On Natural Interest Rate Volatility*

Edouard Challe[†] Mykhailo Matvieiev[‡]

May 16, 2024

Abstract

Episodes of low natural interest rates, even transitory, pose a challenge to monetary policy, by possibly causing the effective lower bound (ELB) on the policy rate to bind. Those episodes are more likely to occur not only when the natural rate is low on average but also when fluctuations around its average level are large. We study the responsiveness of the natural interest rate to structural aggregate shocks affecting the aggregate supply of and demand for savings. Using a quantitative overlapping-generations model, we trace back this responsiveness to the slopes of aggregate savings supply and demand curves and argue that both curves have likely flattened over the past four decades in the US This implies a greater sensitivity of the natural interest rate to structural shocks affecting the supply of and demand for aggregate savings – making it more likely, all else equal, that it fall into negative territory.

Keywords: Natural interest rate; Intertemporal income effects; Overlapping-generations.
JEL Codes: E21, E25, E32.

*We are grateful to two anonymous referees for their constructive comments, the editors for their guidance, and Jesus Bueren for his contribution to the early stages of this project. This work was supported by Agence Nationale de la Recherche (grants ANR-20-CE26-0018-01 and ANR-17-EURE-0020) and by the Excellence Initiative of Aix-Marseille University – A*MIDEX.

[†]Paris School of Economics and CEPR; Email: edouard.challe@psemail.eu.

[‡]Aix-Marseille University, CNRS, EHESS, Centrale Marseille, AMSE, France; Email: mykhailo.matvieiev@univ-amu.fr.

1 Introduction

In this paper, we study the responsiveness of the "natural interest rate" – the equilibrium real interest rate that would prevail in a hypothetical flexible-price economy – to structural aggregate shocks affecting the aggregate supply of and demand for savings. Our motivation for doing so is the widespread concern that, in a low-interest-rate environment, central banks may not be able to effectively track this natural interest rate because of a binding effective lower bound (ELB) on the policy rate – thereby de facto forcing the central bank into an inefficiently contractionary monetary policy when the natural interest rate plunges. While much of the recent literature (reviewed below) has focused on the downward trend in riskless interest rates worldwide and the risk it poses to monetary policy implementation, we stress that an equal concern should be the *volatility* of the natural interest rates around trend, as the latter also – and on an equal footing with the trend – determines how frequently and deeply the natural rate is likely to fall below the ELB.

There are two broad reasons why the volatility of the natural interest rate may evolve over time. First, the frequency, amplitude (or even type) of aggregate shocks hitting the economy may change. Second, the propagation of those shocks to the demand for and supply of aggregate savings, and ultimately to the equilibrium interest rate, may change. In this paper, we entirely focus on the second question and ask: What are the structural factors ultimately determining the response of the natural rate to a given set of aggregate shocks, and how have these factors likely evolved over the past few decades? Accordingly, we consider generic structural shocks of normalised size shifting the supply of aggregate savings (namely, discount-factor shocks) or the demand for such savings (aggregate productivity shocks). In richer environments, other shifters of aggregate savings demand and supply could be considered, but those two are the natural place to start.[1]

We theoretically identify, and tentatively quantify, three factors having likely *increased* the responsiveness of the natural interest rate to underlying structural shocks affecting the supply of and demand for aggregate savings in the US economy over the past forty years: (i) increased out-of-pocket (OOP) health spending in old age (as a fraction of social-security income); (ii) decreased good market competition; and (iii) increased public debt. Increased health spending in old age magnifies the intertemporal income effects associated with changes in the expected return on assets, making households less willing to buffer these changes by adjusting savings. Formally, the aggregate savings supply curve, which expresses aggregate household savings as a function of the expected return on assets, flattens (and can, in principle, even revert). Intuitively, households facing a fall in their expected return on assets will dissave less to secure future consumption when more of their future non-asset income is already committed to health expenses. Because household savings respond less to changes in interest rates, the response of the equilibrium interest rate to aggregate shocks that shift the savings supply or demand curves is magnified.

While old-age health spending affects the shape of the savings supply curve, the last two factors (goods market competition and public debt) affect the shape of the savings *demand* curve – i.e. the responsiveness of borrowers as a whole (firms and the government) to changes in the interest rate that they are facing. First, all else equal, lesser goods market competition reduces the responsiveness of firms' demand for capital

[1] Of course, since the mid-2000s, the world economy has been hit by shocks that were larger than during the Great Moderation – from the Great Financial Crisis to COVID-19 to the Russian invasion of Ukraine. However, while the nature and size of these shocks are relatively well understood, the determinants of their propagation to interest rates are less.

to its cost – formally, it reduces the elasticity of the capital stock to the real interest rate, i.e., the capital demand curve flattens. Second, to the extent that government liabilities are less sensitive to the real interest rate than firms', a relative increase in the contribution of public debt to the overall stock of assets available to households tends to reduce the sensitivity of the total demand for savings to change in the interest rate. That is, both the sensitivity of firms' capital to the interest rate and the contribution of this capital to the overall stock of assets matter for the slope of the savings demand curve and, ultimately, for the responsiveness of the equilibrium interest rate to shifts in the savings demand and supply curves. Both factors, we argue, have pushed towards a flattening of the aggregate savings demand curve over the past forty years in the US. This also implies a greater response of the equilibrium interest rates to exogenous shifters of the aggregate savings supply or demand curves.

We clarify and measure those three forces in two steps. First, we lay out (in Section 2) a simple, tractable overlapping-generations (OLG) model with two period-lived households wherein the three ingredients outlined above (old-age health spending, imperfect competition, and public debt) are operative. Household consumption choices over their life cycle determine the supply of aggregate savings. More specifically, households receive labour earnings when active and social security in old age, complemented by the payoff on their asset portfolio. Total income in old age is partly spent inelastically on health expenses, with what is left being spent on non-durable consumption. Life-cycle saving behaviour generates intertemporal income effects in response to changes in expected asset returns, which become more potent as the share of social security income devoted to health spending increases. On the other side of the market for aggregate savings stand firms and their demand for capital, which equates the marginal product of capital with its cost, as well as the government rolling over a given stock of public debt. We derive approximate analytical expressions for the aggregate savings supply and demand curves in the vicinity of the steady state and use those expressions to uncover the underlying determinants of the slopes of these curves. Finally, we use those slopes to provide analytical formulae for the elasticities of the equilibrium interest rates to discount-factor shocks (which shift the aggregate supply of savings) and total-factor productivity shocks (which shift the aggregate demand for savings while holding supply fixed on impact). In doing so, we arrive at a tight analytical characterisation of how changes in the deep parameters of the model, ultimately reflecting changes in the key structural features of the economy, affect those elasticities. In particular, we show how greater health spending (as a function of social security income), lower competition, and greater public debt all tend towards raising the elasticities of the equilibrium interest rate with respect to both aggregate shocks under consideration.

Guided by the intuitions provided by our analytical model, we next construct (in Section 3) a fully-fledged quantitative OLG model with multiple life periods and a realistic life-cycle pattern of individual labour earnings. While the multiple-period OLG model nests the two-period model as a special case, the former aggregates the heterogeneous saving behaviour of many cohorts. We calibrate the model to two time periods: the early 1980s ("1980" for short) and the recent years ("2020"). We compute numerical aggregate savings supply and demand curves – the quantitative analogues of the analytical curves derived within the two-period model – and confirm that both have been flattening: the savings supply curve is increasing in the interest rate and has a lower slope in 2020 than in 1980, while the savings demand curve is decreasing in the interest rate and has a higher slope in 2020 than in 1980 (this is shown in Figure 3). This again implies that aggregate shocks shifting those curves have a greater impact on the equilibrium interest rate, which we confirm by computing impulse-response functions for this rate after discount factor and productivity shocks. Under our baseline calibration, the dive of the equilibrium interest rate after a positive discount-factor shock (i.e., a positive shock to the aggregate supply of savings) or a negative productivity shock (i.e., a negative shock to the aggregate demand for savings) is about 40% larger in 2020 than in 1980 (holding the size and persistence of the underlying shocks the same) – see Figure 4 for a summary of this pattern. The response of the interest rate to discount-factor shocks is also much more persistent in 2020 than in 1980. We finally run a series of counterfactual experiments in order to decompose the role of the different factors outlined above in magnifying the interest-rate response to aggregate shocks. Namely, we compare our baseline impulse-response functions for "2020" to alternative functions whereby health cost, markups, and public debt are (one by one) set to their value in "1980". Ultimately, we find a moderate role for the rise in old-age health spending in explaining the magnification, and a more substantial role for the other two factors.

Lastly, we extend our baseline model (in Section 4) to incorporate a number of features that may also have played a role in evolution of the key elasticities that we are interested in. In particular, we consider changes across the two time periods under consideration in household longevity, retirement age, population growth, retirement replacement ratio, and private debt. These factors, it turns out, do not significantly alter the main lessons that we draw from our baseline analysis (if anything, they add to the amplification of the interest-rate response relative to our baseline model without those model features.)

1.1 Literature review

Our analysis complements the recent strand of the literature that examines the trends in the *levels* of interest rates – but typically disregards the evolution of their *volatility* around the trend. That literature has solidly established that empirically, actual *riskless* real interest rates have been trending downwards worldwide, at least over the past four decades (Del Negro et al., 2019; Obstfeld, 2023) and perhaps even for centuries (Rogoff, Rossi and Schmelzing, 2022). Interestingly, this downward trend in riskless rates has not been accompanied by a similar trend in the return to capital, at least in the US – if anything, the latter has increased over the past forty years, creating an upward trend in the rate-of-return spread between safe and risky assets (Farhi and Gourio, 2018; Marx, Mojon and Velde, 2021).[2] On the side of quantitative theory, models of "secular stagnation" have attempted to rationalize these facts by tracing them back to the underlying evolution of the structural features of the economy, such as the decline in productivity and population growth, the rise in firms' market power, and an increased preference for safety (Rachel and Summers, 2019; Eggertsson, Mehrotra and Robbins, 2019; Platzer and Peruffo, 2022; Bailey et al., 2022). In particular, Farhi and Gourio (2018) and Marx, Mojon and Velde (2021) insist on the importance of changes in household attitudes towards risk in explaining the currently low value of the riskless rate relative to the return to capital. Relative to those studies, our paper focuses on the responsiveness of the interest rate to aggregate shocks – as opposed to trends in interest rates. We share with some models of secular stagnation the emphasis on firm markups – and, more precisely, the role of their evolution over time. We share with other models of secular stagnation – e.g., Del Negro et al. (2017) – a foundation for the return spread between Treasuries and capital claims based on a *convenience yield* on Treasury debt – see also Krishnamurthy and Vissing-Jorgensen (2012), Fisher (2015), Anzoategui et al. (2019) and

[2]Those studies use the updated return-to-capital series of Gomme, Ravikumar and Rupert (2011).

more recently Chiang and Zoch (2022) for a similar approach. This allows us to match the evidence of a relatively stable return to capital in the long run coupled with a dwindling return on government bonds. Finally, while we do not explicitly model a binding ELB (since we focus on the flexible-price equilibrium and the interest-rate fluctuations it implies), our model indirectly relates to the analysis of monetary-fiscal interactions in the vicinity of the ELB. For example, Campos et al. (2024) argue that in non-Ricardian economies with a demand for public debt (for consumption smoothing, in their framework), raising the public debt increases the natural interest rate and thereby creates more monetary policy space for the central bank. Our analysis qualifies this mechanism since, as we argue, greater public debt also increases the amplitude of natural interest-rate fluctuations – holding the size of the underlying aggregate shock unchanged –, in addition to raising its average level.

2 A simple model of intertemporal income effects and imperfect competition

In this section, we lay out a simple variant of the OLG model of Diamond (1965) and Barro (1974), wherein the three effects discussed above can be decomposed analytically; the quantitative model of Section 3 will be a direct generalization of this simple model. In every period, a mass 1 of households is born. Every household supplies inelastically one unit of labour to firms when "young", and lives of asset payoffs and social security when "old". Households save in the form of capital and government debt, which are imperfect substitutes because the latter incorporates a convenience yield à la Krishnamurthy and Vissing-Jorgensen (2012). We first present firms' behaviour as well as fiscal policy before turning to the households.

2.1 Firms

There is monopolistic competition in the goods market and two production layers: competitive final goods firms produce the final good out of a continuum of intermediate goods, each of which is supplied by a single firm. The production function for final goods is

$$Y_t = \left(\int_0^1 (Y_t(f))^{\frac{\theta-1}{\theta}} df \right)^{\frac{\theta}{\theta-1}} \tag{1}$$

where $Y_t(f)$ denotes the quantity of intermediate good $f \in [0, 1]$ used in production and $\theta > 1$ the cross-partial elasticity of substitution between intermediate goods. Calling $p(f)$ the price of intermediate good f in terms of the final good, the optimal demands for inputs by the final good sector are $Y_t(f) = (p_t(f))^{-\theta} Y_t$, where the $p(f)$ satisfy $\int_0^1 (p_t(f))^{1-\theta} = 1$. In symmetric equilibrium, all intermediate goods firms produce and sell to final goods firms the same quantities and at the same price so that $p_t(f) = 1$ and $Y_t(f) = Y_t$ for all (f, t).

Intermediate good firm f produces by means of the production function:

$$Y_t(f) = \left[\alpha K_t(f)^{\frac{\rho-1}{\rho}} + (1-\alpha)(A_t L_t(f))^{\frac{\rho-1}{\rho}} \right]^{\frac{\rho}{\rho-1}} \tag{2}$$

where $K_t(f)$ is capital, $L_t(f)$ is labor, $\rho > 0$ is the elasticity of substitution between the two, $A(t)$ is labour-augmenting productivity, and $\alpha \in (0, 1)$ a production parameter. We assume in this section that capital fully depreciates within the period (this will be relaxed in Section 3). Labour-augmenting productivity $A(t)$ is subject to permanent shifts, which we assume here to be one-off and unexpected (i.e., they are "MIT" shocks). That is, $A_t = A_{t-1}\mathcal{E}_{At}$, where $\mathcal{E}_{At}$ has mean 1, and the perfect-foresight dynamics prevail after the occurrence of the shock.

In what follows, we characterise the production side of the economy in terms of detrended, "intensive-form" variables, i.e. we define $k_t(f) = K_t/(L_t(f)A_{t-1})$ and $y_t(f) = Y(t)/(L_t(f)A_{t-1})$. Exploiting the symmetry of intermediate goods firm behaviour, we drop f argument and accordingly write the intensive-form aggregate production function as:

$$y_t = \left[\alpha k_t^{\frac{\rho-1}{\rho}} + (1-\alpha)\mathcal{E}_{At}^{\frac{\rho-1}{\rho}}\right]^{\frac{\rho}{\rho-1}} \tag{3}$$

Given the demand curves faced by intermediate goods firms, they charge final goods firms the markup $\mu = \theta/(\theta-1)$ over real marginal cost. Factor prices are thus distorted and lie strictly below the corresponding marginal product; namely, the gross interest rate R_t and the wage w_t are given by:

$$R_t = \frac{\alpha}{\mu}\left[\alpha + (1-\alpha)\left(\frac{k_t}{\mathcal{E}_{At}}\right)^{\frac{1-\rho}{\rho}}\right]^{\frac{1}{\rho-1}} \qquad \text{and} \qquad w_t = \frac{1-\alpha}{\mu}A_t\left[\alpha\left(\frac{k_t}{\mathcal{E}_{At}}\right)^{\frac{\rho-1}{\rho}} + 1 - \alpha\right]^{\frac{1}{\rho-1}} \tag{4}$$

As a consequence, imperfect competition also distorts factor shares. Calling $\kappa^L = w_tL_t/Y_t = w_t/(A_ty_t)$ and $\kappa^K = K_tR_t/Y_t = R_tk_t/y_t$ the labour and capital shares along a balanced growth path without aggregate shocks, we have $\kappa^L + \kappa^K = 1/\mu < 1$. Equivalently, there is a positive pure profit share $(\mu-1)/\mu$ to be distributed to the households along with labour and capital incomes.

2.2 Government

In every period t, the government issues debt, collects taxes from working households and pays out social security payments to retirees. Labour earnings are taxed at the proportional rate τ_t, resulting in a total tax collection of $\tau_t w_t$ (since the mass of newborn is constant and equal to 1). Every retiree of time t gets a social security payment ξ_t, which we assume to increase along with aggregate productivity to ensure balanced growth: $\xi_t = \bar{\xi}A_{t-1}$, where $\bar{\xi}$ ultimately determines the income replacement ratio for retirees.[3] The government budget constraint at time t is thus given by $B_t = R_tB_{t-1} + \bar{\xi}A_{t-1} - \tau_tw_t$. We assume that the government lets τ_t adjust so that detrended public debt remains at an exogenous steady-state value $\bar{B} = B_t/A_t$. The tax rate implied by this fiscal policy is:

$$\tau_t = \frac{A_t}{w_t}\left(\frac{\bar{B}R_t + \bar{\xi}}{\mathcal{E}_{At}} - \bar{B}\right) \tag{5}$$

[3] In balanced growth, the replacement ratio is $\frac{\xi_t}{w_t} = \frac{\bar{\xi}A_{t-1}/A_t}{w_t/A_t} = \frac{\bar{\xi}}{(1+g)(w/A)}$. Given w/A (determined by firms' behaviour), $\bar{\xi}$ determines ξ_t/w_t.

2.3 Households

2.3.1 Households' problem

Households are two-period lived and inelastically provide one unit of labour to intermediate-good producers. They have intertemporal utility

$$U(c_{y,t}, b_t, c_{o,t+1}; \beta_t) = u(c_{y,t} + \psi b_t) + \beta_t u(c_{o,t+1}), \tag{6}$$

where $c_{y,t}$ denotes consumption when young, b_t the holdings of government debt when young, $c_{0,t+1}$ (non-health) consumption when old, $\beta_t \geq 0$ is the subjective discount factor, and $u(\cdot)$ is CRRA with risk aversion coefficient $\sigma > 0$. The utility specification in (6) implies that holding government debt yields utility and, hence, a return discount (or "convenience yield") relative to capital claims in equilibrium. This assumption follows much of the literature on the demand for treasury debt, e.g., Krishnamurthy and Vissing-Jorgensen (2012); Fisher (2015); Anzoategui et al. (2019); Chiang and Zoch (2022). The parameter $\psi \geq 0$ controls the relative desirability of government bonds and therefore the equilibrium spread between the two returns. β_t is stationary with mean β and subject, like productivity shocks, to one-off, unexpected (by possibly persistent) shifts. After a β_t-shock has occurred, the perfect-foresight dynamics prevail.

Following Eggertsson, Mehrotra and Robbins (2019), we assume that working households are rebated firms' pure profits, which they get along with labour earnings whilst working.[4] Since the pure profit share is θ^{-1}, every households gets a transfer Y_t/θ and thus the budget constraints of young households is:

$$c_{y,t} + \underbrace{b_t + k^s_{t+1}}_{a^s_{t+1}} = (1 - \tau_t) w_t + Y_t/\theta, \tag{7}$$

where k^s_{t+1} denotes a young household's holdings of capital claims. The household's financial wealth at the end of time t is thus $a^s_{t+1} = b_t + k^s_{t+1}$, but the two wealth components do not yield the same returns since government debt yields utility; accordingly, call R^k_{t+1} the return on capital claims and R^b_{t+1} that on government bonds.

The retirees of date t get a social security payment $\xi_t > 0$ and the portfolio payoff $b_{t-1}R^b_t + k^s_t R^k_t$. Some of this income is spent inelastically on health spending $\tilde{h}_{0,t}$, and the rest is spent on non-health consumption $c_{0,t}$.[5] Consequently, the budget constraint of date-$t+1$ retirees is given by:

$$c_{o,t+1} + \tilde{h}_{o,t+1} = b_t R^b_{t+1} + k^s_{t+1} R^k_{t+1} + \xi_{t+1}$$

Health spending grows with output in a way that ensures balanced growth, i.e. $\tilde{h}_{0,t} = h_0 A_{t-1}$, where the parameter $h_0 \in [0, \bar{\xi})$ ultimately determines the extent of households' health spending. Using the

[4] Eggertsson, Mehrotra and Robbins (2019) assume that rents are distributed to working households proportionally to labour income, and we shall make the same assumption in Section 3. This is immaterial in the present section since all households receiving rents are symmetric.

[5] Our modelling of health spending as entirely inelastic is consistent with the notion that most old-age OOP expenditures are non-discretionary (McInerney et al., 2022). For example, Medicare Part B contributions (which we are going to use to calibrate the model in Section 3) are regulated and hence not chosen by households. Note also that because health consumption is not a decision, it need not feature in the utility functional (though it could.)

expression for ξ_t, we can write retirees' old-age consumption as:

$$c_{o,t+1} = b_t R^b_{t+1} + k^s_{t+1} R^k_{t+1} + \delta A_t \tag{8}$$

where $\delta \equiv \bar{\xi} - h_o \geq 0$ measures the (detrended) size of old-age non-asset income available for non-health consumption.

2.3.2 Aggregate savings

Under our maintained assumption of perfect foresight after an aggregate shock has occurred, the optimal demands for capital claims and government debt must satisfy, respectively:

$$u'(c_{y,t} + \psi b_t) = \beta u'(c^e_{o,t+1}) R^k_{t+1} \tag{9}$$

and

$$(1-\psi)\, u'(c_{y,t} + \psi b_t) = \beta_t u'(c^e_{o,t+1}) R^b_{t+1} \tag{10}$$

Combining (9) and (10), we observe that

$$\psi = \frac{R^k_{t+1} - R^b_{t+1}}{R^k_{t+1}},$$

so that the parameter ψ directly controls the equilibrium spread between the two asset returns. Exploiting the fact that $R^b_{t+1}/R^k_{t+1} = 1 - \psi$, we can then combine equation (7)-(8) to express the household's intertemporal budget constraint as follows:

$$c_{y,t} + \frac{c_{o,t+1}}{R^k_{t+1}} + \psi b_t = (1-\tau_t) w_t + \frac{Y_t}{\theta} + \frac{\delta A_t}{R^k_{t+1}} \tag{11}$$

The latter expression reveals how a change in the return on capital R^k_{t+1}, holding the spread ψ constant, shifts the budget set of the households and, ultimately, their consumption possibilities. On the one hand, when R^k_{t+1} falls, the discounted cost of old-age consumption (from the perspective of a young household) increases. On the other hand, provided that $\delta > 0$, the present value of income also increases. Crucially, the latter effect is scaled by δA_t, i.e., retirement benefits net of health expenses. In a low-δ economy (say, because health expenses are high), the present value of income is less responsive to change in interest rates than in a high-δ economy, and therefore the first effect (on the cost of future consumption) is relatively stronger.

Under our assumed fiscal policy of constant detrended steady-state debt, and given the fact that the holders of government bonds are symmetric, market clearing for government bonds requires that $b_t = \bar{B} A_t$. Then, using equations (7), (9) and (11), we find that (detrended) individual household savings are:

$$\tilde{a}^s_{t+1} \equiv \frac{a^s_{t+1}}{A_t} = s_t \Omega_t - (1-s_t) \frac{\delta}{R^k_{t+1}} + \psi \bar{B}, \tag{12}$$

where s_t denotes the fraction of households' present value of income (the right-hand side of equation (11)) devoted to savings and is given by:

$$s_t = \left[1 + \beta_t^{-\frac{1}{\sigma}} (R_{t+1}^k)^{\frac{\sigma-1}{\sigma}}\right]^{-1} \tag{13}$$

while Ω_t is a household's detrended first-period disposable income:

$$\Omega_t = \frac{(1-\tau_t)w_t + Y_t/\theta}{A_t} \tag{14}$$

Equations (12)-(13) define a "savings supply curve" summarizing the dependence of workers' end-of-first-period savings $\tilde{a}_{t+1}$ on the expected return on capital claims in the next period, R_{t+1}^k. In the $(R^k, \tilde{a})$-plane, this curve is shifted by the underlying structural shocks, namely to preferences β_t (which move s_t) and to productivity $\mathcal{E}_{At}$ (which move Ω_t). Equations (12)-(13) show that the return on capital R_{t+1}^k affects household savings via two channels: through the saving rate s_t, and through the δ/R_{t+1}^k ratio. When $\delta = 0$, the competition between the intertemporal income and substitution effects on savings only work through s_t, and we recover the traditional result that an increase in R_{t+1}^k raises $\tilde{a}_{t+1}$ (i.e., intertemporal substitution effects dominate) if and only if $\sigma < 1$. When $\delta = 0$ and in addition $\sigma = 1$, so that the intertemporal income and substitution effects of R_{t+1}^k exactly offset each other, then the savings supply curve is flat – i.e. insensitive to R_{t+1}^k. However, even in the case where $\sigma = 1$, if $\delta > 0$, then the curve is upward-sloping, and all the more so that δ is large. That is, a large value of δ *reinforces* intertemporal substitution effects, while a low value of δ *weakens* them. We argue below that δ has fallen over the past few decades, thereby weakening intertemporal substitution and flattening the savings supply curve. In general equilibrium, this implies a stronger equilibrium response of R_{t+1}^k to the underlying structural aggregate shocks.

2.4 Equilibrium

2.4.1 Aggregate dynamics

As in Diamond (1965), the aggregate dynamics of the model is derived from the bonds market-clearing condition at time t. The total demand for bonds (by firms and the government) is $K_{t+1} + \bar{B}$. The total supply of bonds from the households is a_{t+1}. Equating the two, dividing by A_t and rearranging using equations (3) (5), (12) and (13), the aggregate dynamics of the model can be summarised by the following difference equation:

$$\left(k_{t+1} + (1-\psi)\bar{B}\right)\left[1 + \beta_t^{-\frac{1}{\sigma}} (R_{t+1}^k)^{\frac{\sigma-1}{\sigma}}\right] + \delta\left(\beta_t R_{t+1}^k\right)^{-\frac{1}{\sigma}}$$
$$= \bar{B} + \frac{w_t}{A_t} - \frac{(1-\psi)R_t^k \bar{B} + \bar{\xi}}{\mathcal{E}_{At}} + \left[\alpha\left(\frac{k_t}{\mathcal{E}_{At}}\right)^{\frac{\rho-1}{\rho}} + (1-\alpha)\right]^{\frac{\rho}{\rho-1}} \tag{15}$$

The left-hand side of (15) features k_{t+1} and R_{t+1}^k as endogenous variables. However, R_{t+1}^k is a function of $(k_{t+1}, \mathcal{E}_{At+1})$ (see equation (4)), while $\mathcal{E}_{At+1}$ plays no role under our assumption of perfect foresight after the time-t MIT shocks (i.e. $\mathcal{E}_{At+1} = 1$).[6] Therefore, the left-hand side of (15) implicitly defines a function $F(k_{t+1}, \beta_t)$. On the other hand, the right-hand side of (15) features w_t/A_t and R_t^k, both of which are

[6]Put differently, in equation (15) R_{t+1}^k is the ex-ante return on capital in the next period, while R_t^k is the realised, ex-post at time t. Both are conditional on the date-t aggregate shock.

functions of $(k_t, \mathcal{E}_{At})$ only. Therefore, the right-hand side of (15) defines a function $H(k_t, \mathcal{E}_{At})$. Eventually, we can rewrite equation (15) more compactly as

$$F(k_{t+1}, \beta_t) = H(k_t, \mathcal{E}_{At}) \tag{16}$$

The steady-state value of k_t, if it exists, is the solution k^* to $F(k^*, \beta) = H(k^*, 1)$. If the steady state is unique, a dynamic equilibrium exists in the vicinity of that steady state provided that the dynamics $k_{t+1} = F^{-1}[H(k_t, 1), \beta]$ converges towards k^* for an initial condition k_0 sufficiently close to k^*. For plausible calibrations like those adopted in Section 3 below – suitably adjusted for the different number of life periods across the two models –, a unique stationary equilibrium exists, i.e. there is exactly one steady state around which local convergence of k_t towards k^* occurs. This is depicted in Figure 1.

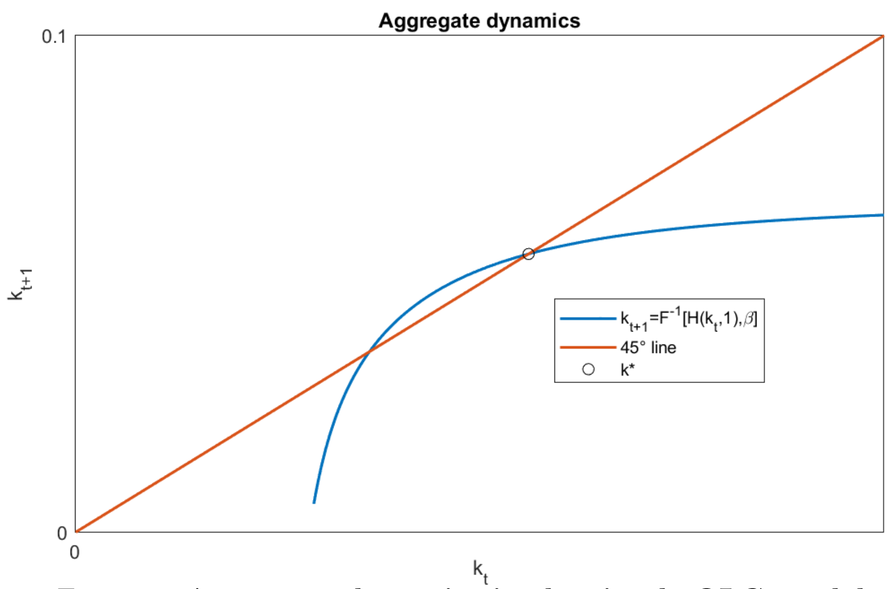

Figure 1: **Aggregate dynamics in the simple OLG model**

2.4.2 Savings demand curve

Denote by $a^d_t = k_t + \bar{B}$ the total demand for funds by the firms and the government. Log-linearising this equation around k^* and using firms' capital demand (4), we obtain the following (date-t) savings demand curve:[7]

$$\hat{a}^d_t = -\left(\frac{k^*}{\bar{B} + k^*}\right)\left(\frac{\rho}{\mu\kappa^L}\right)\hat{R}^k_t + \left(\frac{k^*}{\bar{B} + k^*}\right)\hat{\mathcal{E}}_{At} \tag{17}$$

where hats denote proportional deviations from the steady state. The savings demand curve (depicted in Figure 2a) is decreasing in $\hat{R}^k_t$ and shifted by $\hat{\mathcal{E}}_{At}$. The slope of the curve is the elasticity of aggregate bond

[7]The curve is derived as follows. First, from equation (4) we get:

$$\hat{R}^k_t = -\frac{1}{\rho}\left(\frac{1-\alpha}{\alpha {k^*}^{\frac{\rho-1}{\rho}} + 1 - \alpha}\right)(\hat{k}_t - \hat{\mathcal{E}}_{At}) = -\frac{1}{\rho}\left(1 - \frac{\alpha}{\alpha + (1-\alpha){k^*}^{\frac{1-\rho}{\rho}}}\right)(\hat{k}_t - \hat{\mathcal{E}}_{At}),$$

where the term $\alpha/(\alpha + (1-\alpha){k^*}^{\frac{1-\rho}{\rho}})$ is equal to $\mu\kappa^C$, i.e. the product of the capital share and the aggregate markup. It follows that the elasticity of the demand for capital with respect to the interest rate is $-\rho/\mu\kappa^L$. Next, the log-linearization of $a^d_t = k_t + \bar{B}$ gives $(k^* + \bar{B})\hat{a}^d_t = k^*\hat{k}_t$. Combining those expressions yields (17) demand with respect to the interest rate. Holding everything else constant (including the labour share κ^L), greater firm market power μ reduces the sensitivity of bond demand to the interest rate – i.e. it flattens the savings demand curve. If the labour share also changes, then it is the product of the markup and the labour share that matters for the slope of the savings demand curve. In Section 3 below, we argue that this product has fallen between 1980 and 2020, pushing the slope of the savings demand curve downwards.

The second factor affecting the slope of the savings demand curve is the steady state shares of public debt versus capital over total assets. As the share of public debt increases (and consequently $k^*/(\bar{B} + k^*)$ falls), the sensitivity of total asset demand to interest rates falls, essentially because government issuance is less sensitive to those rates than firms'. It follows that a greater value of this ratio also contributes to flattening the savings demand curve. To the extent that both the share of public debt in total assets and the product of the aggregate markup to the labour share are greater in 2020 than in 1980, the savings demand curve has flattened over the past decades, as is depicted in Figure 2a.[8]

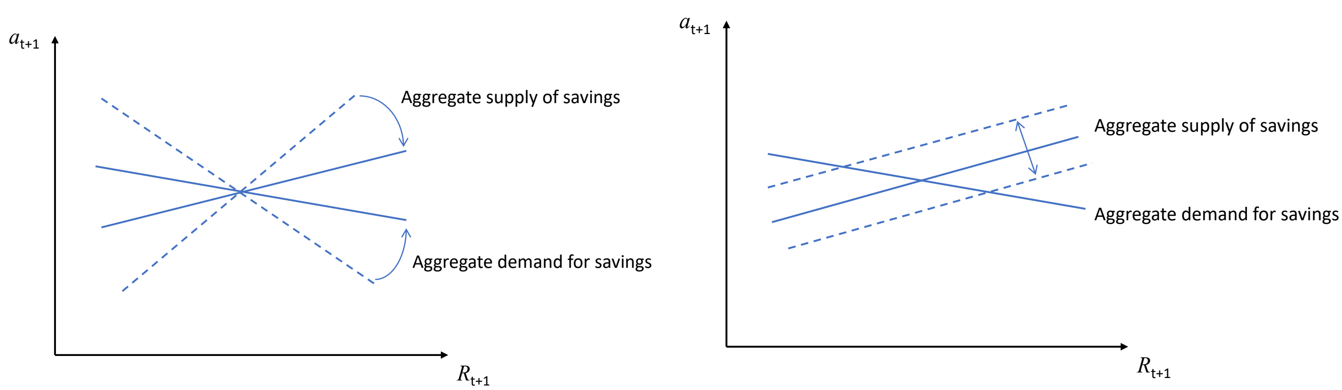

(a) **Savings supply and demand** (b) **Impact of a β_t-shock**

Figure 2: **Bond market equilibrium**

2.4.3 Savings supply curve

We similarly derive a linearised savings supply curve, based on the nonlinear curve derived in Section **??**. For the sake of expositional clarity, we confine our attention to the log utility benchmark, i.e., the case where $\sigma = 1$. This illustrates that nothing in the mechanisms that we highlight in this paper crucially rely on a low elasticity intertemporal substitution $(1/\sigma)$ – in fact, our purpose is precisely to stress that high old-age health spending relative to income strengthens intertemporal income effects regardless of the exact value of the intertemporal elasticity of substitution. Moreover, to the extent that the period length is long in an OLG model with two-period-lived households, it does make sense to calibrate the elasticity of intertemporal substitution to a relatively high value. In Section 3, when the period length is one year, we adopt a realistically lower calibration for the elasticity of intertemporal substitution.

[8]The last determinant of the slope of the savings demand curve is the capital-labour elasticity ρ. However, and unlike the other factors, there is no evidence that it has substantially evolved over time (see Section 3 for how we calibrate this parameter).

In the case where $\sigma = 1$, the log-linearisation of equation (12) gives:

$$\hat{a}^s_{t+1} = \underbrace{\left[1 + \frac{\Lambda}{1+\beta} - \frac{\psi \bar{B}}{\tilde{a}^s}\right] \widehat{\Omega}_t}_{\text{determined at } t-1} + \underbrace{\left[\frac{1+\Lambda - \psi\bar{B}/\tilde{a}}{1+\beta}\right] \hat{\beta}_t}_{\text{shocked at } t} + \underbrace{\left[\frac{\Lambda}{1+\beta}\right] \widehat{R}^k_{t+1}}_{\text{clearing at } t+1} \tag{18}$$

where $\hat{a}^s_{t+1} \equiv (\tilde{a}_t - \tilde{a})/\tilde{a}$ and $\Lambda = \delta/(\tilde{a}R)$. Note that in the $\left(R^k, \tilde{a}\right)$-plane, the slope of the savings supply curve is

$$\frac{\Lambda}{1+\beta} = \frac{\delta}{(1+\beta)(k^* + \bar{B})R^k}$$

where R^k is the steady state return to capital, i.e. the marginal product of capital divided by the aggregate markup (μ). All else equal, the savings supply curve flattens when δ falls and/or when either term in the denominator rises. As we argue in Section 3, δ has fallen over the past decades due to the rise in medical expenditures of old-age individuals. While μ has increased substantially over the past decades (see, e.g., De Loecker, Eeckhout and Unger, 2020), this has not translated into a lower return to capital, which may even have increased slightly (Gomme, Ravikumar and Rupert, 2011; Farhi and Gourio, 2018; Marx, Mojon and Velde, 2021). Finally, public debt has almost quadrupled, while the capital stock has not fallen. Therefore, if anything, the denominator has increased, reinforcing rather than overturning the effect of health spending. Through the lens of the model, this implies that the savings supply curve has flattened, as is depicted in Figure 2a.

2.4.4 Elasticities with respect to structural shocks

The slopes and shifts of the savings supply and demand curves ultimately determine the response of the interest rate to aggregate shocks. Note that the timing of shocks matters here, as is reflected in the fact that equations (17) and (18) do not have the same time index. To be more specific, a *preference shock* (to β_t) occurring at time t affects the supply of savings at that time and hence the capital stock of time $t+1$ (equation (18)). How the inflow of savings will materialize into more capital versus a lower interest rate depends on the slope of the savings demand curve (17), evaluated at time $t+1$ (see Figure 2a). The fact that both curves have flattened implies, all else equal, that a given shift in the savings supply curve produces a greater response of R_{t+1} to the underlying shock. This is depicted in Figure 2b for the case of discount-factor shocks.

To compute the elasticity of R_{t+1} with respect to β_t, equate $\hat{a}^d_{t+1}$ in (17) with $\hat{a}_{t+1}$ in (18) and solve for R_{t+1}. We get:

$$\epsilon_\beta = -\frac{\partial \log R_{t+1}}{\partial \log \beta_t} = \frac{\Lambda + 1 - \frac{\psi \bar{B}}{\tilde{a}}}{\Lambda + \left(\frac{k}{k^* + \bar{B}}\right)\left(\frac{\rho}{\mu \kappa^L (1-s)}\right)}, \tag{19}$$

where again $\Lambda = \delta/\tilde{a}R^k$. To understand how this elasticity is affected by the parameters of interest, focus first on the case where $\psi = 0$, i.e., there is no convenience yield on Treasury debt. In this case, equation (19) becomes:

$$\epsilon_\beta = \left[1 - \frac{1}{1+\Lambda}\left(1 - \frac{\rho k^*}{\mu \kappa^L (1-s)}\left(\frac{k^*}{k^* + \bar{B}}\right)\right)\right]^{-1}$$

Because the elasticity of substitution between labour and capital ρ is small (as low as 0.3 according to Gechert et al. 2022), the term in brackets in the denominator of (19)) is positive under any plausible parameterization of the model. If we further assume that this term is not too large so that ϵ_β has the standard sign – meaning positive, so that an increase in desired savings causes a fall in R_{t+1} –, then we can confirm that the elasticity ϵ_β is larger (all else equal) when (i) health spending relative to social security income is large – so that δ (and hence Λ) is small; (ii) the aggregate markup (i.e. μ) is large; and (iii) the share of public debt in total assets is large (so that $k^*/(k^* + \bar{B})$) is small. Those properties remain valid provided that ψ is not too large. Importantly, they will remain valid in the quantitative model of Section 3, when we compute the impulse response to a persistent β-shock.

Let us now turn to the elasticity with respect to productivity shocks. Suppose a productivity shock unexpectedly hits the economy at time $t+1$. At that time, the aggregate supply of savings has already been formed (based on the expected value of R^k_{t+1} as of time t) and is inelastic. This means that, in terms of deviations from steady state, we must have $\hat{a}^d_{t+1} = \hat{a}^s_{t+1} = 0$. It then follows from equation (17) that the elasticity of the interest rate with respect to the shock is:

$$\epsilon_{\mathcal{E}} = \frac{\partial \log R_{t+1}}{\partial \log \mathcal{E}_{At+1}} = \frac{\mu \kappa^L}{\rho} \tag{20}$$

and is thus greater when the product $\mu\kappa^L$ is larger, or when the elasticity ρ is smaller. As discussed above, while there is no evidence that ρ has changed over time, the evidence on the joint evolution of aggregate markups and the labour share suggests that $\mu\kappa^L$ has become larger over time – thereby increasing the elasticity $\epsilon_{\mathcal{E}}$.

3 Quantitative model

This section presents a quantitative overlapping-generations model, which nests the two-period setup of Section 2 as a special case. As before, the economy incorporates households of different ages, a vertically integrated two-layer production sector, and the government. To account for the effect of income and asset profiles, we now extend the life cycle of households to include multiple periods and age-specific productivities.

3.1 Setup

Households. A large number of households with identical preferences populate the economy. Each household has a finite planning horizon limited by the lifespan J. Every period, the oldest cohort of age J exits the economy, and a new cohort enters. Thus, households of different ages coexist in the economy at any point in time. Households become economically active and enter the economy upon reaching age 26.

Households born at time t choose their consumption path and asset portfolio to maximize expected lifetime utility:

$$\max_{\{c_t\}} E_t \sum_{j=26}^{J} \beta_t^{j-26} \frac{1}{1-\sigma} \Big(c_{j,t+j-26} + \psi b_{j+1,t+j-26} \Big)^{1-\sigma}, \tag{21}$$

where $c_{j,t}$ and $b_{j,t}$ represent consumption and bond holdings, respectively, of households of age j at time t. As before, β_t denotes the time discount factor at t, while ψ represents preferences for government bonds. Additionally, $\sigma > 0$ is the coefficient of relative risk aversion. Households choose consumption $c_{j,t}$ and asset portfolio subject to the budget constraint:

$$c_{j,t} + b_{j+1,t} + k_{j+1,t} = (1-\tau_t)\chi_j w_t + (1+r_t^b)b_{j,t-1} + (1+r_t^k)k_{j,t-1} + d_{j,t} + \delta_t I_{j \geq J_r} \tag{22}$$

In equation (22), $b_{j,t-1}$ and $k_{j,t-1}$ denote bond and capital holdings of j-year old households available at time t, $(1+r_t^b)b_{j,t-1}$ and $(1+r_t^k)k_{j,t-1}$ are corresponding gross payoffs (with $r_t^k \equiv R_t^k - 1$ and $r_t^b \equiv R_t^b - 1$), $d_{j,t}$ is age-specific profit payment, w_t is the wage rate (subject to the tax rate τ_t), and χ_j stands for age-specific productivity. Before age J_r, households remain in the labour force and supply one unit of labour inelastically. When households reach the retirement age J_r, they drop out of the labour force and $\forall j \geq J_r, \chi_j = 0$, which implies zero labour earnings. Similar to our baseline model, retired households receive retirement benefits ξ_t and bear health costs h_t, and $\delta_t = \xi_t - h_t$ denotes social security payments minus inelastic health spending. The age indicator $I_{j \geq J_r}$ equals unity if the condition in the subscript is true and zero otherwise.

Firms. The production side of this economy is the same as in Section 2, except that we no longer assume full depreciation of capital. Consequently, we rewrite the optimal demand for capital as:

$$r_t^k + \delta_k = \frac{\alpha}{\mu}\left(\frac{Y_t}{K_t}\right)^{\frac{1}{\rho}} \tag{23}$$

where δ_k is the depreciation rate. Aggregate pure profits D_t are output minus factor payments:

$$D_t = Y_t - (r_t^k + \delta_k)K_t - w_t L_t \tag{24}$$

Finally, following Eggertsson, Mehrotra and Robbins (2019), we assume that profits are distributed proportionally to labour income:

$$d_{j,t} = \frac{\chi_j}{\sum_{j=26}^{J_r-1}\chi_j} D_t \tag{25}$$

Government, aggregation and exogenous processes. As in the baseline two-period model, the government provides retirement social security. The government imposes a proportional tax on labour earnings τ_t and issues one-period bonds B_t to finance its expenditures. The government budget constraint is:

$$B_t + \tau_t w_t L_t = \Xi_t + (1+r_t^b)B_{t-1} \tag{26}$$

where Ξ_t is the total government spending on retirement benefits, calculated as $\Xi_t = \xi_t \cdot (J - J_r + 1)$ using the fact that all cohorts are of the same size (normalised to 1). Equal cohort size also implies that aggregate labor supply L_t is

$$L_t = L = \sum_{j=26}^{J_r-1} \chi_j \tag{27}$$

Furthermore, clearing of the asset market requires:

$$K_{t+1} = \sum_{j=26}^{J} k_{j+1,t} \tag{28}$$

$$B_t = \sum_{j=26}^{J} b_{j+1,t} \tag{29}$$

Finally, the process for the preference and log of productivity (β_t and $\ln A_t$) are:

$$\ln A_t = \rho^A \ln A_{t-1} + \varepsilon_t^A \tag{30}$$

$$\beta_t = (1-\rho^\beta)\bar{\beta} + \rho^\beta \beta_{t-1} + \varepsilon_t^\beta \tag{31}$$

where $\bar{\beta}$ is the steady-state value of the subjective discount factor and ρ^A and ρ^β are the persistence coefficients for the productivity and time-preference shocks. ε_t^A and ε_t^β are innovations, where $\varepsilon_t^A \sim \mathcal{N}(0, \sigma_A^2)$, $\varepsilon_t^\beta \sim \mathcal{N}(0, \sigma_\beta^2)$.

3.2 Calibration

To quantitatively assess the responsiveness of the equilibrium interest rate to aggregate shocks affecting the aggregate savings demand and supply, we calibrate our model to match key moments of aggregate US data. We consider two different sets of parameters, corresponding to the two steady states under consideration: the US economy in "1980" and "2020". We first discuss the parameters that are identical across steady states and then turn to the parameters that differ across steady states: parameters mapping to markups, factor shares, government debt, and health expenditures.

One period in the model corresponds to a year. The duration of the economically active life of households is 60 years ($J = 85$). Households work in their first 40 years of life ($J_r = 66$) and then stay in retirement for the next 20 years. We choose the age-specific productivity components χ_j to match the evidence on the life-cycle profile of labour earnings from PSID data. We take the cubic regression estimates of labour income on age from Heathcote, Storesletten and Violante (2010). The retirement benefit ξ matches the 40% replacement ratio. We set the value of the relative risk aversion σ to 4, as is standard in quantitative OLG models (Auerbach and Kotlikoff, 1987; Ríos-Rull, 1996). Finally, following the meta-analysis of Gechert et al. (2022), we calibrate the elasticity of substitution between capital and labour (the parameter ρ in the production function) to 0.3.

Next, we turn to the set of parameters whose value changes across the "1980" and "2020" steady states (see Table 1 for a summary): the elasticity of substitution between intermediate goods (θ) and the coefficient before the capital stock (α) in the production function for final goods (equation (2)), the depreciation rate δ_k, steady-state public debt ($\bar{B}$) and health expenditure (h). Those are calibrated to match the following five moments: aggregate firm markups, the labour share, the investment-output ratio, the public debt-output ratio, and the ratio of health expenditure to the retirement benefit. First, we follow Eggertsson, Mehrotra and Robbins (2019) in setting the elasticity of substitution θ to 5 in our most recent steady state (2020 here, 2015 in their paper). This corresponds to a 25% markup rate over marginal costs. Consistent with the evidence that aggregate markups have considerably increased over the past decades (e.g., Farhi and Gourio, 2018; De Loecker, Eeckhout and Unger, 2020), we set θ so that the aggregate markup in 1980 is only 10%. We calibrate the production function parameter α to match the labour share, which according to Autor et al. (2020) has dropped from 67% in 1980 to 59% in 2020. As discussed before, Farhi and Gourio (2018) and Marx, Mojon and Velde (2021) pointed out (using the basis of the updated series of Gomme, Ravikumar and Rupert (2011)), that, unlike the riskless rate, the return to capital has not declined over the past 40 years (and may even have increased slightly). Accordingly, we calibrate β in either steady state to target an annual return to capital ($r^k = MPK/\mu - \delta_k$ in the model) of 5%, roughly in line with the average post-tax return on all capital over the period in Gomme, Ravikumar and Rupert (2011). We choose the value of ψ that matches the real return on government bonds, given the return on capital. For "1980" we target the real average market yield on US Treasury Securities at a 10-year constant maturity, i.e. the nominal yield minus average inflation over 1980-1984; this gives us a target of 4.5% annually in real term, slightly below the return on capital. For "2020", the same computation gives us an annual real return on Treasuries of 1% (our results are almost insensitive to plausible variations around these numbers.) We calibrate public debt to match the ratio of gross federal debt to GDP (from the FRED database), which was 30% in 1980 and hovered around 120% after the Covid-19 crisis.[9]

As discussed in Section 2 in the context of the two-life-period model, the potency of intertemporal income effects shapes the savings supply curve. Importantly, greater constrained health spending in old age magnifies intertemporal income effects, dampening households' response to changes in asset returns and flattening the savings supply curve. In practice, most old-age OOP health expenditures are non-discretionary (see, e.g., McInerney et al., 2022 for a discussion). We adopt a model-consistent (and conservative) calibration by considering only that part of old-age health expenditures that are strictly inelastic – because they are regulated– namely, Medicare Part B premia.[10] In 1980, the annual Medicare Part B premium was \$115.2, while the average annual Social Security benefit was \$3337.7, implying that retirees were paying about 3.5% of their benefits in Medicare Part B premia on average. In 2022 the Part B premium had risen to \$2040 annually, and the average Social Security benefit was \$17526.5, implying an average share of healthcare spending of 12%.[11] We calibrate the steady state values of δ_t in equation (21) (i.e., social security minus constrained health spending) to match those figures. Note that considering other potential sources of increases in OOP health expenditure would quantitatively magnify the effects that we are stressing.[12]

Finally, when drawing impulse-response functions, we calibrate the persistence of the productivity processes to $\rho^A = 0.99$ (close to the unit-root process postulated in Section 2, and also to the value typically estimated in the data) and the persistence of the preference shock to the much lower value $\rho^\beta = 0.90$. This is consistent with the (New Keynesian) view that the latter shocks drive fluctuations in aggregate demand and, as such, are significantly less persistent than productivity shocks.

[9]The ratio actually peaked at a value near 130% in 2020, but the peak was short-lived.

[10]Unlike Medicare Part A, participation in Medicare Part B requires the payment of a monthly premium. Whilst it is possible to opt-out, in practice, retirees do not once they have signed in, making the premium payments non-discretionary.

[11]Medicare Part B premia up to 2016 are in Table 2.C1 of Social Security Administration (2016), and those for more recent years at the Centers for Medicare & Medicaid Services (www.cms.gov). Average annual social security benefits are obtained by dividing total social security payments (FRED series W823RC1) by the number of enrollees obtained from the Social Security Administration (www.ssa.gov/oact/STATS/OASDIbenies.html).

[12]For example, Platzer and Peruffo (2022) argue that the price of discretionary OOP spending has increased over the last decades and calibrate their model accordingly.

Parameter	Meaning	Targeted Moment	"1980"	"2020"
θ	Elasticity of substitution	Average markup rate	10%	25%
α	Prod. function parameter	Labor share	67%	59%
B	Public debt	Debt-GDP ratio	30%	120%
β	Subjective discount factor	Return to capital	5%	5%
ψ	Preference for bonds	Return to bonds	4.5%	1%
δ_k	Depreciation rate	Investment-to-output ratio	15.9%	15.9%
h	Health spending	Medicare Part B premium over social security benefit	3.5%	12%

Table 1: **Parameters and data targets**

3.3 Results

We begin by plotting the savings demand and supply curves generated by our quantitative model, i.e., the quantitative analogues of equations (17) and (18). Because firms are homogeneous, the savings *demand* curve can be computed analytically, just like Section 2. We numerically construct the savings *supply* curve by aggregating the individual saving decisions of households of different ages in the range of plausible interest rate values. The two curves and their intersection are depicted in Figure 3, for the "1980" and "2020" economies. Just as predicted in the context of the two-life-period model of Section 2, the increase

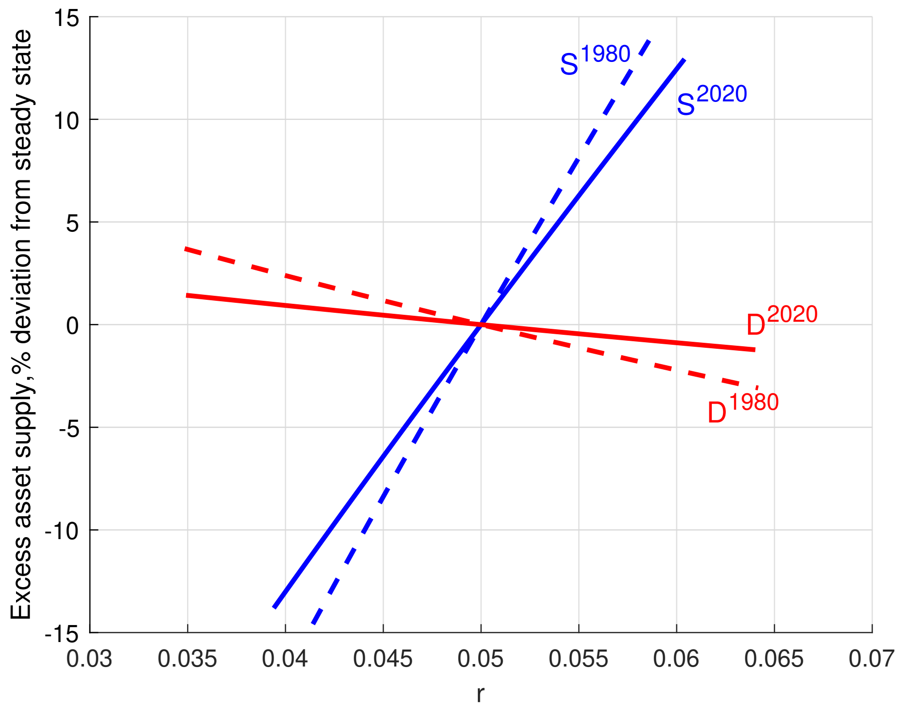

Figure 3: **Quantitative savings demand and supply curves**

in OOP health spending flattens the savings supply curve in the quantitative model (i.e., it decreases its slope). On the other hand, the rise in firms' market power and the share of public debt in total assets flatten the savings demand curve (i.e., they increase its slope).

Next, we measure the dynamic impact of a given structural shock on the equilibrium interest rate, comparing impulse responses across the 1980 and 2020 calibrations. To produce a meaningful comparison, we subject the economy to aggregate shocks that are normalised to generate a trough of 100 basis points in the equilibrium interest rate under the 1980 calibration; the size of the trough under the 2020 calibration (for the very same shock) then tells us what amplification is brought about by the change in parameters – this is the quantitative analogue of comparing elasticities in the context of the simple model of Section 2.

First, consider a positive shock to the subjective discount factor β that hits the economy in period $t = 0$. The left panel of Figure 4 shows the interest rate adjustment following the shock. Note that the aggregate supply of savings at time t is predetermined (it follows from savings decisions made in period $t = -1$), while the demand for savings is not affected by the shock other than via the interest rate.[13] It follows that the interest rate stays still at time $t = 0$. However, as the shock hits at time $t = 0$, households seek to shift their consumption to future periods, and end-of-period savings go up, creating downward pressure on the interest rate that gradually builds up from time $t+1$ onward. Consistent with our analysis of the slopes of the savings supply and demand curve, we observe a significantly more pronounced (and persistent) decline in the equilibrium interest rate for the 2020 economy. In particular, the interest-rate trough is magnified by 48bp.

Next, consider the dynamic impact of a shock to productivity (right panel of Figure 4). Again, the shock is normalised to produce a 100bp drop under the 1980 calibration. While the aggregate supply of savings is predetermined when the shock hits, the demand for savings collapses on impact, creating immediate downward pressure on the equilibrium interest rate. The impact is maximum at the time of the shock, after which the interest rate rises, overshoots, and eventually converges back to a steady state. In the case of a productivity shock, the interest rate trough is magnified by 40bp in 2020, relative to 1980.

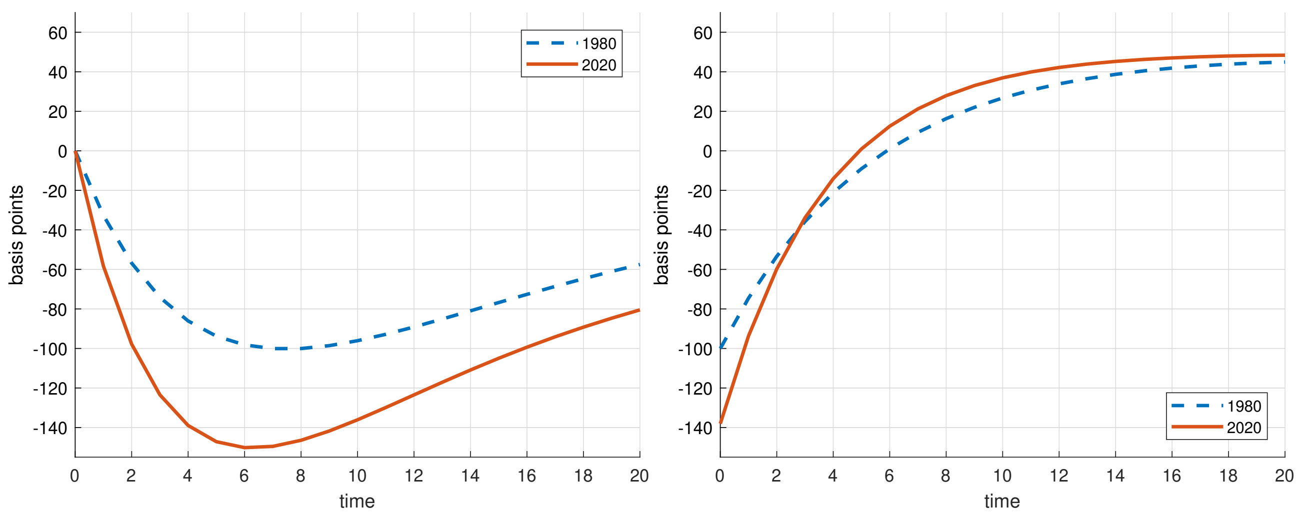

(a) Response to a β_t-shock

(b) Response to a A_t-shock

Figure 4: **Interest rate responses to aggregate shocks**

Finally, in order to quantify the contribution of the different forces at play to the disparity in the interest-rate responses between the 1980 and 2020 economies, we consider three counterfactual sets of calibrated parameters, isolating, in turn, the impact of health spending, goods market competition, and government debt. In these counterfactual calibrations, we repeat our earlier quantitative experiment of

[13]In the context of the two-life-period model, equation (18) makes it clear that the time t β-shock affects $\hat{a}^s_{t+1}$, not $\hat{a}^s_t$.

perturbing the economies with productivity and time-preference shocks. Then, we gauge the contribution of each channel to the amplification of the interest-rate response by comparing the disparities with the baseline calibration.

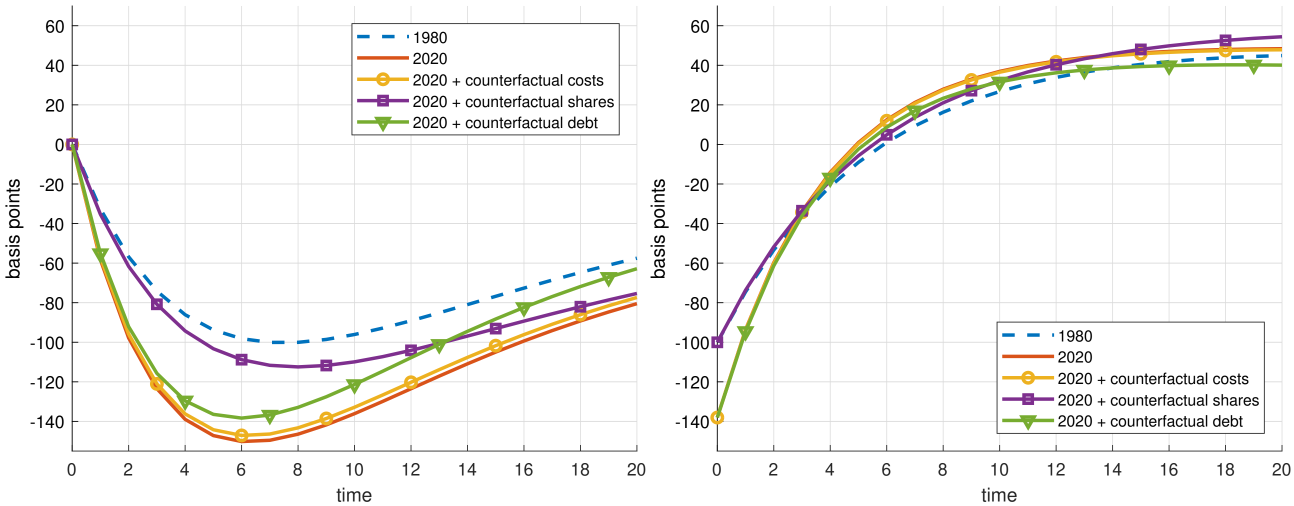

(a) Impact of a β_t-shock

(b) Response to a A_t-shock

Figure 5: **Counterfactual interest-rate paths**

The left panel in Figure 5 shows the decomposition of the interest rate response following the time-preference shock. First, consider an economy with the same markup, labour share and government debt as in the 2020 calibration, but with counterfactual health costs of 3.5% of retirement benefits, i.e., corresponding to its value in 1980. We observe that the interest rate response to the time-preference shock in this economy is less pronounced than under the baseline 2020 calibration. Quantitatively, health costs account for 5-10% of the difference between the responses in 1980 and 2020. Next, we isolate the effect of goods market competition by fixing the health cost and government debt to match the 2020 moments, but setting the aggregate markup and the labour share to their calibrated value of 1980. The interest-rate response in this counterfactual economy shows that goods market competition accounts for most of the quantitative difference across our baseline scenarios of 1980 and 2020. Finally, to evaluate the role of government debt, we examine the interest rate response in the 2020 economy with government debt set to its counterfactual 1980 value. This shows that the increase in government debt is quantitatively significant and accounts for up to 20% of overall amplification.

Next, we run the same counterfactual experiments to decompose the interest-rate response to the productivity shock (right panel of Figure 5). The impact response of the interest rate is more substantial in 2020, and virtually all quantitative differences can be fully attributed to decreased goods market competition. The contribution of the other forces, namely health costs and government debt, is negligible. This is consistent with our theoretical results established in Section 2 – see the formula for the $\epsilon_{\mathcal{E}}$ elasticity. To summarize, the joint effect of increased old-age health spending and public debt and decreased goods market competition significantly amplifies the response of the equilibrium interest to aggregate shocks, as predicted by the simple model of Section 2. The amplification is predominantly driven by the evolution of pubic debt and markups, with a more minor role for changes in old-age health spending.

4 Robustness

Our baseline quantitative model is, in many ways, minimal and accordingly abstracts from several structural changes potentially affecting the responsiveness of the natural interest rate to aggregate shocks. In this section, we incorporate two sets of structural changes that have occurred between "1980" and "2020", namely (i) changes in population structure (population growth, life expectancy, retirement age, and retirement replacement ratio) and (ii) changes in private debt and inequality.

4.1 Population structure and consumer debt

First, we alter population ageing relative to the baseline model as follows. Just as before, households enter the economy at the age of 26 and are certain to pass away by the age of 85. However, we now introduce the possibility that agents can die earlier, facing an idiosyncratic risk of early death. The probability of surviving until age $j+1$, conditional on being alive in age j, is denoted by subscript s_j. The unconditional probability of surviving until age $j+1$ is denoted by superscript s^j. Since nobody survives after age J, the corresponding survival probabilities beyond $J+1$ are $s_J = 0$ and $s^J = 0$. We also allow for a change in the retirement age, from the age of 65 in 1980 to the age of 67 in 2020.

Households now incorporate the possibility (and associated probability) of an early death into their optimization problem. The expected lifetime utility of a household is now:

$$\max_{\{c_t\}} E_t \sum_{j=26}^{J} s^j \beta_t^{j-26} \frac{1}{1-\sigma} \Big(c_{j,t+j-26} + \psi b_{j+1,t+j-26} \Big)^{1-\sigma} \tag{32}$$

Given that the population consists of a large number of households, there is no uncertainty regarding the size of cohorts at any time t. Following Ríos-Rull (1996), agents self-insure against idiosyncratic mortality risk via fair one-period annuity contracts. Consequently, the budget constraint of a household is now given by:

$$c_{j,t} + b_{j+1,t} + k_{j+1,t} = (1-\tau_t)\chi_j w_t + (1+r_t^b)\frac{1}{s_j} b_{j,t-1} + (1+r_t^k)\frac{1}{s_j} k_{j,t-1} + d_{j,t} + \delta_t I_{j \geq J_r} \tag{33}$$

Second, we assume that the economy features a constant exogenous rate of population growth, denoted by g_n. Having normalised the size of the youngest cohort to unity, we rewrite the aggregate labour supply and the asset market equilibrium as

$$L_t = L = \sum_{j=26}^{J_r-1} \frac{s^j}{(1+g_n)^{j-26}} \chi_j \qquad K_{t+1} = \sum_{j=26}^{J} \frac{s^j}{(1+g_n)^{j-26}} k_{j+1,t} \qquad B_t = \sum_{j=26}^{J} \frac{s^j}{(1+g_n)^{j-26}} b_{j+1,t} \tag{34}$$

Finally, to match realistic values of consumer debt, we assume that young households (those before 35) have a lower time discounting factor $\beta^y < \beta$, such that $(1+r_t)\beta^y < 1$. Additionally, young households face an (ad-hoc) borrowing limit d^b:

$$a_{j+1,t} \geq d^b \tag{35}$$

In this setting, during the most productive years of the life cycle, households save not only for retirement but also to repay debts accumulated during the earlier stages of their life, as in Eggertsson, Mehrotra and Robbins (2019), and Jimeno (2019).

4.2 Parametrization

As before, we examine two distinct parameter sets associated with the two steady states: the US economy in 1980 and 2020. Now, in addition to the seven data targets reported in Table 1, we calibrate the model with five more targets: the growth rate of the population, the survival probabilities at different ages, the consumer debt to output ratio, the retirement age, and the retirement replacement ratio (see Table 2 for details.)

Parameter	Meaning	Targeted Moment	*1980*	*2020*
θ	Elasticity of substitution	Average markup rate	10%	25%
α	Production function parameter	Labor share	67%	59%
B	Public debt	Debt-GDP ratio	30%	120%
β	Subjective discount factor	Return to capital	5%	5%
ψ	Preference for bonds	Return to bonds	4.5%	1%
δ_k	Depreciation rate	Investment-to-output ratio	15.9%	15.9%
h	Health spending	Medicare Part B premium over social security benefit	3.5%	12%
d^b	Borrowing constraint	Consumer-debt-GDP ratio	4.2%	6.3%
ξ	Unemployment benefit	Replacement rate	50%	40%
J_r	Retirement age	Full retirement age	65	67
g_n	Population growth	Annual population growth	0.960%	0.455%
s^j	Survival probabilities	Actual survival probabilities		

Table 2: **Additional parameters and data targets**

The population growth rates we target are in line with the estimates of the US Census Bureau, which were 0.96% in 1980 and 0.455% in 2019 (the last pre-pandemic year). Our survival probabilities s^j are the actual survival rates in 1980 and 2019, sourced from the Human Mortality Database. We set the unconditional probability of surviving until age 26 to unity ($s^{26} = 1$) and set all the remaining survival probabilities according to the mortality tables.

We calibrate the ad hoc borrowing limit d^b to match the consumer debt to output ratio. Following Eggertsson, Mehrotra and Robbins (2019), we consider an increase in consumer debt from 4.2% in 1980 to 6.3% in 2020.

Finally, we allow for changes in two parameters related to retirement: the pension replacement ratio and the retirement age. According to the US Social Security Administration, the retirement social security benefit in 2020 replaced about 40% of the labour income; the size of the pension ξ matches this value in the model calibrated for 2020. We calibrate ξ to match 50% of labour income in 1980, corresponding to the replacement ratio of medium earners in the early 1980s (Diamond and Gruber, 1999). Regarding the retirement age, we consider an increase in the age of full retirement from 65 in 1980 to 67 in 2020.

We summarize all twelve parameters that vary across models in Table 2. The parameters related to relative risk aversion, age-specific productivity, and the elasticity of substitution between capital and labour are the same as in Section 3.

4.3 Interest-rate responses

We reexamine the interest-rate sensitivity to aggregate time-preference and productivity shocks in the 1980 and 2020 economies, allowing for the additional features just described. Hereafter, we refer to the economies that match all corresponding targets in Table 2 as "1980" and "2020". Additionally, we renormalize the size of shocks to generate a trough of 100 basis points following a β_t-shock and 100 basis points on-impact interest rate drop after a negative A_t-shock under the 1980 calibration. Figure 6 presents the interest-rate response following the shocks.

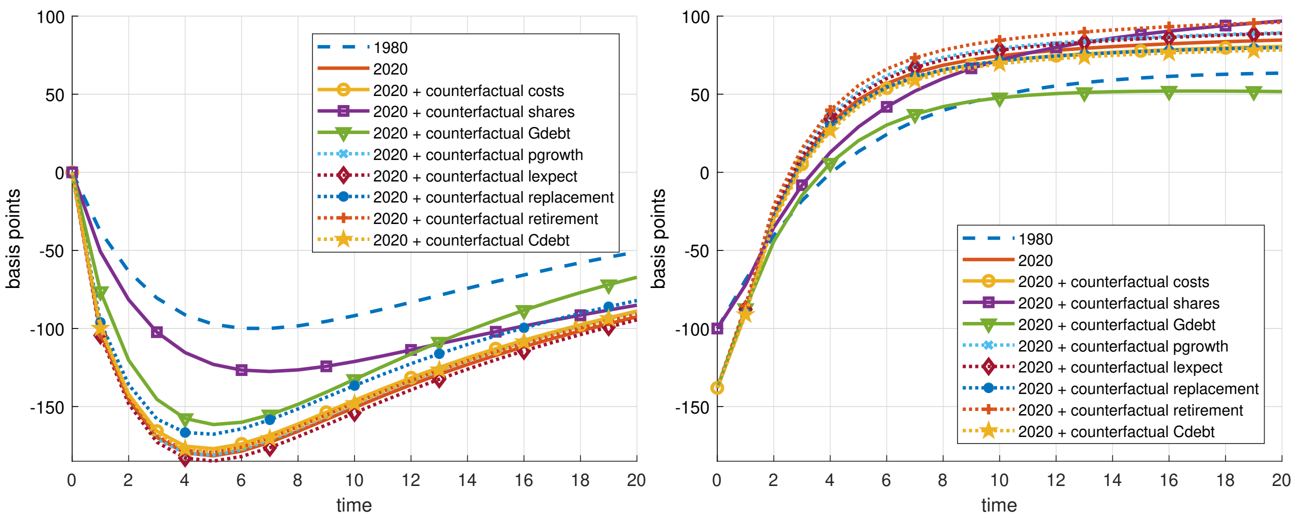

(a) Impact of a β_t-shock (b) Impact of a A_t-shock

Figure 6: **Interest-rate paths in the extended model.**

The left panel of Figure 6 illustrates that accounting for additional data moments does not qualitatively alter the interest response to a preference shock. Comparing the responses in 2020 to those in 1980, the economy experienced a much more pronounced decline in the equilibrium interest rate following a shock of the same size. Quantitatively, the interest rate response is magnified by 85 basis points. Moreover, when considering a negative productivity shock, the economy responds with a 140 basis points decline in interest rates in 2020, contrasting with the 100 basis points drop observed when matching the 1980 moments.

This confirms that our findings in Section 3 are robust to the relaxation of our baseline assumptions regarding population structure and borrowing constraints. Specifically, a decrease in goods market competition and an increase in government debt remain the most important drivers of the sensitivity of the natural interest rate to structural aggregate shocks. However, the increase in health spending contributes only minimally to the overall amplification. In the case of a productivity shock (right panel of Figure 6), the quantitative difference between the interest-rate responses is entirely attributed to the change in the degree of goods market competition.

Finally, we quantify the specific contribution of population growth, life expectancy, consumer debt, and social security to the disparity in the interest-rate responses between the 1980 and 2020 economies. We find that a decrease in the retirement replacement rate accounts for about 10% of the overall amplification, while other forces play a quantitatively small role. Additionally, we document that the aforementioned factors do not affect the savings demand curve and are hence irrelevant for the amplification of aggregate productivity shock.

5 Concluding remarks

In this paper, we have attempted to clarify the fundamental forces affecting the size of the response of the equilibrium interest rate to macroeconomic shocks affecting the demand for and supply of aggregate savings – focusing on the factors that have likely changed between the early 1980s and the recent years, namely (i) the increases in old-age out-of-pocket medical expenditures, (ii) firm markups, and (iii) public debt. The model incorporates the most minimal set of frictions necessary to capture those forces, namely: (i) imperfect competition and (ii) an overlapping-generations structure generating life-cycle savings and the non-neutrality of public debt. Our quantitative framework suggests that changes in markups and public debt (i.e., the demand side of the market for aggregate savings) significantly magnified the interest-rate response to shocks, leaving a more moderate role for old-age medical spending (i.e., the supply side of this market).

References

Anzoategui, Diego, Diego Comin, Mark Gertler, and Joseba Martinez. 2019. "Endogenous technology adoption and R&D as sources of business cycle persistence." *American Economic Journal: Macroeconomics*, 11(3): 67–110.

Auerbach, Alan J, and Laurence J Kotlikoff. 1987. *Dynamic Fiscal Policy.* Cambridge University Press.

Autor, David, David Dorn, Lawrence F Katz, Christina Patterson, and John Van Reenen. 2020. "The fall of the labor share and the rise of superstar firms." *The Quarterly Journal of Economics*, 135(2): 645–709.

Bailey, Andrew John, Ambrogio Cesa-Bianchi, Marco Garofalo, Richard Harrison, Nick McLaren, Sophie Piton, and Rana Sajedi. 2022. "Structural change, global R* and the missing-investment puzzle." *Bank of England Staff Working Paper No. 997.*

Barro, Robert J. 1974. "Are government bonds net wealth?" *Journal of political economy*, 82(6): 1095–1117.

Campos, Rodolfo G, Jesús Fernández-Villaverde, Galo Nuño, and Peter Paz. 2024. "Navigating by falling stars: monetary policy with fiscally driven natural rates." National Bureau of Economic Research.

Chiang, Yu-Ting, and Piotr Zoch. 2022. "Financial Intermediation and Aggregate Demand: A Sufficient Statistics Approach." Federal Reserve Bank of St. Louis.

Del Negro, Marco, Domenico Giannone, Marc P Giannoni, and Andrea Tambalotti. 2017. "Safety, liquidity, and the natural rate of interest." *Brookings Papers on Economic Activity*, 2017(1): 235–316.

Del Negro, Marco, Domenico Giannone, Marc P Giannoni, and Andrea Tambalotti. 2019. "Global trends in interest rates." *Journal of International Economics*, 118: 248–262.

De Loecker, Jan, Jan Eeckhout, and Gabriel Unger. 2020. "The rise of market power and the macroeconomic implications." *The Quarterly Journal of Economics*, 135(2): 561–644.

Diamond, Peter A. 1965. "National debt in a neoclassical growth model." *The American Economic Review*, 55(5): 1126–1150.

Diamond, Peter, and Jonathan Gruber. 1999. "Social security and retirement in the United States." In *Social security and retirement around the world*. 437–473. University of Chicago Press.

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.

Farhi, Emmanuel, and François Gourio. 2018. "Accounting for macro-finance trends: Market power, intangibles, and risk premia." National Bureau of Economic Research.

Fisher, Jonas DM. 2015. "On the structural interpretation of the smets–wouters "risk premium" shock." *Journal of Money, Credit and Banking*, 47(2-3): 511–516.

Gechert, Sebastian, Tomas Havranek, Zuzana Irsova, and Dominika Kolcunova. 2022. "Measuring capital-labor substitution: The importance of method choices and publication bias." *Review of Economic Dynamics*, 45: 55–82.

Gomme, Paul, B Ravikumar, and Peter Rupert. 2011. "The return to capital and the business cycle." *Review of Economic Dynamics*, 14(2): 262–278.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L Violante. 2010. "The macroeconomic implications of rising wage inequality in the United States." *Journal of political economy*, 118(4): 681–722.

Jimeno, Juan F. 2019. "Fewer babies and more robots: economic growth in a new era of demographic and technological changes." *SERIEs*, 10(2): 93–114.

Krishnamurthy, Arvind, and Annette Vissing-Jorgensen. 2012. "The aggregate demand for treasury debt." *Journal of Political Economy*, 120(2): 233–267.

Marx, Magali, Benoît Mojon, and François R Velde. 2021. "Why have interest rates fallen far below the return on capital?" *Journal of Monetary Economics*, 124: S57–S76.

McInerney, Melissa, Matthew S Rutledge, Sara Ellen King, et al. 2022. "How Much Does Health Spending Eat Away at Retirement Income?" Center for Retirement Research.

Obstfeld, Maurice. 2023. "Natural and Neutral Real Interest Rates: Past and Future."

Platzer, Josef, and Marcel Peruffo. 2022. *Secular Drivers of the Natural Rate of Interest in the United States: A Quantitative Evaluation.* International Monetary Fund.

Rachel, Łukasz, and Lawrence H Summers. 2019. "On secular stagnation in the industrialized world." National Bureau of Economic Research.

Ríos-Rull, José-Víctor. 1996. "Life-cycle economies and aggregate fluctuations." *The Review of Economic Studies*, 63(3): 465–489.

Rogoff, Kenneth S, Barbara Rossi, and Paul Schmelzing. 2022. "Long-Run Trends in Long-Maturity Real Rates 1311-2021." National Bureau of Economic Research.

Social Security Administration. 2016. "Annual Statistical Supplement to the Social Security Bulletin, 2015." Office of Research, Evaluation, and Statistics.