2	-0.03	1.2
Public debt	0.31	-14.3	0.26	-13.5	0.38	-17.3
OOP	-0.14	6.6	-0.15	7.4	-0.31	14.0
Labor share	0.11	-5.0	0.05	-2.5	-0.26	11.7
Exog gov spending	0.03	-1.4	0.02	-0.9	0.01	-0.6
Interactions	-0.06	3.0	-0.02	1.1	0.06	-2.7

Table 5: Change in r^* going from 1975 to 2015 steady state, and decomposition in three versions of the model. “Baseline” refers to fully nonhomothetic, “Homo C” assumes $\sigma_g = \bar{\sigma} \forall g$, and “Homothetic” is explained in the text. Rows below “Total” change one driver of r^* at a time. Columns 2, 4, and 6 show the change in r^* in percentage points (pp); columns 3, 5, and 7 show the relative contribution of each row to the total. “Exog gov spending” refers to exogenous public spending (G) and does not include social security. Rows of individual drivers sum to Total. “Inequality” changes both permanent productivity and persistent shock variance.

(or a third in Homo C) made up for by a larger contribution of demographics. Demographics is the most important driver, ahead of TFP growth, in the fully homothetic model. OOP medical expenses make a larger contribution in the fully homothetic model, with -0.31 pp – more than double the amount compared to the other two versions.

	Baseline		Homo C		Fully Homothetic	
	1975	2015	1975	2015	1975	2015
r^* (%)	2.66	0.55	2.70	0.49	2.45	0.51
Top 10% share (%)	26.6	34.8	26.6	34.8	26.6	34.8
ϕ_{PI}		0.70		1.05		1.01

Table 6: Selected variables going from 1975 to 2015 steady state in three versions of the model. “Baseline” refers to fully nonhomothetic, “Homo C” assumes $\sigma_g = \bar{\sigma} \forall g$, and “Fully Homothetic” is explained in the text. “Inequality” changes both permanent productivity and persistent shock variance. ϕ_{PI} is derived using method described in appendix A.5.

The reason we choose the model with heterogeneous consumption preferences as the baseline model is that it is able to match the data on consumption response to permanent income changes – that is, the elasticity ϕ_{PI} . Table 5 shows steady-state interest rates and top 10% share for the three economies. We can see that all three economies can account for the increase in inequality from 1975 to 2015 seen in the data. We recalibrate each version of the model, so all hit the targets for r^* and top 10% labor-income share in the 2015 steady state. However, the baseline model is more consistent with the data than the other two versions in the following respect. The last row in table 6 shows the value of ϕ_{PI} . The baseline model was calibrated to hit the data estimate of 0.7. The fully homothetic model by definition should yield a value of 1, and the model comes very close at 1.01. Eliminating nonhomothetic consumption preferences increases the estimate to 1.05.³³

³³The estimate of ϕ_{PI} in Homo C goes above 1 because in the model capital income increases in a perfectly correlated way with permanent income. Higher capital income leads to more consumption, ceteris paribus. As the estimate shows, the effect of higher capital income is so strong that the ϕ_{PI} goes above 1. We report ϕ_{PI} without controlling for capital income to stay consistent with the baseline data estimate in Straub (2019).

These results show that the impact of permanent productivity changes on r^* is sensitive to the degree of nonhomotheticity in the model. Working in a fully homothetic model, one would conclude that permanent productivity differences and resulting inequality can't account for any decline in r^* . But such a model can't account for the higher savings rates of high-permanent-productivity agents either. This is precisely the reason why we prefer the results from the baseline version of the model. We also see that nonhomothetic consumption preferences can help to match the ϕ_{PI} estimated from the data and make the permanent-inequality channel more powerful.

5.1.3 Inequality Due to Permanent Productivity versus Income Risk

Our model can generate within-age inequality through two mechanisms: differences in permanent productivity e_i and persistent productivity shocks $\zeta_{i,t}$. In this subsection we take a closer look at which mechanism is more relevant for the natural-rate decline.

	Total	Inequality	Permanent	Idiosyncratic	Interaction
PANEL A: Change Both					
Δr^*	-2.16	-0.70	-0.58	-0.10	-0.02
% of inequality		100.00	82.90	14.80	2.20
% of total	100.00	32.40	26.80	4.80	0.70
Top 10% share	26.60	24.70	25.70	34.10	
σ earnings	0.46	0.44	0.48	0.55	
PANEL B: Only Permanent					
Δr^*	-2.08	-0.65	-0.65		
% of inequality		100.00	100.00		
% of total	100.00	31.30	31.30		
Top 10% share	26.60	24.80	24.80		
σ earnings	0.48	0.46	0.46		
PANEL C: Only Idiosyncratic					
Δr^*	-1.61	-0.19		-0.19	
% of inequality		100.00		100.00	
% of total	100.00	11.60		11.60	
Top 10% share	34.60	33.30		33.30	
σ earnings	0.53	0.52		0.52	

Table 7: Change in r^* and selected statistics going from 1975 to 2015 steady state, and decomposition by inequality source in baseline model. Panel A adjusts both permanent productivity e_H and variance of idiosyncratic shock σ_e under the “Inequality” driver. Panel B only changes permanent productivity, and panel C only changes idiosyncratic shock variance under the “Inequality” driver. Column “Total” changes all drivers to 1975 steady state. Column “Inequality” changes permanent productivity and/or persistent shock variance. Column “Permanent” solely changes permanent productivity; Column “Idiosyncratic” solely changes persistent shock variance. Row 1 in each panel shows the change in r^* in percentage points. Rows “Top 10% share” and “ σ earnings” show the model implied top 10% labor income share and the standard deviation of labor income at the 1975 steady state, respectively.

The baseline experiment of section 5.1.1 was conducted under the assumption that the variance of the permanent productivity draw and the variance of the idiosyncratic productivity innovation changed by the same magnitude in order to match the top 10% labor income share in 1975. To understand which source of inequality is more relevant for the r^* decline, we change each productivity term in isolation,

holding the other one constant. This is the same approach as in section 5.1.1, where we also changed one driver at a time and looked at the response of r^* .

Table 7, panel A shows the results. We see that the change in permanent productivity accounts for most of the contribution of inequality. The total effect of inequality on r^* is 0.70 pp. Of this effect 83%, or 0.58 pp, is attributable to permanent productivity movements. Only 14.8%, or 0.10 pp, is due to the increase in variance of the idiosyncratic component. Interactions are small. We can also see that changes in permanent productivity account for most of the change in the top 10% labor-income share. The top 10% share increases from 26.56% in 1975 to 34.87% in 2015, or 8.31 pp. The change in the idiosyncratic-income term accounts for less than 1 pp of this increase. Changing innovation variance σ_ϵ to the implied 1975 value only reduces the top 10% share to 34.1%.

Since productivity is unobserved, it is not clear whether our approach of changing the variance of both productivity terms by the same factor is the correct way to proceed. To get an idea of the sensitivity of the results to this assumption, we rerun the experiments under two extreme scenarios. The first attributes all change in inequality to a change in permanent productivity, holding the idiosyncratic variance constant; the second holds permanent productivity constant and solely changes the variance of the persistent shock (idiosyncratic variance). Table 7, panels B and C show the implications.

The results confirm our previous finding: it is the change in permanent productivity that matters for r^* . If we only change permanent productivity to hit the top 10% labor-income share of 1975, then the contribution of inequality in absolute (-0.65) and relative (31.3%) terms is of a similar magnitude to that under the baseline. Solely changing idiosyncratic productivity leads to a very different picture. Even if we completely eliminate the risk – that is, set the variance to zero – inequality only declines by 1.3 pp and r^* by 0.19 pp. A change in the idiosyncratic term is not nearly enough to account for the change in the top 10% labor-income share observed in the data. This also means that most of the inequality in labor income comes from permanent productivity differences to begin with. Remember that permanent productivity differences are still present in this experiment, but, different to the experiment of panel A, we leave permanent productivity unchanged when going to the 1975 steady state.

Finally, we also show the standard deviation of labor earnings for each scenario. This number is in the range of 0.45 to 0.55 throughout and is 0.58 in the 2015 steady state, while in the data this number is about 0.8 in 1980 and increases to 0.92 by 2015. The model is thus not able to match the absolute size of the standard deviation of labor earnings, but it can generate the size of the relative increase.³⁴

5.2 Transition-Path Analysis

The transition-path analysis starts at a hypothetical 1950 steady state. In 1951 the full path of parameter changes is revealed to our model agents, who act with perfect foresight. The parameters that are changing are the ones pertaining to the drivers of table 3. Details on the changes have been outlined above; figure

³⁴We confirm in unreported results that the qualitative conclusions do not change if we increase the variance of the idiosyncratic-productivity term to a relatively high value of $\sigma_\epsilon = 0.1$.

15 in appendix A.12 shows the transition path of the variables. By 2060 all shocks reach their terminal steady-state value, with demographic variables and debt-to-output ratio being the last to converge. It takes a very long time for the model to converge to a new steady state. r^* starts oscillating around the terminal steady-state value starting around 2100, but the echo from demographic change doesn't disappear for several hundred more years.

Figure 2: Transition path of r^* and top 10% labor-income share. Data for top 10% labor-income share are from Piketty and Saez (2003). Values on vertical axes in both panels are in percentages. Note that the horizontal axes, denoting time, are different in the two panels.

Figure 2a shows a transition path for r^* . We see that r^* stands at 4% in 1950 and starts to decline immediately. Note the drastic decline after 2000. This is driven by the continuation of the decline in productivity growth after the brief IT-revolution-led boom in the '90s and by the baby boom generation's reaching the age of peak asset holdings.

r^* hits a low in 2031 at about 0.38%, from which point on a gradual increase starts. It stabilizes at 1.04%. Although the model implies an increase in r^* after 2030, the increase is modest. Even though public debt increases substantially to 200% of output in the terminal steady state, debt-servicing costs only rise modestly.³⁵

The rise in r^* after 2030 is mainly driven by two factors. First, the dip around 2030 is really the effect of the baby boom generation's hitting the age of peak asset holdings and retiring. We see a similar pattern about thirty years later when the children of the baby boom generation, another comparatively large cohort, reach the same stage of life.

Second, starting in 2008 public debt starts a steep upward trajectory (see figure 12 in appendix A.2). The debt-to-output ratio rises from about 40% in 2008 to 200% by 2060. Although we find that r^* is not as sensitive to increases in debt in our model as it is in other studies, the sheer size of the increase in the

³⁵In fact, increasing the level of public debt is a way for the government to raise resources and to lower the tax-to-output ratio at this level of debt. We return to this point in section 5.3.

debt-to-output ratio still pushes r^* up by a substantial amount. The echo in the r^* dynamics comes from demographics; more specifically, the baby boom generation is large. Baby boomers' children are also a relatively large cohort, and in turn their children's children are too. According to the model, this effect dissipates slowly. Figure 2b shows inequality dynamics. The top 10% labor-income share implied by the model lines up well with the data. Note that inequality is an endogenous object in the model. The transition path for additional variables can be found in figures 23 to 25 in appendix A.12.2.

Figure 21 in appendix A.12.2 shows the transition path for r^* implied by the model together with the average and range of estimates of r^* from the literature, while figure 22 shows the model-implied path of r^* together with a smoothed short-term real rate from the data. The model-implied path lines up well with the estimates. In particular, both model and estimates seem to suggest an accelerated decline starting around 2000. The magnitude of the decline between 1970 and 2015 is of a comparable size.

5.2.1 Scenario Analysis

The transition path in figure 2a relies on assumptions about the future path of driving variables such as demographics, productivity growth, the earnings process, and public debt. Here we look at the sensitivity to these assumptions. One question of interest is how low or high r^* can become under still-plausible assumptions about underlying drivers.

Table 8 shows values of the respective drivers in the terminal steady state under “low” and “high” scenarios. Appendix A.12.4 gives the full transition paths of the drivers and further details on the source and construction of the data.

	Baseline	Scenario “Low”	Scenario “High”
Population growth rate (%)	0.46	−0.64	1.04
Public debt (%)	200.00	50.00	350.00
$\ln \frac{e_H}{e_L}$	1.43	1.64	1.22
σ_e^2	0.04	0.05	0.03
TFP growth (%)	0.60	0.10	1.10

Table 8: Parameters in terminal steady state of transition path for three different scenarios: “baseline,” “low,” and “high.” Values of earnings process, $\ln \frac{e_H}{e_L}$ and σ_e^2 , under low scenario lead to decline in r^* when compared to baseline and therefore fall into the low-scenario category. All other parameters are held at the baseline transition-path calibration. See appendix A.12.4, figures 26 to 29 for full path of each variable under scenarios low and high and for details on construction of the paths.

A few comments on the scenarios chosen are in order. For population growth we rely on UN forecasts (United Nations, 2019). They present several scenarios regarding fertility, mortality, and migration. The low scenario assumes a lower fertility rate together with no migration. As a consequence of this assumption, population growth turns negative around 2040 (see figure 32 in appendix A.12.3). In the high scenario, population growth rate increases to levels observed around 2000 because of higher fertility. For public debt we have a low scenario that stabilizes the ratio of public debt-to-output at the postwar average of 50% and a high scenario with a ratio of 350%. The labor income tax parameter λ_0 adjusts to ensure the government's budget constraint holds. We include scenarios under which variables of the

earnings process continue to change after 2015. As in the baseline, we let the permanent productivity in the high state, e_H , and variance of the innovation to the persistent earnings shock, σ_ϵ , adjust. In the low scenario we assume that the pre-2015 dynamics of these variables continue for another twenty-five years. In the high scenario, they revert over the course of the next twenty-five years. For the TFP growth rate we contemplate a rise or fall of 0.5 pp.³⁶ All other drivers are held constant at baseline values. Note that the changes in low scenarios all depress the natural rate further, at least in the terminal steady state, while the opposite holds in the high scenarios.

Scenario	r^* (%)	
	“Low”	“High”
Baseline	1.04	
Population growth rate	0.94	1.14
Public debt	0.10	2.01
$\ln \frac{e_H}{e_L}$ and σ_ϵ^2	0.70	1.41
TFP growth rate	0.39	1.73
ALL	-1.00	3.17

Table 9: r^* in terminal steady state under various scenarios for driving variables. “Baseline” refers to baseline calibration introduced in previous section. Rows other than “ALL” change one driver at a time, either to the “low” value or “high” value of table 8. “ALL” changes all parameters together, to value of either category low or high in table 8. Values are in percentages.

Table 9 summarizes the results by showing the natural rate in the terminal steady state in all scenarios. Figure 3 shows the implied path for r^* when we change all four drivers together (“ALL” scenario). The transition path when changing the drivers individually can be found in appendix A.12.3, figures 32 to 35.

We see that under the “ALL low” scenario, the natural rate continues its downward trend and goes negative around 2035. At the terminal steady state, the natural rate is negative at -1.0%. This means that our model can produce a permanent negative-natural-rate environment under plausible values for the driving variables. Figure 30 in appendix A.12.3 shows that the top 10% labor-income share increases to about 42% in this case. Under the “ALL high” scenario, the natural rate converges to 3.17%, a level comparable to that in the 1960s in our model.

The largest contributor to the difference in the terminal steady-state value of r^* in the high and low scenarios is public debt, followed by the TFP growth rate and drivers of the earnings process. A debt-to-output ratio of 150 pp above (below) our baseline steady-state value of 200% increases (decreases) r^* by about 1 pp.

The low population growth rate does not add much further downward pressure on r^* – only about 10 bp. Appendix A.9 shows that for the 1975 to 2015 period, changes in survival probabilities are more important for the decline in r^* than changes in the population growth rate. The elasticity of r^* with respect to the population growth rate seems to be even smaller under the population structure projected for the next decades. It should also be noted that life expectancy seems to have peaked in the U.S. in recent years, so it is unclear whether we should expect large increases of life expectancy in the future.

³⁶These assumptions follow Congressional Budget Office (2019).

Figure 3: Transition path r^* implied by model under baseline calibration and under “ALL low” and “ALL high” scenarios. The parametrization of the terminal steady state is explained in table 8. We plot a linear transition from 2100 to the terminal steady-state value (dotted segments of the lines).

5.3 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 of households to supply savings and labor. 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 want to propose the idea of 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 regarding the optimal level of a 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. A cost we can account for is the impact of policy on the steady-state level of output, which we report as well in this section. We leave a more thorough normative analysis for future research.

We undertake a comparative-statics exercise at the baseline 2015 steady state. Table 10 shows the results. The first and second columns list the policy variables and their values at the baseline calibration. Column 3 shows the response of r^* (in basis points) to a 1 pp change in the policy parameter. Column 4 shows the response of steady-state output. Since output is normalized to 1 in the initial steady state, the number

Policy Parameter (x)	Baseline Value	$\frac{dr^*}{dx} \times 100$	$\frac{dy}{dx}$	$\frac{dT^{inc}}{dx}$
Consumption tax τ_c (%)	5.00	-3.25	0.17	-0.47
Profit tax τ_{corp} (%)	25.00	-1.25	0.17	-0.12
Capital income tax τ_k (%)	40.00	2.30	-0.16	-0.09
Estate tax τ_{beq} (%)	10.00	0.47	0.00	-0.05
Tax progressivity λ_1	0.181	4.51	-0.82	0.05
Basic income T (%)	0.00	-0.52	-0.56	0.68
Public coverage of med expenses T^M (%)	0.00	0.38	-0.02	0.01
Retirement age $\bar{g} + 25$	66.00	0.79	0.75	-0.55
Replacement rate ϕ (%)	100.00	1.66	-0.11	0.05
Public debt B (%)	74.20	0.64	-0.01	-0.02
Exog public spending G (%)	13.50	10.11	0.19	0.88

Table 10: Response of selected variables to change in policy parameter 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 steady-state output in effective-labor terms, in levels, to a 1 pp change in the policy parameter. Column 5 shows response of labor income tax receipts relative to output in percentage points to change in parameter by 1 pp. The only exception in all three columns is “Retirement age,” which denotes response to change in retirement age of one year. Note that since households enter the model at biological age twenty-six, retirement age in data corresponds to $\bar{g} + 25$.

also reads as the percentage change of output. We let the level parameter of the income tax λ_0 adjust to ensure that the budget constraint holds at the new steady state. Column 5 shows the impact on labor income tax revenue. Table 17 in appendix A.12.4 shows additional statistics.

The table reads as follows: A change in the consumption tax rate τ_c by 1 pp from the baseline value of 5% to 6% leads to a decline in r^* by 3.25 bp. The response of output to the same change is an increase of 0.17%. The response of r^* and output turns out to be fairly linear for moderate changes. Consequently, given that the steady-state value of τ_c is 5%, eliminating the consumption tax would increase r^* by roughly 15 bp and lower output by about 0.85%.

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% would lead to a rise in the natural rate of about 23 bp. If we were to use the additional tax receipts to fund public coverage of 90% of medical expenses, r^* would increase by another 34 bp, making for a total of 57 bp. To put this number in perspective, this policy would reverse about 80% of the impact on r^* of demographic change over a span of forty years (compare section 5.1.1). Both policies would reduce incentives to save: the tax would lower the post-tax return; coverage of medical expenses would reduce expenses in old age and thus the need to accumulate savings.

Higher capital income taxes and interest rates crowd out capital investment, so the policies come with an output cost. The proposed change in the capital income tax would lower steady-state output by about 1.6%. Adding medical coverage would bring the total fall in output to 3.4%. Note that these numbers are level effects, as output growth is exogenous in our model. In the following, we want to highlight the impact of additional policy proposals.

Figure 4: Left panel shows response of r^* to change in ratio of public debt-to-output $\frac{B}{Y}$ at the steady state. Right panel shows response of labor income tax receipts relative to output (dashed line) and "output" (solid line). Values for r^* and $\frac{T_{inc}}{Y}$ are in percent, "output" is in levels and refers to output per effective labor ($\tilde{y}$). Labor income tax level parameter λ_0 adjusts to clear the government's budget. All other parameters are held at the baseline 2015 steady-state value.

Public Debt: Figure 4, panel (a) shows the response of r^* to changes in the steady-state ratio of public debt-to-output, $\frac{B}{Y}$. Panel (b) shows the effect on labor income tax revenues and output.³⁷ We see that the response of r^* to changes in $\frac{B}{Y}$ is basically linear and is modest: it would take a value of $\frac{B}{Y}$ of almost 500% to raise r^* to a level of 3%.

Note that increasing the steady-state ratio of public debt-to-output is a way for the government to *raise* resources.³⁸ Figure 4, panel (b) shows that the labor income tax receipts needed to clear the government's budget in the steady state are lower as the debt ratio increases. We can get intuition for this result if we assume for a moment that r^* was negative in the steady state. In this case, lenders basically pay the government to borrow from them. In a model with population and productivity growth, the relevant number to look at is $r^* - n - \gamma$, with n and γ being the rates of population and productivity growth, respectively. In our baseline steady state, the natural rate is low, 0.5%, but population and productivity still grow at a positive rate.

Naturally, increasing public debt can't be a source of finance indefinitely, as higher debt levels increase the natural rate. Higher r^* means that debt-servicing costs on all public debt go up. There is thus an inflection point at which raising the debt level implies that tax revenues have to go up in order to clear the government's budget. In our model, this happens at a debt-to-Y ratio of around 400%. The output cost of moving to such a high debt level is surprisingly low. At the inflection point of 400%, output only falls by about 3% from the initial steady state, even though debt rises by 325 pp.³⁹

³⁷Here we again make the assumption that labor-tax parameter λ_0 adjusts to clear the government's budget.

³⁸This point has been made before – for example, in Mehrotra (2018).

³⁹Note that labor supply is endogenous in our model, likely contributing to the relatively low output cost of high debt levels. As the debt level rises, the fall in the labor income tax counteracts the crowding-out effect of higher interest rates since households are more willing to supply labor, ceteris paribus.

Tax-and-transfer policy: The response of the natural rate to the different tax rates included in our model varies considerably. While increasing the consumption tax or the profit tax lowers the natural rate, a higher capital income tax and estate tax go in the other direction. The tax changes affect the incentives to save differently. To understand the effect of the profit tax, keep in mind that most of the dividends go to retired households, and disproportionately to high-income households. Since the latter rely a lot on savings to finance consumption when old, lower post-tax dividends encourage saving and thus lead to a lower natural rate.

Parameter λ_1 affects tax progressivity. Changes in the parameter itself are hard to interpret, especially when λ_0 changes at the same time. It is more intuitive to compare marginal tax rates (MTRs). Luckily, it turns out that increasing λ_1 by 1 pp roughly implies a 1 pp increase in the MTR at the 90th percentile of the labor-income distribution.⁴⁰ We find that raising λ_1 by 10 pp, to 0.281, increases the MTR at the 90th percentile by 10.03 pp.⁴¹ The resulting change in the natural rate is an increase of 48 bp. Here we again let λ_0 adjust. This policy also raises average tax rates (ATRs) at the 90th percentile (from 17.77% to 20.28%), but the ATR declines for the median worker (from 6.86% to 4.85%). The policy is thus truly redistributive.⁴² The effect on output is -8.2%, a rather large decline per increase in r^* when compared with other policies.

A basic-income program for workers, T^W , does not raise the natural rate. The transfer is not well targeted: since it goes to all workers, no matter their income, households tend to save the additional income. This especially holds true if the policy is financed through a rise in the income tax, $T_y(\cdot)$. Note that households in retirement also pay income tax, and the lower post-tax income in retirement induces additional savings. The impact of covering medical expenditures is discussed in the next section.

Figure 5: Response of r^* to a change in replacement-rate parameter ϕ . Labor income tax-level parameter λ_0 adjusts to clear the government's budget. All other values are at the baseline 2015 steady state. Values for r^* and replacement-rate parameter ϕ are in percentages. ϕ is 100% under the baseline calibration.

⁴⁰This holds true under the assumption of λ_0 adjusting to clear the budget in response to a change solely in λ_1 . Appendix A.12.4 discusses the marginal and average tax rates implied by our model and plots those rates.

⁴¹From 32.65% to 42.68%.

⁴²The top 10% post-tax income share falls 2.63 pp (from 31.8% to 29.1%) and the top 10% wealth share by 3.7% (from 57.2% to 53.5%) in response to the policy.

Social security: Figure 5 shows the responsiveness of r^* to a change in replacement-rate parameter ϕ . Parameter ϕ shifts the retirement benefit of all agents proportionally. In the baseline calibration we set $\phi = 1$. Thus, cutting ϕ to 0.5 would reduce the retirement benefit of each household by 50%.

The figure highlights the drastic effects of moving from a pay-as-you-go retirement system to a fully funded system – that is, one with $\phi = 0$. Without a retirement benefit, agents would have to accumulate enough savings when young to cover all their spending in retirement. This would shift out the savings supply considerably. The natural rate would fall below -1.5%. A pay-as-you-go system does not create additional savings. The benefits of the elderly are paid by contributions of the young directly. At the baseline calibration, the sensitivity of r^* to a 1 pp increase in ϕ is 1.66 bp, and the response is fairly linear at that point.

We find that r^* is insensitive to changes in the mandatory retirement age. Increasing the retirement age by four years to seventy has an impact on r^* of less than 4 bp.⁴³

5.4 Sensitivity

We want to see how sensitive the results are to our choice of parameter values taken from the literature and Type II parameters taken from table 2. Table 11 shows the results if we choose low and high values for the median value of the coefficient of relative risk aversion $\bar{\sigma}$ (in columns 3 and 4), the depreciation rate δ (columns 5 and 6), the Frisch elasticity $1/\psi$ (columns 7 and 8), and the scaling term in the consumption-utility function o (columns 9 and 10). We include the baseline in column 2 for easy comparison. We show the total change in r^* going from a 1975 steady state to a 2015 steady state, the total change going from a 2015 steady state to the terminal steady state of section 5.2, and the decomposition into individual drivers going from 1975 to 2015.

	Base	$\bar{\sigma} = 1.167$	$\bar{\sigma} = 2.1$	$\delta = 5.6\%$	$\delta = 10.7\%$	$1/\psi = 0.91$	$1/\psi = 3.37$	$o = 0.25$	$o = 0.35$
Δr^* 1975-2015	-2.16	-2.30	-1.87	-2.25	-2.03	-2.19	-2.13	-2.18	-2.00
Δr^* 2015-term	0.46	0.30	0.68	0.58	0.45	0.46	0.72	0.36	0.71
TFP growth	-1.00	-1.06	-0.89	-0.91	-1.06	-0.99	-1.01	-1.00	-0.97
Demographics	-0.71	-0.80	-0.48	-0.85	-0.58	-0.71	-0.59	-0.77	-0.56
Inequality	-0.70	-0.69	-0.72	-0.76	-0.64	-0.67	-0.81	-0.65	-0.74
Public debt	0.31	0.30	0.29	0.37	0.27	0.31	0.33	0.29	0.33
OOP	-0.14	-0.17	-0.11	-0.19	-0.11	-0.15	-0.12	-0.15	-0.13
Labor share	0.11	0.12	0.07	0.12	0.11	0.06	0.11	0.12	0.09
Exog G spending	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.03
Interactions	-0.06	-0.02	-0.07	-0.06	-0.05	-0.06	-0.07	-0.05	-0.06

Table 11: Change in r^* , and decomposition under various parametrizations of the model. “Base” has $\bar{\sigma} = 1.5$, $\delta = 8\%$, $1/\psi = 2$, and $o = 0.3$. Row “ Δr^* 1975-2015” denotes change from 1975 to 2015 steady state. Row “ Δr^* 2015-term” denotes change from 2015 to terminal steady state. Terminal steady state is described in section 5.2. Rows 4-11 change one driver of r^* at a time between 1975 and 2015. First row shows changes in parameters. All others are held at baseline. The model is recalibrated for each column to hit targets of table 1. Values are in percentage points. “Exog G spending” refers to exogenous public spending (G) and does not include social security. Rows of individual drivers sum to “ Δr^* 1975-2015”. “Inequality” changes both permanent productivity and persistent shock variance.

The overall total change in r^* going from 1975 to 2015 and from 2015 to the terminal steady state remains

⁴³On the one hand, a higher retirement age implies a shorter expected duration in retirement and an additional year of labor income. This tends to reduce savings, putting upward pressure on r^* . On the other hand, the implied reduction in the labor income tax due to lower social security spending increases savings. Appendix A.12.4, figure 37 shows the response of r^* , labor income tax receipts, and output.

fairly constant among the different versions. The historical change, 1975 to 2015, lies between -1.87 pp and -2.3 pp. The projected change going forward, from 2015 to the terminal steady state, lies in a range of 0.3 pp to 0.72 pp. The forecast for r^* going forward therefore lies between 0.8% and 1.22%. Looking at the individual decomposition, we see that TFP growth is the number-one driver throughout. The second most important driver is either demographic change or the rise in inequality, depending on the parametrization. Our preferred interpretation is that both are of about equal importance. The remaining drivers and interactions are stable throughout.

6 Conclusion

In this paper we developed a heterogeneous-agent, OLG model with nonhomothetic preferences that includes the most important determinants of the natural rate (r^*) suggested in the literature. The model can account for a 2.2 pp decline in r^* between the 1975 and 2015 steady states. Rising income inequality is an important driver and can account for about as much of the decline in r^* as demographic change or the slowdown of productivity growth, an important finding. Public debt is the major force going against the trend, pushing r^* up. We find permanent income inequality to be of greater importance compared with inequality due to uninsurable income risk. Matching the degree of nonhomotheticity in consumption and savings behavior in the model to empirical estimates is essential for this result. We find a moderate role for OOP health expenditures in explaining the historical decline, but a sizable effect of moving to public coverage of OOP expenses. Our model predicts a reversal of the downward trend by 2030, but r^* will stabilize at the low level of 1%. Policy can have considerable impact on the level of r^* through several levers. This motivates our proposal to treat the level of the natural rate as a choice variable for policy makers. We want to give that proposal a more thorough treatment in future research.

A Appendix

A.1 Data on Interest Rates in U.S.

Figure 6: Various interest rates and trend of short-term interest rate. Details see text.

Figure 6 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.⁴⁴ The figure also shows NBER recessions.

A.2 Appendix to Sections 3 and 4

Demographics: We model a closed economy in which time is discrete, $t = 0, 1, 2, \dots$, and there are G^d household generations at any point in time. We will refer to a household of generation g of being of “age” g . The total population is large so we consider each generation to be an interval N_t^g on the real line. Total population size at time t is denoted by N_t . The transition between generations is deterministic but death is random, so that a household of generation g at time t will move to generation $g + 1$ in time $t + 1$ or die. A household belonging to generation g at time t survives into period $t + 1$ with probability $p_{g,t}$. The law of large numbers thus ensures that the mortality rate is given by $1 - p_{g,t}$. There is a predetermined “oldest” generation ($g = G^d$) in the model for which the probability to die is one, that is, $p_{G^d,t} = 0$.

⁴⁴The data are taken from [Bauer and Rudebusch \(2020\)](#). Inflation expectations are based on surveys.

Figure 7: Population growth rate and probability to survive until age 99 conditional on age 26 along the transition path. Source: US Census Bureau.

Total population N_t is the sum of all generations N_t^g :

$$N_t = \sum_g^{G^d} N_t^g$$

We define generation shares as $n_t^g \equiv \frac{N_t^g}{N_t}$. Denote by $N_t^{g,M}$ net migrant inflows of age g at t . We assume that migrants are exact replicas of the native population: this holds true for their productivity state z , the health state h and for asset holdings a . For the latter, to avoid having to assume capital inflows from foreign, we let the resources for the migrants be drawn from the pooled bequest stock within the period.

In every period, $M_t = \tilde{n}_t N_{t-1}$ individuals enter the economy, with $\tilde{n}_t$ denoting a fertility rate (including migration flows of age $g = 1$).

The law of motion for generations is

$$N_t^g = p_{g-1,t-1} N_{t-1}^{g-1} + N_t^{g,M} = \bar{p}_{g-1,t-1} N_{t-1}^{g-1} \quad \text{if } g > 1 \quad (14)$$

and

$$N_t^1 = M_t \quad (15)$$

Here $1 - \bar{p}_{g-1,t-1}$ is the mortality rate adjusted for migration flows.

Population growth rate is given by n_t such that $N_{t+1} = (1 + n_t)N_t$. Note that given survival probabilities $\{\bar{p}_{g,t}\}$, a given $\tilde{n}_t$ implies a n_t and vice versa. Demographic change can happen due to an increase or decline in $\tilde{n}$ and/or a change in survival probabilities $\{\bar{p}_{g,t}\}$, where the latter includes changes in migration flows.

As stated in the main text, we use population data from the Census Intercensal Population Estimates provided by the US Census Bureau (U.S. Census Bureau, 2016, 2017). The data consists of head counts of individuals in each age group for each year. From this we calculate migration adjusted survival probabilities $\bar{p}_{j,t}$ and population growth rates n_t . This means that these values include migration flows. Since these flows are substantial in the U.S. over the time period we consider, we view this as a benefit. We calculate individual survival probabilities $p_{j,t}$ assuming $p_{j,t} = \min(\bar{p}_{j,t}, 1)$ if $g > 16$, and set $p_{j,t} = 1$ otherwise.

Figure 7 plots population growth rate n_t and conditional survival probability s_{74} along the transition path. An important fact to keep in mind is that given our choices on model age, particularly that individuals enter the model at biological age 26, the baby boom appears in the model at around the year 1975.

For population growth rate n in the steady-state exercise we use five year averages around 1975 and 2015, respectively. Population growth rate is of population in the range of model age, that is, 26 to 99. Survival probabilities $\{\bar{p}_g\}$ and $\{p_g\}$ in the steady state are from year 1975 and 2015, respectively.

For the transition path, we fix migration adjusted survival probabilities $\bar{p}_g$ and population growth rate n_t at 2060, the last year for which projections are available. This means that population dynamics before this point in time will still be propagated beyond the year 2060, leading to echo effect in several endogenous variables in the model, including r^* . To facilitate convergence of the transition path algorithm we smooth population growth rate implied by the Census data with an MA(7) smoother.

Utility functions: Disutility from labor supply is parametrized by two parameters, ϕ and χ , see equation (9). We set $\phi = 0.5$ under the baseline, implying a Frisch elasticity of 2. This is at the lower bound of the macro estimates for the Frisch elasticity, see Whalen and Reichling (2017). χ is included in the numerical calibration routine to ensure that average output is equal to 1.

In the bequest utility function, equation (10), we set $b_{1,g} = \sigma_g \forall g$. Parameters b_0 and k_b are included in the numerical calibration routine. They are important to match the targets of bequest-to-output ratio as well as a moment from the distribution of bequests. We target a bequest-to-output ratio of 5%, between the 2% reported in Hendricks (2001) and 8% in Alvarado, Garbinti and Piketty (2017). Regarding the distributional moment, De Nardi, French and Jones (2010) target that 30% of households have bequests below 6.75% of average output, that is, the share of population with virtually no bequest. This moment is hard to hit in our model, and it applies to single households. Our households are not restricted to be single, so we use the statistic reported in De Nardi, French and Jones (2010) that includes all households and target that 30% of the population have bequests below 18.5%.

Bequests are distributed equally among households of model age 9 and 39 (biological age 35 and 65).

Public Policy variables: We use data that supplements Congressional Budget Office (2020). Debt is "Debt Held by the Public". Outlays are "Total" - "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 . For steady-state values, we take 5 year averages around 1975 and 2015. In the data "Social Security" + "Federal Civilian and Military Retirement" amount to 5.0% and 5.8% in 1975 and 2015, respectively.

Inequality: Inequality data is taken from [Piketty and Saez \(2003\)](#). We use wage income shares, table B2 from the updated file February 2020, from Emmanuel Saez' webpage⁴⁵. For 1975 we use a 5 year average around 1975, for 2015 we use the latest available data point, 2011.

Figure 8: Deterministic productivity profile $h_z^g(e_i)$ by permanent productivity type

Productivity: The deterministic age profile $h_z^g(e_i)$ follows a quadratic polynomial as in [Guvenen et al. \(2019\)](#):

$$h_z^g(e_i) = \frac{x_1(g-1) - (g-1)^2}{x_2} e_i^{lc}, \quad g \in [1, \bar{g}-1]$$

In the case of $e_i = e_H$, peak productivity is reached at biological age 54, implying $x_1 = 56$ and $x_2 = 2 \times 28^2$. For $e_i = e_L$ we have $x_1 = 46$ and $x_2 = 2 \times 23^2$, implying a peak at biological age 49. The peak increase at age 54 corresponds to the maximum of the estimated profile in [Guvenen et al. \(2015\)](#), figure C.20. The peak is reached earlier for the median worker, as can be seen in their figure C.21.

e_i^{lc} governs the size of the increase in the productivity profile. We take the maximum increase in life-cycle earnings from [Guvenen et al. \(2015\)](#), figure 11. For the low permanent type we take a 47% maximum income growth over the life-cycle ($e_L^{lc} = 0.47$), slightly below the median 60%. The high type increase we set to 110% ($e_H^{lc} = 1.1$). The profiles are plotted in figure 8. For lack of data, we keep the deterministic age profile unchanged between steady states and along the transition path.

For the transition path of e_H and σ_e , we feed in a linear transition between the steady-state values, where the transition starts in 1965 and ends in 2015 in the baseline exercise. We choose 1965 as starting date of the transition for e_H because it turns out to make inequality dynamics in the model to align better with

⁴⁵See <https://eml.berkeley.edu/~saez/>, accessed 8/5/2020.

Figure 9: Transition path of e_H and σ_ϵ . Units for e_H correspond to % difference to e_L on log scale.

the data. We smooth around starting and end date using a MA(5). Figure 9 shows the transition path of e_H and σ_ϵ .

Production: We want to start by pointing out that the setup described in the main text is equivalent to one in which there is a representative intermediate good producer, and a layer of retailers that buy the intermediate good, and differentiate it into several goods whose consumption translate into utils following a CES utility function.

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

$$Y_t = \left[\int_0^1 y_{it}^{\frac{\theta_t-1}{\theta_t}} \right]^{\frac{\theta_t}{\theta_t-1}} \quad (16)$$

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

The production side implies following firm FOCs:

$$w_t = \frac{1}{\mu_t} (1 - \alpha) \frac{y_{it}}{l_{it}} \quad (17)$$

$$r_t^k = \frac{1}{\mu_t} \alpha \frac{y_{it}}{k_{it}} \quad (18)$$

where r_t^k the real user cost of capital.

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

$$K_{t+1} = (1 - \delta)K_t + I_t \quad (19)$$

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

The markup is important to match a specific capital-to-output ratio in the steady state, compare table

1 and text. The values for the $\frac{K}{Y}$ target used in the literature vary substantially. [Straub \(2019\)](#) targets $\frac{K}{Y} = 3.5$, [Rachel and Summers \(2019\)](#) $\frac{K}{Y} = 2.5$. The model directly links $\frac{K}{Y}$ and $\frac{I}{Y}$ in steady state.⁴⁶ [Eggertsson, Mehrotra and Robbins \(2019\)](#) target an investment-to-output ratio of $\frac{I}{Y} = 16.8\%$, implying $\frac{K}{Y} = 1.7$ under our baseline calibration. We conduct sensitivity analysis with respect to depreciation rate which moves around the investment-to-output ratio.

The value for the depreciation rate δ is set to 0.08 in the baseline calibration, which is in the range of estimates of [Nadiri and Prucha \(1996\)](#). For the TFP growth rate γ we rely on estimates from [Fernald \(2014\)](#), covering the full postwar period.⁴⁷ The TFP growth rate γ is set to 0.7% in the 2015 steady state. The transition path of γ is shown in figure 15. For the transition path we smooth the original series using a weighted linear least-squares regression with span 0.4.

No-arbitrage on asset supply side implies the standard condition:

$$1 + r = 1 + r^k - \delta \quad (20)$$

Labor share: We use data from the NIPA tables from [BEA \(2020\)](#). Our labor share measures attributes all of proprietors' income to labor. For steady-state values we use 5 year averages around 1975 and 2015, respectively.

Profits: Profits are distributed according to the following rule: within permanent type, profits are distributed according to the age profile of capital income reported in [Piketty, Saez and Zucman \(2018\)](#). Figure 14 shows the distribution of capital income by age. Profits monotonically increase in age with a slight decline after age 85. Among permanent types, we split according to bottom 90% to top 10% distribution of capital income (excluding net interest payments) in the data. Households of type e_L receive 35.7% of post-tax capital income (excluding interest payment).

OOP expense data: 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 $h_m^{z_i, \delta}$ (details below). Households with high permanent productivity type e_H have the highest profile. Households of e_L type and that enter retirement with low ζ_i realizations have a low $h_m^{z_i, \delta}$ profile, the remaining e_L households an intermediate profile. Figure 10 shows $\exp(h_m^{z_i, \delta})$ if in bad health state, by productivity state z_i .⁴⁸

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

⁴⁶The link comes from the law of motion for capital, equation (19). In the steady state, it holds that $\frac{I}{Y} = [(1 + \gamma)(1 + n) - (1 - \delta)] \frac{K}{Y}$.

⁴⁷An updated TFP series is provided at <https://www.frbfs.org/economic-research/indicators-data/total-factor-productivity-tfp/>

⁴⁸Note that the productivity state z_i combines realizations of both e_i and ζ_i .

Figure 10: Profile of out-of-pocket (OOP) health expenditure under bad health state. “High e type” refers to all households with $e_i = e_H$. “Low e type low” refers to households with $e_i = e_L$, and $\zeta_i \in \{\zeta_1, \zeta_2\}$, that is, the two low realizations of ζ . “High e type high” refers to households with $e_i = e_L$, and $\zeta_i \in \{\zeta_3, \zeta_4\}$.

expenditure by about 13.5 times. About a third of retired households finds themselves in bad health condition.

For the transition path, we assume a linear transition between 1975 and 2015, smoothed using MA(5). The transition path is plotted in figure 15.

$h_m^{z,g}$ profiles: For high type we take quintile 5 estimate of the deterministic profile in Kopecky and Koreshkova (2014). For the low 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 empirical estimates of old-age poverty levels. Thus, as baseline we use quintiles 1 and 2. Kopecky and Koreshkova (2014) mention that Q1 and Q2 are qualitatively different than others, likely because individuals in these groups rely more on government-provided services. We can then think of low 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 equation (21), 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.

$$\Lambda_{hh} = \begin{bmatrix} 0.7918 & 0.2082 \\ 0.3610 & 0.6390 \end{bmatrix} \quad (21)$$

OOP price: 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. For steady-state values we use 5 y averages around 1975 and 2015, respectively.

Figure 11: Income tax system and social security benefit payment as function of income. “ T_{ss} ” refers to social security benefit before tax, “ xi ” to social security benefits after tax (ξ). Derived using parametrization of 2015 steady state.

Government: Labor income tax function $T_y(\cdot)$ is given by equation (22), implying post-tax labor income given by (23):

$$T_y(y) \equiv A \left[\frac{y}{A} - \lambda_0 \left(\frac{y}{A} \right)^{1-\lambda_1} \right] \quad (22)$$

$$wz_l - T_y(wz_l) = A_t \lambda_0 \left(\frac{wz_l}{A} \right)^{1-\lambda_1} \quad (23)$$

Here, λ_1 regulates the progressivity of tax rates, with higher values indicating a more progressive tax system.

We set $\lambda_1 = 0.181$ based on [Heathcote, Storesletten and Violante \(2017\)](#). λ_0 will be endogenously determined in the model to ensure that the government’s budget constraint, equation (8), is satisfied.

Next we turn to the social security payment ξ . The social security system is a piece-wise linear schedule

following the Old Age and Survivor Insurance component of the U.S. social security system (Huggett and Ventura (2000), De Nardi and Yang (2014)). Pre-tax social security payment T^{SS} is given by

$$T^{SS}(x, W) = \phi \left\{ 0.9 \min(x, 0.2W) + 0.32 \max[0, \min(x, 1.24W) - 0.2W] + \right. \quad (24)$$

$$\left. + 0.15 \max[0, \min(x, 2.47W) - 1.24W] \right\} \quad (25)$$

where x is the transfer base and W is average labor income in the economy, ϕ a parameter than allows to shift social security benefit of all agents proportionally. Ideally, we would want to track average lifetime labor income for each agent and use it as the transfer base, but this is computationally infeasible. We instead select $x_{i,g,t} = w_t z_{i,g,t} \bar{l}_t$ as an approximation. $\bar{l}$ is the average labor supply in the economy. There is no income risk in retirement: agents keep the realization of z from the last period before retirement until they die. Making the benefit formula dependent on average labor income W implies that payments are indexed to productivity growth, again ensuring the existence of a balanced growth path.

In accordance with the U.S. system, social security benefits are taxed. Post-tax social security benefits ζ are calculated using the tax formula (22), adjusted for payroll taxes

$$\zeta(z, W) = T^{SS} - \left[T^{SS} - A\lambda_0 \left(\frac{T^{SS}}{A} \right)^{(1-\lambda_1)} - \tau_{pr} T^{SS} \right] \quad (26)$$

The τ_{pr} -term accounts for the fact that retirees do not pay payroll tax. Figure 11b plots the social security benefit function, both pre- and post-tax, as a function of the benefit base $wz\bar{l}$. We set $\tau_{pr} = 0.1$ under the baseline calibration, following Kitao (2014). Figure 13 shows the effective replacement rate implied by the social security system. The effective replacement rate is defined as benefit paid relative to benefit base ($wz\bar{l}$).

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.

We will calibrate the model to some debt-to-output ratio DtoY. It holds that

$$DtoY_t \equiv \frac{B_{t+1}}{Y_t} = (1 + \gamma_t)(1 + n_t) \frac{\tilde{B}_{t+1}}{\tilde{Y}_t} \quad (27)$$

In the steady state, the debt-to-output ratio is constant. Figure 12 shows the debt-to-output ratio along the transition path. Note that the debt-to-output ratio is exogenous in the baseline experiments. Figure 15 in appendix A.2 shows the exogenous path for G .

Figure 12: Public debt along the transition path, in percent of output. The debt-to-output ratio is exogenous in the baseline experiment.

As mentioned in the text, as a default we feed in a path for $\lambda_{0,t}$ along the transition path and let G_t adjust to ensure that the government's budget constraint, equation (8), holds. We construct the path as follows: to start, note that in steady state, we set G to the empirically observed value and let λ_0 adjust. We calculate the value for λ_0 at four steady states: initial, terminal, 1975 and 2015. Then, we assume a linear transition between these values. The only exception is after 2015: we assume that λ_0 stays constant until 2050. The idea is that we do not take a stance on changes in the tax code from the present to 2050. After that date, we assume that budget consolidation happens through adjusting taxation, using the CBO's 2050 forecast of G_t as terminal value (Congressional Budget Office, 2019).

This strategy is used for two reasons, one economical and the other numerical: first, we want to avoid counterfactual fluctuations in marginal tax rates that would arise if we let λ_0 adjust along the transition path. Given endogenous labor in our model, a change in λ_0 induces a behavioral response, so we want to avoid too much fluctuations in marginal tax rates implied by changes in λ_0 . Tax rates only move slowly in the data, due to the legislative process connected to adjusting tax rates. In a similar vein, our decision to hold λ_0 constant from 2015 to 2050 amounts to the assumption that tax rates stay constant for the foreseeable future, an assumption we prefer in lack of clear evidence to the contrary. We are more comfortable in having G_t adjust on the transition path since the only impact in the model is the respective resource cost. The second reason is numerical: we found that convergence is slow and fragile if we feed in a path for G_t and let $\lambda_{0,t}$ adjust on the transition path.

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.

The transfer T^M is given as:

$$T^M = 0, m = M \quad (28)$$

$$\text{if: } (1 + \tau_c)\underline{c} + m \leq (1 + (1 - \tau_k)r)[a + beq] + \xi + T^M + (1 - \tau_d)d \quad (29)$$

$$T^M = 0, m = \max\left(\underline{M}, \min\left(M, (1 + (1 - \tau_k)r)[a + beq] + \xi + T^M + (1 - \tau_d)d - (1 + \tau_c)\underline{c}\right)\right) \quad (30)$$

$$\text{if: } (1 + \tau_c)\underline{c} + \underline{M} < (1 + (1 - \tau_k)r)[a + beq] + \xi + T^M + (1 - \tau_d)d \leq (1 + \tau_c)\underline{c} + M \quad (31)$$

$$T^M = \max\left(0, (1 + \tau_c)\underline{c} + \underline{M} - ((1 + (1 - \tau_k)r)[a + beq] + \xi + T^M + (1 - \tau_d)d)\right), m = \underline{M} \quad (32)$$

$$\text{if: } (1 + (1 - \tau_k)r)[a + beq] + \xi + T + (1 - \tau_d)d \leq (1 + \tau_c)\underline{c} + \underline{M} \quad (33)$$

where $\underline{M}$ is a basic OOP medical expense. What these conditions are saying is that transfer T^M is means-tested and only available when $a + beq = 0$ 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.

Figure 13: Effective replacement rate implied by social security system.

Figure 14: Share of profits going to each age group. Source: [Piketty, Saez and Zucman \(2018\)](#).

Figure 15: Transition path for shock variables under baseline transition-path experiment.

A.3 Market Clearing Conditions

In the equilibrium, households maximize their problem (2) or (4), firms maximize profits, the generations evolve according to (14), and the government's debt evolves according to (8), taking the (path of) real interest rates and real wages as given. Let $\lambda(z, h, a, g)$ be a measure of states, this measure sums to one.

The market clearing condition for assets is:

$$B_{t+1} + K_{t+1} = \sum_g \sum_{z,h,a} g_{a,t}(z, h, a, g) \lambda_t(z, h, a, g), \quad (34)$$

where $g_a(\cdot)$ is the policy function for savings. Asset demand (or equivalently savings supply) is the aggregated savings of households. Asset supply (equivalently savings demand) in this economy consists of government debt and capital. The real interest rate is determined by the intersection of the asset demand and the asset supply curve. Factors that have an impact on supply and demand of asset will therefore be of great importance for our study.

Goods market clearing is given by:

$$Y_t = \sum_g \sum_{z,h,a} [g_{c,t}(z, h, a, g) + p_m m_g(z, h)] \lambda_t(z, h, a, g) + I_t + G_t, \quad (35)$$

where $g_c(\cdot)$ is the policy function for consumption. The goods market clearing is standard once we take into account out-of-pocket medical expenses.

The market clearing condition for effective labor is:

$$L_t^D = \sum_{g \leq \bar{g}} \sum_{z,h,a} z g_{l,t}(z, h, a, g) \lambda_t(z, h, a, g), \quad (36)$$

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.

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

$$\lambda_{t+1}(z', h', a', g') = \sum_{z,h,a,g \text{ s.t. } a'=g_a(z,h,a,g)} P(z', h', g' | z, h, g) \lambda_t(z, h, a, g) \quad (37)$$

Distributed profits have to fulfill

$$\sum_{z,h,a,g} d_t(z, h, a, g) \lambda_t(z, h, a, g) = \left(1 - \frac{1}{\mu_t}\right) Y_t \quad (38)$$

where $d_t(z, h, a, g)$ are profits going to individual of corresponding state. The right-hand-side is simply the sum of firm profits.

Bequest clearing implies

$$\sum beq_{t+1} = (1 - \tau_{beq}) \sum \lambda_t^{death} g_a(z, h, a, g) \quad (39)$$

where λ_t^{death} is the distribution of the deceased.

A.4 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 $\{r_t, w_t, d_t\}$, distribution $\{\lambda_t\}_{t=0}^{\infty}$, paths for policy functions $\{g_{a,t}(z, h, a, g), g_{c,t}(z, h, a, g), g_{l,t}(z, h, a, g)\}_{t=0}^{\infty}$ that jointly satisfy:

1. Household first-order conditions (41), (43), (44)
2. Household budget constraints (3) and (5)
3. Firm first-order conditions (18) and (17)
4. Government budget constraint (8)
5. Asset market clearing (34)
6. Goods market clearing (35)
7. Labor market clearing (36)
8. Law of motion for the measure $\lambda(\cdot)$ (37)
9. Distribution of profits (38)
10. Bequest clearing (39)

A.5 Additional Material for Section 4.7

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.⁴⁹ 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 adjusted for measurement error, $\hat{y}_{i,g}^{post}$, and age dummies:

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

⁴⁹We simulate data for 7,500 model households and find the estimate of ϕ_{PI} to be reasonably robust to increasing the number of simulations.

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.

Choice of weighting matrix W: The numerical calibration routine includes a weighting matrix W , see equation (13). We can change the “importance” of each target by setting an appropriate W . We can also take into account that the moments are expressed in different units. It turns out that it is not easy to hit all seven targets of table 1. We set the weighting matrix to prioritize the natural rate r^* , top 10% labor income share and elasticity of consumption out of permanent income, ϕ_{PI} . We think this is appropriate given we put a focus on the inequality results in the main text. We experiment with several different choices of the weighting matrix, using guess and verify, to bring the model moments close to the data moments.

A.6 Homothetic Economy

In section 5.1.2 we list the assumption made under version “Fully Homothetic” of the model. Here we give additional details.

Assumption (2) states that the social security system is linear. This means that pre-tax social security payment changes from T^{SS} , equation (25), to $T^{SS,lin}$

$$T^{SS,lin} = \phi_{lin} w z \bar{l} \quad (40)$$

where ϕ_{lin} is the replacement rate. We set ϕ_{lin} equal to the effective replacement rate at the average $w z \bar{l}$ in the baseline model (that is, $\phi_{lin} = \frac{\bar{w z l}}{w z l}$). We set $\tau_{pr} = 0$.

Assumption (5) states that labor supply is exogenous. In this case every household in the economy is endowed with $\bar{l}$ labor, which we calibrate to hit average output equal to 1.

Assumption (7) states homogeneous age-productivity profile. In this case we have $h_z^g(e_i) = h_z^g(e_L) \forall g, i$, where the calibration of $h_z^g(e_L)$ is left unchanged to baseline (see section 4.3).

Assumption (8) states that OOP medical expenses are proportional to permanent productivity type. Here we set the deterministic profile $h_m^{z,g}$ equal to the Q1 series of [Kopecky and Koreshkova \(2014\)](#) for all households of permanent type e_L . For households of permanent type e_H it holds that $h_m^{z(e_H),g} = e_H h_m^{z(e_L),g}$. The shock h is left unchanged to baseline. In this economy, we also eliminate the consumption floor and means-tested transfer T^M .

A.7 First-order Conditions

Worker first-order conditions (FOCs):

$$\begin{aligned} \frac{1}{1 + \tau_c} u_c(c, l, g) &= \beta p_{g,t+1} \mathbb{E}_{(z', h') | (z, h)} [V'_a(z', h', a', g) + V'_a(z', h', a', g + 1)] \\ &\quad + \beta(1 - p_{g,t+1}) b_0 (1 - \tau'_{beq}) \frac{1}{o'} \left(k_b + (1 - \tau'_{beq}) \frac{a'}{o'} \right)^{-b_{1,g}} \end{aligned} \quad (41)$$

$$V_a(z, h, a, g) = \frac{1}{1 + \tau_c} (1 + (1 - \tau_k)r) u_c(c, l, g) \quad (42)$$

$$u_l(c, l, g) = -\frac{1}{1 + \tau_c} \left[A \lambda_0 (1 - \lambda_1) \left(\frac{wz l}{A} \right)^{-\lambda_1} \frac{zw}{A} \right] u_c(c, l, g) \quad (43)$$

Retiree FOCs:

$$\begin{aligned} \frac{1}{1 + \tau_c} u_c(c, g) &= \beta p_{g,t+1} \mathbb{E}_{(z', h') | (z, h)} [V'_a(z', h', a', g) + V'_a(z', h', a', g + 1)] \\ &\quad + \beta(1 - p_{g,t+1}) b_0 (1 - \tau'_{beq}) \frac{1}{o'} \left(k_b + (1 - \tau'_{beq}) \frac{a'}{o'} \right)^{-b_{1,g}} \end{aligned} \quad (44)$$

Utility functions first derivatives:

$$\frac{\partial u}{\partial c_t} = c_t^{-\sigma_g} \left(\frac{1}{o_t} \right)^{1-\sigma_g} \quad (45)$$

$$\frac{\partial u}{\partial l_t} = -\chi l_t^\psi \quad (46)$$

$$\frac{\partial v}{\partial a} = b_0 (1 - \tau_t^{beq}) \frac{1}{o_t} \left[k_b + (1 - \tau_t^{beq}) \frac{a_t}{o_t} \right]^{-b_{1,g}} \quad (47)$$

A.8 Extension with Exogenous International Capital Flows

A potential driver for the decline of r^* often mentioned in the literature is inflows of foreign capital into the U.S. (Bernanke, 2005; Rachel and Smith, 2015; Rachel and Summers, 2019). In our framework, the natural rate balances the demand and supply for savings. A greater supply of savings due to net capital inflows exhibits additional downward pressure on r^* . The Net International Investment Position (NIIP) of the U.S., which is the accumulation of net capital inflows, declined from 4.3% of GDP in 1976 to -40.9% of GDP in 2015 (see figure 16).

In order to assess the potential impact of capital flows on r^* within our framework, we extend our model

Figure 16: U.S. Net International Investment Position (NIIP) from 1976 to 2020, in percent of GDP. Source: BEA

and assume the availability of an additional, exogenous, stock of savings.⁵⁰ We match the size of this additional stock of savings to the NIIP of the U.S.

Formally, assume there is an additional entity active in the economy, interpreted as a foreign country, with the following budget constraint:

$$A_{t+1}^* = (1 + (1 - \tau_k)r_t)A_t^* + Y_t^* \quad (48)$$

Here, the asterisks denotes variables of the foreign country, with Y^* denoting net income inflows into the domestic economy from foreign.

For simplicity of exposition, we assume that foreign can only invest in government bonds. This is without loss of generality, given that a no-arbitrage condition holds between investing in capital and government bonds. The government budget constraint has to be adjusted to take into account foreign investment. Equation (8) is then replaced by

$$B_{t+1}^D + B_{t+1}^* = (1 + r_t)[B_t^D + B_t^*] + (\Xi_t - \Lambda_t). \quad (49)$$

Here B_t^D and B_t^* denote bond holdings by domestic households and foreign, respectively. We continue to denote total government debt by $B_t = B_t^D + B_t^*$. The budget constraint of domestic households stays unchanged, except for the updated notation on domestic households' bond holdings, b_t^D .

The condition for asset market clearing, equation (34) in the baseline model, is now given by

$$K_{t+1} + B_{t+1} = \sum g_{a,t} \lambda_t + A_{t+1}^*. \quad (50)$$

⁵⁰ A full treatment of foreign capital flows, modeling asset demand and supply in foreign countries in a similar manner as in our baseline model, is beyond the scope of this paper. A model in this spirit has been developed in Barany, Coeurdacier and Guibaud (2018).

Market clearing condition for goods, (35) in the baseline, is now given by

$$Y_t = \sum [g_{c,t} + p_m m_g] \lambda_t + I_t + G_t - Y_t^* \quad (51)$$

The asset market clearing condition shows that an additional stock of savings is available in this economy. If, starting from a given equilibrium situation, A^* increases, the natural rate has to decline in order to discourage savings by households and incentivise firms to use more capital in production, in order for a new equilibrium to set in.

A.8.1 Foreign Capital Flows: Numerical Analysis

We set A^* to the NIIP of the U.S. We recalibrate the model to match the same targets as in table 12.

If we decompose the change in the natural rate from 1975 to 2015 as in section 5, we get the result shown in table 12.

	Δr^*	%
Total r^* change	-2.54	100.0
TFP Growth	-0.97	38.1
Inequality	-0.75	29.7
Demographics	-0.67	26.4
Public debt	0.34	-13.3
NIIP	-0.32	12.7
OOP	-0.15	6.0
Labor Share	0.09	-3.6
Exog G Spending	0.03	-1.2
Interactions	-0.14	5.3

Table 12: Change in r^* going from 1975 to 2015 steady state and decomposition in model with NIIP.

The total decline of r^* increases to 2.54 pp. The change in the U.S. NIIP between 1975 and 2015 adds another 0.32 pp to the decline in r^* , relative to a situation in which A^* had stayed constant. According to this exercise, foreign capital flows are a non-negligible driver when it comes to the decline in the natural rate, but quantitatively not as important as the top three explanations, TFP growth, inequality, and demographic change. Interestingly, the size of the effect coming from NIIP is very close to the impact of public debt, although with opposite sign.

A.9 Decomposing Demographic Change

The driver “Demographics” in the baseline decomposition – table 4 of section 5.1.1 – is the result both of a change in fertility $\bar{n}_t$ as well as survival probabilities $\{\bar{p}_g\}$.⁵¹ We now examine the impact coming from either change in isolation. To do this, we use the same approach as for our baseline steady-state results: in one case, we only change fertility $\bar{n}_t$ while holding all other driving variables, including survival probabilities $\{\bar{p}_g\}$, constant. This gives us the individual impact of the change in fertility on r^* . In the

⁵¹As discussed in appendix A.2, for a population in a steady state the fertility rate $\bar{n}_t$ implies a particular population growth rate n , and vice versa. As in all other results reported, we match the population growth rate n_t in a particular year and calculate the implied fertility rate $\bar{n}_t$ that supports this population growth rate in the steady state.

other case, we solely change survival probabilities $\{\bar{p}_g\}$, but hold all other drivers constant. This gives the individual impact of changes in survival probabilities on r^* , denoted “Demo Mortality.” Table 13 shows the results.⁵²

	Δr^*	%
Demographics	-0.71	100
Demo Fertility	-0.25	35.2
Demo Mortality	-0.49	69.0
Interaction	+0.03	-4.2

Table 13: Change in r^* going from 1975 to 2015 steady state due to demographic change and sub-components. Values in column Δr^* are in percentage points. See text for details on decomposition.

We see that changes in survival probabilities account for 69% of the total change and are therefore more important than changes in fertility over the time period of 1975 to 2015. Interactions are very small in magnitude.

Changes in fertility $\bar{n}_t$ and survival probabilities $\{\bar{p}_g\}$ impact population dynamics, as can be seen from equations (14) and (15) in appendix A.2, which govern the law of motion of the population distribution. In addition, there is a direct impact of changes in survival probabilities on households’ decisions, seen from Euler Equations (41) and (44) in appendix A.7. A change in survival probabilities implies a change in life expectancy, which is directly relevant for households’ optimal consumption and savings, as well as labor supply decisions. A change in population dynamics (and thus fertility) affects households in our model only indirectly, through general equilibrium effects - for instance, due to changes in interest rates, wages, and the tax and transfer system.

The fundamental difference in the way changes in distinct determinants of population dynamics affect total savings motivates an alternative decomposition of demographic change: first, an effect coming from a change in survival probabilities $\{p_g\}$ that enter households’ Euler Equations, which we refer to as the “behavioral” effect. Second, an effect coming only due to changes in the composition of the population (from fertility $\bar{n}_t$ and survival probabilities $\{\bar{p}_g\}$). We refer to this effect as “compositional.”⁵³

Table 4 shows the results for this decomposition. For the results in “Demo Behavior” we only change $\{p_g\}$, while holding $\{\bar{p}_g\}$, $\bar{n}_t$, and all other driving variables constant. For the results in “Demo Composition” we solely change $\bar{n}_t$ and $\{\bar{p}_g\}$.⁵⁴

We see that most of the change is coming from the “behavioral” component, with a -0.6 pp change or 84.5% of the total effect of demographic factors on r^* . The expectation of a higher life expectancy leads households to accumulate more savings when they are still in the labor force, because they expect to

⁵²For the results in table 13 we change both $\{p_g\}$ and $\{\bar{p}_g\}$ simultaneously, as in the main text. The distinction between the two variables is explained in appendix A.2.

⁵³The notation developed in appendix A.2 gives rise to the interpretation that holding $\{p_{g,t}\}$ constant but changing $\{\bar{p}_{g,t}\}$ means that all change in the law of motion of the population distribution, apart from $\bar{n}_t$, is due to migration flows. While we prefer the interpretation that we hold constant the *perceived* survival probabilities of households, while actual survival rates and immigration rates change according to the data, both interpretations imply the same result given the way we treat migration in our model.

⁵⁴Note that we always let prices change so that all markets clear (see appendix A.3), so general equilibrium effects will be present in either case.

	Δr^*	%
Demographics	-0.71	100
Demo Composition	-0.11	15.5
Demo Behavior	-0.60	84.5
Interaction	0	0

Table 14: Change in r^* going from 1975 to 2015 steady state due to demographic change and sub-components. Values in column Δr^* are in percentage points. See text for details on decomposition.

spend a longer period in retirement. Changes in r^* due to dynamics in the population distribution are of much smaller magnitude.⁵⁵

A.10 The Role of Out-of-Pocket Medical Expenses

PANEL A: Baseline		
	Δr^*	%
Total r^* change	-2.161	100.0
OOP	-0.14	6.6
ϕ_{PI} w/ OOP	0.70	
ϕ_{PI} w/o OOP	0.67	
PANEL B: Homothetic Economy		
	Δr^*	%
Total r^* change	-2.194	100.0
OOP	-0.31	14.0
PANEL C: w/o OOP		
	Δr^*	%
Total	-2.002	100.0
TFP growth	-1.03	51.6
Demographics	-0.69	34.6
Inequality	-0.70	35.2
Public debt	0.29	-14.6
Labor share	0.12	-6.0
Exog gov spending	0.03	-1.4
Interactions	0.02	-0.9

Table 15: Change in r^* going from 1975 to 2015 steady state, and decomposition in baseline model and model without out-of-pocket (OOP) medical expenses. Panel A reproduces results from table 4, together with ϕ_{PI} estimate from baseline model and version with OOP expenses set to zero, respectively. Panel B shows results from a fully homothetic economy as in section 5.1.2. Panel C shows decomposition in an economy in which OOP expenses are set to zero. Rows change one driver of r^* at a time. Column 2 shows the change in r^* in pp; column 3 shows the relative contribution of each row to the total. “Exog gov spending” refers to exogenous public spending (G) and does not include social security. Rows sum to “Total.” “Inequality” changes both permanent productivity and persistent shock variance.

One of the contributions of this paper is to add OOP medical expenses to the analysis of drivers of the r^* decline. Regarding the related literature, De Nardi, French and Jones (2010) ask whether health expenses in late stages of life can explain the slow decline of asset holdings of people of old age. Kopecky and Koreshkova (2014) explores the role of nursing home expenses on savings. Both find health expenses to be an important determinant of savings. We draw heavily from these papers for modeling and parametrizing

⁵⁵This result is interesting in relation to Mian, Straub and Sufi (2021): using a shift-share analysis, they find that demographic factors are not as important as inequality in explaining changes in r^* . The benefit of our model relative to a shift-share analysis is that we can isolate the impact of a rise in life expectancy on r^* . The importance of changes in life expectancy for r^* that we find can reconcile our result of a large effect of demographic change with their findings.

our OOP health-expense process. These papers do not focus on the natural rate, and to the best of our knowledge we are the first to explore the contribution of health expenses to the r^* decline.

The impact of OOP medical expenses can come in two ways: first, directly through the increase in the relative price of OOP medical expenses p_m between 1975 and 2015; second, because including OOP medical expenses in our model changes the impact of other drivers.

Panel A of table 15 again shows baseline results of the steady-state analysis. As mentioned, the impact of more than doubling the price of medical goods and services p_m has a moderately negative impact on the natural rate: a decline of 0.14 pp, or 6.6% of the total, going from the 1975 steady state to the 2015 steady state. OOP expenses as a share of output rise from 0.54% to 1.5% over the same time frame.

It is noteworthy that the conclusion is somewhat different if we look at the fully homothetic economy. The impact of OOP medical expenses increases by more than 100% both in absolute and relative terms. OOP expenses account for a fall of r^* by 0.31 pp, or 14% of the total (see panel B in table 15).

An interesting result is obtained when looking at the ϕ_{PI} estimate in the baseline model once we set OOP medical expenses to zero and recalibrate all but σ_{slope} to hit the usual targets.⁵⁶ Contrary to our expectations, ϕ_{PI} falls from the calibrated value of 0.7 in the baseline model to 0.67 (see table 15, panel A). This means including OOP medical expenses leads to a smaller increase in savings out of permanent income for e_H households, at least relative to households with a low permanent productivity draw e_L . This is because according to the data used to calibrate the OOP-medical-expense process, OOP expenses are quite high for the lower quintiles in the early stages of retirement, particularly relative to permanent income differences. It is only in later stages of retirement that permanent-income-rich households face considerably higher OOP medical expenses (see figure 10 in appendix A.2). The former effect seems to dominate, inducing e_L -type agents to increase their savings rates by more than e_H -type agents, at least in relative terms. A result from our analysis is thus that OOP medical expenses alone are not a good candidate to explain why permanent-income-rich households have higher savings rates in their working years – that is, why the ϕ_{PI} estimate is considerably below 1.

Finally, does the exclusion of OOP medical expenses have an impact on the other drivers? The answer is *no*. Panel C in table 15 shows the results if we set OOP medical expenses to zero but otherwise calibrate the model to the same targets as in the baseline (the target for OOP is obviously excluded). Comparing the numbers with table 4, we see that the total r^* decline is of similar magnitude and the various drivers account for about the same shares as before.

The conclusion of this section is that, at least in the way we implement OOP medical expenditures, they do not seem to be a major force in explaining the decline in r^* in models rich enough to account for other key data, such as elasticity ϕ_{PI} .

---

⁵⁶When we set OOP medical expenses to zero we also eliminate the consumption floor.

Figure 17: Response of r^* to a change in means-tested transfer T^M . Labor income tax-level parameter λ_0 adjusts to clear the government's budget. All other values are at the baseline 2015 steady state. The means test is different from that under the baseline calibration: all out-of-pocket (OOP) medical expenses up to level $\bar{m}$ are covered. $\bar{m}$ is set to the value of full OOP expenses of e_L type. r^* is in percentages. T^M is in percentage of full medical expenses of e_L type.

Policy: Figure 17 shows the effect of increased coverage of medical expenses on r^* . The horizontal axis shows the coverage of OOP medical expenses by the public. We assume that only basic OOP medical expenses are covered: the maximum coverage is the full amount paid by e_L types.⁵⁷ We see that were the public to cover all basic procedures that are now paid OOP (100% coverage), the natural rate would go up by about 35 bp. Income tax receipts to cover the expenses would have to go up by about 1% (see figure 36 in appendix A.12.4).

A.11 Productivity Growth and Nonhomothetic Consumption Preferences

An assumption made in the baseline scenario is that the term o_t in the utility function grows at the same rate as TFP. This assumption ensures that a steady state along a balanced growth path exists. If we were to keep o_t constant indefinitely, savings rates of households would converge to one. This is a direct consequence of working with our nonhomothetic consumption preferences, as they imply that savings rates increase in income. However, we can still assume that there is an intermediate period in which o_t grows at a different rate than A_t . In this section we explore the consequences of such assumptions.

The interpretation we want to give to such an experiment is the following: nonhomothetic preferences imply that richer households have higher savings rates. But productivity growth usually implies that many households get richer as time goes by. We can even make the extreme assumption that over the years, inequality doesn't change at all but the income of each household grows at the same positive rate. Then the model would imply that as households earn higher and higher income, they also save relatively more and more of that income. This implies higher aggregate savings rates and downward pressure on r^* . While we are not claiming that this is the whole story when it comes to savings rates as economies

⁵⁷ An interpretation of this assumption is that the public does not cover expenses that can be seen as luxury goods consumed by e_H types, such as single rooms in nursing homes. Setting public OOP coverage to 100% covers about two-thirds of all OOP medical expenses in the baseline calibration.

develop and grow, we consider this to be a plausible story with the potential to account for some of the decline in r^* .

	Baseline		Constant o_t term	
	Δr^* pp	%Total	Δr^* pp	%Total
Total	-2.161	100.0	-3.467	100.0
Inequality	-0.70	32.4	-0.41	11.8
o term	n.a.	n.a.	-1.61	43.5
o term: fix inequality	n.a.	n.a.	-1.15	33.2
Interactions	-0.06	3.0	-0.15	4.3

Table 16: Change in r^* going from 1975 to 2015 steady state, and selected decomposition in baseline versions of the model and under assumption of constant o_t term. Baseline refers to fully nonhomothetic version. Rows below “Total” change one driver of r^* at a time. Columns 2 and 4 show the change in r^* in percentage points (pp); columns 3 and 5 show the relative contribution of each row to the total. Other drivers are the same as in table 4 for both versions of the model. “Inequality” changes both permanent productivity and persistent shock variance.

We make the extreme assumption that the o_t term doesn’t grow between 1975 and 2015 while TFP term A_t has an average growth rate of 1.1%. This implies that the o term is about 55.5% higher in 1975 relative to baseline. This might be counterintuitive, but since the model variables are expressed in effective-labor terms, holding the o term constant actually implies that we have to *change* the value of o when we move from the 2015 steady state to the 1975 steady state. Then, holding all else constant, for a given amount of income, households will want to consume *more* when they are young if o_t is higher. In the long run the term o_t grows with TFP growth rate γ , as in the previous experiments. Table 16 shows the implications for our results.

We see that the implied decline in r^* increases from 2.16 pp under the baseline to 3.48 pp in the version with a constant o_t term. Just the effect of not changing the o_t term amounts to a decline in r^* of 1.61 pp, or 43.5% of the total decline. The contribution of changing the variance of both permanent and persistent productivity declines from -0.7 pp in the baseline to -0.41 pp in version “Constant o_t term,” or from 32.4% to 11.8% in relative terms. This result is partially explained by the fact that just the effect of not changing the o_t term already lowers the top 10% labor-income share to 30.5% from the 34.8% in the 2015 steady state. This is due to the response of endogenous labor supply. Consequently, variances of the productivity terms don’t have to change that much to account for the full change in inequality between the 1975 and 2015 steady states. Note that the contributions of all other drivers stay unchanged relative to table 4 because the 2015 steady state stays unchanged.

If we adjust the productivity term z together with the o_t term, such that the observed top 10% share stays *unchanged*, what is the decline in r^* ? This is another way to look at the question of how much an unchanged o_t term itself can explain the fall in r^* . In the previous paragraph, we held the variance of productivity terms constant. Now we hold *observed* inequality (measured by the top 10% labor-income share) constant. The o_t term itself changes by the same amount in both cases. In this experiment, the o_t term only accounts for a 1.15 pp decline in r^* , which is still 33.2% of the total observed decline. If we measure the contribution of inequality from this point, we are back to a sizable contribution of inequality of 0.87 pp, or 25%.

Our interpretation of this result is that according to our model, and under the assumption that utility parameters do *not* adjust with TFP growth, a sizeable share of the decline in r^* can be accounted for by everyone's enrichment due to TFP growth. Here, as everyone has more income, all households want to consume more when retired, thereby increasing savings and pushing down r^* . This result holds true even if we hold observed inequality constant. An implication of this is that as economies develop and all households' incomes in an economy are lifted, r^* is bound to decline. Note that according to this reasoning, in a world with integrated capital markets, global development has important consequences for r^* in every country, a point worth further exploring in future research.

A.12 Additional Tables and Figures

A.12.1 Steady State

Figure 18: Lorenz curves from model in baseline 2015 steady state and data. Horizontal axes show percentile in the distribution, vertical axes show shares. Panel top-left shows Lorenz curve for assets. Panel top-right shows Lorenz curve for labor income and social security transfer. Here, labor income is pre-tax, except for a proportional tax to finance the social security transfer. Panel bottom-left shows Lorenz curve for total pre-tax income and social security transfer. Panel bottom-right shows Lorenz curve for post-tax total income and post-tax social security transfer. The data is taken from [Piketty, Saez and Zucman \(2018\)](#).

Figure 18 shows selected Lorenz curves. Note that the only targeted moment across all four panels is the top 10% labor income share.⁵⁸ Figure 19 shows model-implied profiles for assets, consumption and pre-tax labor and social security income. Figure 20 plots life-cycle asset profiles against data and shows asset profiles by percentiles.⁵⁹

⁵⁸As can be seen from the top-right panel of figure 18, the top 10% labor-income share and the top 10% share of the combination of labor income and social security transfer overlap.

⁵⁹The data are from the Survey of Consumer Finances' 2016 and 1989 waves ([Board of Governors of the Federal Reserve System, 2014, 2020](#)).

(a) Assets

(b) Consumption

(c) Pre-tax labor and social security income

Figure 19: Life-cycle profiles in baseline 2015 steady state, by permanent productivity type. Panel (a) shows asset holdings in effective-labor terms, $\tilde{a}$. Panel (b) shows consumption in effective-labor terms, $\tilde{c}$. Panel (c) shows pre-tax labor and social security income in effective-labor terms, $\tilde{w}z\hat{l} + T^{SS}$. High permanent productivity type is denoted by “L”, low permanent productivity type denoted by “H”. All profiles are averages over respective type. The horizontal axes show age, the vertical axes values relative to average income in per capita terms, which is normalized to 1.

(a) Mean assets

(b) Share of assets by age

(c) Asset profiles by percentile

Figure 20: Various life-cycle asset profiles in baseline 2015 steady state, model and data. Panel (a) shows model-implied mean asset holdings, and profile of average asset holdings from the data. The profiles are set relative to output. Note that each profile has a separate vertical axis. Panel (b) shows share of total assets held by each age group. Panel (c) shows model-implied percentiles of asset holdings, in effective-labor terms, for each age group. The data are from the 2016 Survey of Consumer Finances (panel (a) and (b)) and 1989 SCF (panel (b)).

A.12.2 Transition Path

Figure 21: Transition path r^* model, and average and range of estimates of r^* from literature.

Figure 22: Transition path r^* model and trend of short-term interest rate. The latter is U.S. Federal Funds Rate minus inflation expectations smoothed using MA(120). Further details on the data in appendix A.1.

Figure 21 shows the model-implied path for r^* together with average and range of estimates from the literature. The range represents estimates of r^* using seven different approaches from five papers. All these approaches have in common that they decompose the real rate into a cycle and a trend component, $r_t = r_t^* + r_t^c$, where r_t^* is a stochastic trend and r_t^c a cycle component. They then either use time-series methods or estimate a simple New Keynesian model to recover r_t^* . In figure 21 we compare the trend component r_t^* to our model-implied natural rate r^* . [Bauer and Rudebusch \(2020\)](#) gives further details on estimation approach and further sources.

Figure 23: Transition path for selected variables under baseline. OOP stands for out-of-pocket. Data is from [Piketty and Saez \(2003\)](#).

Figure 24: Transition path for selected variables under baseline. Path “measured” refers to Solow residual in a growth accounting decomposition. “gamma*” refers to adjusted TFP growth rate $\gamma^* = (1 + \gamma)^{1-\alpha} - 1$. “SocSec/Y” refers to aggregate social security transfers relative to output. “Share retired” refers to share of population that is retired.

Figure 25: Transition path for selected variables under baseline.

A.12.3 Scenario Analysis

Figure 26: Transition path for population growth rate under three different scenarios: baseline, “low r^* ,” and “high r^* .”

Figure 27: Transition path for public debt under three different scenarios: baseline, "low r^* ," and "high r^* ."

Figure 28: Transition path for earnings process parameters e_H (left) and σ_e (right) under three different scenarios: baseline, "low r^* ," and "high r^* ."

Figure 29: Transition path for TFP growth rate γ under three different scenarios: baseline, "low r^* " and "high r^* ".

Figure 30: Implied top 10% labor income share under different scenarios: baseline, "low r^* ," and "high r^* ."

Figure 31: Implied path for exogenous government spending G to output (left) and total tax revenues to output (right) under different scenarios: baseline, "low r^* ," and "high r^* ."

Figure 32: Implied path for r^* under different scenarios for path of population growth rate after year 2020: baseline, low, and high. The path for population growth rate can be seen in figure 26. All other parameters are held the same as in baseline transition-path analysis.

Figure 33: Implied path for r^* under different scenarios for path of public debt after year 2020: baseline, low, and high. The path for public debt can be seen in figure 27. All other parameters are held the same as in baseline transition-path analysis.

Figure 34: Implied path for r^* under different scenarios for path of earnings process parameters e_H and σ_e after year 2020: baseline, low, and high. The path for earnings process parameters can be seen in figure 28. All other parameters are held the same as in baseline transition-path analysis.

Figure 35: Implied path for r^* under different scenarios for path of TFP growth rate γ after year 2020: baseline, low, and high. The path for TFP growth rate γ can be seen in figure 29. All other parameters are held the same as in baseline transition-path analysis.

A.12.4 Policy Analysis

Figure 36: Response of labor income tax receipts T_{inc}/Y and output in response to change in coverage of OOP medical expenses. Labor income tax level parameter λ_0 adjusts to clear the government's budget. All other values are at the baseline 2015 steady state. $\frac{T_{inc}}{Y}$ in percent, output in levels.

Figure 37: Left panel shows response of r^* to change in mandatory retirement age $\bar{g}$ at the steady state. Right panel shows response of labor income tax receipts relative to output, and output. Values for r^* and $\frac{T_{inc}}{Y}$ are in percent, $\bar{g}$ shows model age (biological age minus 26), output is in levels of output per effective labor ($\bar{y}$). Tax parameter λ_0 adjusts to clear the government's budget. All other parameters are held at the baseline 2015 steady-state value.

A.13 Further Results on Policy

Policy Parameter (x)	Baseline Value	$\frac{dr^*}{dx} \times 100$	$\frac{d\bar{y}}{dx}$	$\frac{dr^*}{dT^{inc}} \times 100$
Consumption tax τ_c	0.05	-3.25	0.17	6.86
Profit tax τ_{corp}	0.25	-1.25	0.17	10.27
Capital income tax τ_k	0.40	2.30	-0.16	-24.82
Estate tax τ_{beq}	0.10	0.47	0.00	-10.34
Tax progressivity λ_1	0.181	4.51	-0.82	85.99
Basic income T	0.00	-0.52	-0.56	-0.77
Public coverage of med expenses T^M	0.00	0.38	-0.02	42.57
Retirement age $\bar{g} + 25$	66.00	0.79	0.75	-1.42
Replacement rate ϕ	1.00	1.66	-0.11	34.09
Public debt B	74.20	0.64	-0.01	-25.90
Exo public spending G	13.50	10.11	0.19	11.52

Table 17: Elasticity of r^* in response to change in policy parameter at baseline 2015 steady state. Numbers in forth column denote response of r^* in basis points to change in parameter by one percentage point, except "Retirement age" which denotes response of r^* in basis points to change in retirement age of one year. Column 5 shows response of r^* in basis points if income tax revenue is increased by 1 pp and revenues used to change respective policy parameter to stay budget neutral.

The exercise underlying table 17 changes income tax parameter λ_0 to balance the government's budget. Another way to look at the impact of changes in policy parameters on r^* is by normalizing the effect on the income tax revenue. In table 5 we hold the change in the income tax constant, at 1 percent. The numbers can thus be interpreted as the following: if the government were to lower income tax revenue by 1 pp, and use the receipts to adjust a respective policy parameter to keep the budget balanced, what would be the change in r^* (for a negative change in income tax revenue we have to add a minus sign to the values in column 5).

The results show that introducing spending on public coverage of medical expenses financed by a 1 percentage point increase in the income tax raises the natural rate by 42.57 bp. It is thus a more "cost-effective" way of raising the natural rate then, for example, increasing the replacement rate for social security.

Marginal and Average Tax Rates: Labor income tax is given by (equation (22) setting $A = 1$):

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

Marginal tax rate (MTR) is given by

$$T'_y(y) = 1 - \lambda_0(1 - \lambda_1)y^{-\lambda_1}$$

Average tax rate (ATR) is given by

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

Figure 38 plots MTR and ATR functions at the baseline 2015 steady state.

Figure 38: Marginal tax rate (MTR) and average tax rate (ATR) as a function of pre-tax labor income wzl at baseline 2015 steady state. Vertical dotted lines show 50th and 90th pre-tax labor income percentile under baseline 2015 steady state.

A.14 Overview Related Literature

Table 18 shows results across several important papers in the r^* literature. The acronyms refer to:⁶⁰

- EMR: Eggertsson, Mehrotra and Robbins (2019)
- GJL: Gagnon, Johannsen and Lopez-Salido (2016)
- AR: Auclert and Rognlie (2018)
- RS I: Rachel and Summers (2019) model 1
- RS II: Rachel and Summers (2019) model 2
- S: Straub (2019)
- PP: Platzer and Peruffo

⁶⁰Rachel and Summers (2019) contains two models, one with an OLG structure but without intragenerational inequality and one with idiosyncratic income risk but no OLG structure. We refer to them as models 1 and 2, respectively.

	EMR	GJL	AR	RS I^1	RS II^1	S	PP
Time							
Start	1970	1980	1980	1970	1970	1970	1975
End	2015	2016	2013	2017	2017	2014	2015
Change in r^*							
Total	-4.02	-1.7*	?	-1.6	?	?	-2.16
Growth	-1.90	-0.45*	n.a.	-1.8	n.a.	?	-1.00
Demographics	-3.66	-1.25	n.a.	-1.8	n.a.	?	-0.71
Permanent Inequality	n.a.	n.a.	?	n.a.	n.a.	-1.00	-0.58
Idiosyncratic Inequality	n.a.	n.a.	-0.54	n.a.	-0.7	?	-0.10
Public Debt	2.11	n.a.	?	0.8	0.4	?	0.31
Exog Gov Spending	?	n.a.	?	0.0	?	?	0.03
Labor Share	-0.52	n.a.	?	?	?	?	0.11
OOP health	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.	-0.14
Interactions	0.26	?	?	-1.1	?	?	-0.08

Table 18: Overview of change in r^* across several papers in literature. Rows “Start” and “End” refer to start- and end-date of window over which the r^* change is calculated. The unit is percentage points. A cell with “n.a.” means that this element is not in the model and thus the number can’t be reported. A cell with “?” means the model does in principle allow one to produce this number but it is not reported in the paper. 1: for all advanced economies. * ... Results from model with TFP growth. **For definitions of the acronyms, see previous page. See respective papers for further details on respective model.

A.15 Algorithm

The algorithm consists of an inner and an outer loop. For the inner loop, we take aggregate variables, that is, r^* , Y , λ_0 , average labor supply $\bar{l}$ and $\sum beq$ as given and solve for the stationary distribution. We use the endogenous grid method from [Carroll \(2006\)](#) to derive the policy functions. The outer loop finds the aggregate variables that clear all markets. Here we employ a version of Broyden’s algorithm to solve the sizeable model in reasonable time.

References

- Abel, Andrew B.** 2003. "The Effects of a Baby Boom on Stock Prices and Capital Accumulation in the Presence of Social Security." *Econometrica*, 71(2): 551–578.
- 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 Matt Rognlie.** 2020. "Wealth and Demographics in the 21st Century."
- 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.
- Bernanke, Ben S.** 2005. "The Global Saving Glut and the US Current Account Deficit." Remarks at the Homer Jones Lecture, St. Louis, MO.
- BLS, Bureau of Labor Statistics, U.S. Department of Labor.** 2020. "Consumer Price Index (CU) Tables." Online Edition.
- Board of Governors of the Federal Reserve System.** 2014. "1989 Survey of Consumer Finances."
- Board of Governors of the Federal Reserve System.** 2020. "2016 Survey of Consumer Finances."
- Boppart, Timo, and Per Krusell.** 2020. "Labor Supply in the Past, Present, and Future: A Balanced-Growth Perspective." *Journal of Political Economy*, 128(1): 118–157.
- 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.
- Carroll, Christopher D.** 2006. "The method of endogenous gridpoints for solving dynamic stochastic optimization problems." *Economics Letters*, 91(3): 312–320.

- Carvalho, Carlos, Andrea Ferrero, and Fernanda Nechio.** 2016. "Demographics and real interest rates: inspecting the mechanism." Federal Reserve Bank of San Francisco Working Paper Series 2016-5.
- 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.
- 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.
- Fernald, John G.** 2014. "A quarterly, utilization-adjusted series on total factor productivity."
- Gagnon, Etienne, Benjamin K. Johannsen, and J. David Lopez-Salido.** 2016. "Understanding the New Normal : The Role of Demographics." Board of Governors of the Federal Reserve System (U.S.) Finance and Economics Discussion Series 2016-080.
- Geanakoplos, Author-Name: John, Michael Magill, and Martine Quinzii.** 2004. "Demography and the Long-Run Predictability of the Stock Market." *Brookings Papers on Economic Activity*, 35(1): 241–326.
- Gourinchas, Pierre-Olivier, and Jonathan A. Parker.** 2003. "Consumption Over the Life Cycle." *Econometrica*, 70(1): 47–89.
- Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song.** 2015. "What Do Data on Millions of U.S. Workers Reveal about Life-Cycle Earnings Risk?" National Bureau of Economic Research, Inc NBER Working Papers 20913.

- 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.
- Hendricks, Lutz.** 2001. "Bequests and Retirement Wealth in the United States." mimeo.
- Ho, Paul.** 2019. "Estimating the Effects of Demographics on Interest Rates: A Robust Bayesian Perspective."
- 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.
- Ikeda, Daisuke, and Masashi Saito.** 2014. "The effects of demographic changes on the real interest rate in Japan." *Japan and the World Economy*, 32(C): 37–48.
- Jones, Callum.** 2018. "Aging, Secular Stagnation and the Business Cycle." International Monetary Fund IMF Working Papers 2018/067.
- Kitao, Sagiri.** 2014. "Sustainable Social Security: Four Options." *Review of Economic Dynamics*, 17(4): 756–779.
- 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.
- Krueger, Dirk, and Alexander Ludwig.** 2007. "On the consequences of demographic change for rates of returns to capital, and the distribution of wealth and welfare." *Journal of Monetary Economics*, 54(1): 49–87.
- Lisack, Noémie, Rana Sajedi, and Gregory Thwaites.** 2019. "Population Ageing and the Macroeconomy."
- Mehrotra, Neil.** 2018. "Implications of Low Productivity Growth for Debt Sustainability." mimeo.
- Mian, Atif, Ludwig Straub, and Amir Sufi.** 2021. "What explains the decline in r^* ? Rising income inequality versus demographic shifts." Proceedings of the 2021 Jackson Hole Symposium.
- Nadiri, M Ishaq, and Ingmar Prucha.** 1996. "Estimation of the Depreciation Rate of Physical and RD Capital in the U.S. Total Manufacturing Sector." *Economic Inquiry*, 34(1): 43–56.

- Papetti, Andrea.** 2019. "Demographics and the natural real interest rate: historical and projected paths for the euro area." European Central Bank Working Paper Series 2258.
- 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.
- Rachel, Lukasz, and Lawrence Summers.** 2019. "On Secular Stagnation in the Industrialized World." National Bureau of Economic Research, Inc NBER Working Papers 26198.
- Rachel, Lukasz, and Thomas Smith.** 2015. "Secular Drivers of the Global Real Interest Rate." Bank of England Staff Working Paper No. 571.
- Sajedi, Rana, and Gregory Thwaites.** 2016. "Why Are Real Interest Rates So Low? The Role of the Relative Price of Investment Goods." *IMF Economic Review*, 64(4): 635–659.
- Straub, Ludwig.** 2019. "Consumption, Savings, and the Distribution of Permanent Income." Revise and resubmit at Econometrica.
- United Nations, Department of Economic and Social Affairs, Population Division.** 2019. "World Population Prospects 2019." United Nations, Department of Economic and Social Affairs, Population Division Online Edition. Rev. 1.
- U.S. Census Bureau, Population Division.** 2016. "National Intercensal Tables 1900 - 2016 (various editions)." Online Edition.
- U.S. Census Bureau, Population Division.** 2017. "National Population Projections 2017." Online Edition.
- Whalen, Charles, and Felix Reichling.** 2017. "Estimates of the Frisch Elasticity of Labor Supply: A Review." *Eastern Economic Journal*, 43(1): 37–42.

PUBLICATIONS

Secular Drivers of the Natural Rate of Interest in the United States: A Quantitative Evaluation
Working Paper No. **WP/22/30**