Public Debt and Interest Rates*

William B. Peterman
Federal Reserve Board

Erick Sager
Federal Reserve Board

February 26, 2025

Abstract

Motivated by the historically large rise in U.S. public debt after 2007, this paper studies the interest rate response to higher public debt. Using a quantitative heterogeneous agent life-cycle model, we measure the interest rate responses to different one-time fiscal shocks that raise the level of debt and find that interest rate movements depend heavily on what policies the debt finances, in contrast to the conventional approach of abstracting from the composition of policies underlying the debt. For instance, a government consumption shock initially crowds out productive capital by roughly one-to-one, leading to a large increase in the interest rate at impact. Consequently, elevated interest rates lead households to accumulate more savings, which creates funds for productive capital and mitigates some of the initial interest rate increase. However, savings accumulating is slow, leading to large effects on longer term yields at the time of the shock. In contrast, a one-time increase in transfers or income tax reduction increases households savings significantly on impact, thus offsetting the crowd-out from the increase in debt and, depending on the specifics of the policy, a larger savings response can actually lower the interest rate at the time of the shock. The interest rate then rises to the new higher steady state level as households consumer these resources over time. Thus, policies that boost household savings generally have a smaller (or even negative) effect on longer term yields. Quantitatively, we find that the elasticity of a 10-year yield with respect to public debt varies significantly across policy types and maturities, with an elasticity between -0.1 to 2.0 basis points. Finally, estimating the elasticity with a reduced-form estimation strategy used in previous research, we find supportive evidence that the effects vary across the term and type of policies generating the debt.

Keywords: Government Debt; Life Cycle; Heterogeneous Agents; Incomplete Markets

JEL Codes: H6, E21, E6

*Correspondence: william.b.peterman@frb.gov or erick.r.sager@frb.gov. The authors thank Eric Engen and Kevin Seliski for insightful comments and discussions, as well as seminar participants at Midwest Macro at VCU, the National Tax Association Annual Meetings, and the Bank of Italy. We thank Octavio Aguilar for excellent research assistance. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Board of Governors or the Federal Reserve System. All errors are our own.

1 Introduction

After having stabilized in the 1990s, U.S. public debt has risen immensely since the onset of the Great Recession from 60% to its current size of 120% as a share of GDP. This large increase in public debt has generated concerns about the extent to which public debt distorts the allocation of resources to the private sector, chiefly due to the potential for crowding out the productive capital stock and inducing higher interest rates. Seminal research attempts to quantify these distortions with a singular elasticity, estimating that a one percentage point increase of public debt as a share of GDP leads to an increase in long term interest rates on US issued securities by 2 to 5 basis points.[1] However, while these prior estimates clearly imply that the rise in public debt would put upward pressure on interest rates, the risk free real interest rate has actually been low and generally trended downward since the 1990s.[2] Given the potential incompatibility between the steep increase in debt with the apparent secular decline in interest rates, this paper revisits the relationship between public debt and the real interest rate. However, instead of representing this relationship with a singular elasticity, we consider how the real interest rate response varies with different debt-financed fiscal policies. Moreover, we also consider the evolution of interest rates over a period of transition after the policy implementation so that we can investigate how the effect of fiscal policies additionally varies over the yield curve. Importantly, we find that this relationship is not well approximated with just one elasticity, but instead that there is substantial heterogeneity in interest rate responses both to different types of debt-financed fiscal policies and at different horizons of the yield curve.

Using a quantitative heterogeneous life-cycle model that is calibrated to match a rich set of empirical observations for the U.S. economy, we conduct a series of counterfactual experiments designed to study interest rate responses at varying maturities to increased public debt from different debt-financed fiscal policies. Specifically, we study transition paths after three different shocks that increase debt by the same amount and converge to a common steady state: (i) a one-period increase in government purchases, (ii) a one-period increase in government transfers to households,

[1] This research considers the five-year forward five-year interest rate. For example, see Gale and Orszag (2004), Engen and Hubbard (2004), Laubach (2009), as well as the more recent update in Gamber and Seliski (2019).

[2] We focus on the risk free rate since it is a rate that is unaffected by temporary moves due to fluctuations in short rates from monetary policy during this period. For estimates see Del Negro et al. (2017), Holston et al. (2017), Johannsen and Mertens (2018), Kiley (2020), Lubik and Matthes (2015), Gomme et al. (2011)). Moreover, long rates which would also be less affected by the business cycle were relatively low over this period. See Bauer and Rudebusch (2016) and Levrero and Deleidi (2019) for a discussion.

(iii) and a one-period reduction in income tax rates. All three fiscal shocks yield an interest rate elasticity of 1.4 basis points with respect to a 1 percentage point increase in public debt between the initial and new steady states that are common to each. Steady state comparisons underlie the conventional wisdom that a single elasticity with respect to public debt – no matter the composition of policies that generated that debt – is sufficient to characterize interest rate responses to fiscal policy. However, we find substantial variation in interest rate responses of varying maturities that arise from transitional dynamics following each of the fiscal shocks. For example, when we focus on the 5-year forward 10-year interest rate – which is typically used to empirically estimate the interest rate elasticity to debt – we find a range of elasticities between -0.1 to 2 basis points depending on the type of fiscal policy shock.

There are two channels by which a fiscal policy affects the interest rate.[3] First, holding household savings constant, public debt directly crowds out productive capital thus raising the interest rate (we refer to this as the *direct crowd-out channel*). Second, household savings decisions need not be held constant in general equilibrium, and the partial equilibrium increase in the interest rate from the direct crowd out channel leads households to save more thereby offsetting some of reduction in productive capital and inducing a smaller increase in interest rates in general equilibrium (we refer to this as the savings channel). While the direct crowd-out channel has the same immediate impact on spot interest rates for any fiscal policy, the savings channel's effect varies with respect to both policy type and along the transition. For example, when the debt finances government expenditures, the savings channel is essentially zero at the onset of the shock since households do not anticipate the fiscal shock. Thus, with no offset to the direct crowd out channel, the spot interest rate increases even more at the onset of the shock and overshoots the eventually higher steady state rate. Over time households accumulate more savings in response to a higher spot interest rate, which puts downward pressure on the spot interest rate. Thus, the longer term yield at the time of the shock is even higher than the new steady state interest rate leading to a larger elasticity on longer yields. In contrast, when the debt finances a lower income tax rate or lump sum transfers to households for one period, households receive resources. Households respond by saving most of the one-time transfer of resources, which strengthens the savings channel at the onset of the shock. Thus, the spot interest rate is lower at the onset of the transition due to this excess savings. As households spend down some of this influx leading the level of savings to normalize at a level that is higher than the initial pre-shock level but lower than the initial elevated level

[3]We consider both spot rates and longer run yields, which are determined by the path of the one period spot interest rates over each transition from a fiscal policy change.

at the onset of the debt shock. As a result interest rate elasticity on longer term yields is lower for these shocks.

Given the importance of the savings channel, next we examine an illustrative example of a more persistent shock that is more consistent with some recent fiscal policy changes. In particular, we consider a 10-year income tax reduction that is calibrated to match the average size and distribution of tax revenues reductions generated by the Tax Cut and Jobs Act (TCJA).[4] We find that for the first ten years the spot interest rate falls as the savings channel builds during the period in which taxes are lower. Although the strength of the savings channel begins to recede after the tax cut ends, it continues to dominate leading to lower spot interest rates for another decade as it takes time for households to spend down their excess savings. Thus, at the realization of the shock long yields fall with an elasticity of 1 percentage point of additional eventual debt leading to a 0.5 basis point fall in the 30-year yield. Thus, even though there has been a substantial increase in debt over the last few decades, these results suggest that the observed decline in interest rates may not be inconsistent but, instead, could be forthcoming as the imminent upward interest rate pressures have yet to materialize.

Next, we revisit the empirical estimation strategy to understand how previous attempts to summarize the dynamics between debt and interest rates with one elasticity relate to our findings. First, we build on the empirical strategy employed by the past literature, but we introduce several methodological updates. First, the past literature, notably in Engen and Hubbard (2004) and Laubach (2009), confronts endogeneity concerns from business cycle shocks that may occur at the same time as changes in public debt (e.g., due to countercyclical fiscal policy) by focusing on the relationship between the *expected levels* of interest rates and public debt as a share of GDP. Specifically, they consider long term expectations using CBO's projections for future debt and the forward rate (i.e. the projection of future interest rates based off the yield curve). By considering these long term expectations, the authors identify the relationship after the component of short term business cycle shock has presumably died out. However, estimating the relationship in levels creates a potential for bias due to non-stationarity in public debt-to-GDP ratios, and thus we alleviate this concern by estimating the relationship in first differences. Additionally, we use newer public debt projections from the CBO that extend the forecast horizon at which we estimate public debt and interest rates (from 5 years in prior work to 10 years in this paper).[5] The

[4]The bulk of the tax cuts in the TCJA were temporary for five years. However, policymakers are currently considering extending those provisions. Thus, we examine a shock somewhat more persistent than the original legislation.

[5]In addition, reformulating the empirical specification in first differences additionally allows us to

ten year projections are more consistent with the overall effect on long run U.S. debt, particularly because deficit financed fiscal policies tend to be in effect past five years. With our updated methodology and data, we find that estimated elasticities of 1 to 1.5 basis points, which are considerably smaller than the previously found estimates of 2 to 5 basis points, primarily due to introducing longer forecast horizons.

This estimated elasticity provides insight into the average response of interest rates with respect to historical changes in public debt. However, as we demonstrate quantitatively, there can be considerable variation in these elasticities depending on the term of the interest rate and the type of fiscal policy shock generating the debt. Leveraging the first differences used by our empirical specification, we can decompose the public debt projections into contributions from different sources of changes in debt (tax revenues, mandatory outlays and discretionary outlays) and directly test whether they have the same effects. Although we are unable to statistically confirm the heterogeneity in part due to limited power from the data, we find considerable differences in point estimates across sources, ranging from -8 to 1.8 basis points. Moreover, we find that the estimates are sensitive to both how far in the future the forward rate is measured along with the length of the maturity. Thus, next we develop a quantitative model that is calibrated to match a large set of micro and macro moments from the U.S. economy in order to explore the mechanisms for these types of heterogeneity in the elasticities across policies and horizons.[6] Thus, taken as a whole both the quantitative and empirical results indicate that it is incorrect to consider one homogeneous interest rate elasticity for long term interest rates with respect to debt and instead one must differentiate across types of debt financed policies along with the terms of the interest rate even at the longer end of the yield curve.

Related Literature. The empirical contributions of this paper build from seminal papers that estimate the elasticity of the interest rate to changes in public debt, such as Engen and Hubbard (2004), Gale and Orszag (2004), and Laubach (2009). These papers devise strategies for isolating the effects of fiscal policy from other macroeconomic shocks that may cause the interest rate to vary, such as business cycle shocks. Their general approach was to regress the level of N-year ahead interest rates on the

decompose overall public debt into changes due to the primary debt and due to net interest payments. We focus only on the primary debt since net interest payments introduce a concern over reverse causality, that higher interest rates mechanically increase government debt by increasing government borrowing costs. Finally, we include additional controls to capture potential correlation with the business cycle and population growth.

[6] Our model follows Peterman and Sager (2022) by assuming that households have a life cycle consisting of decisions over how much to work, when to retire and collect social security, when to consume and how much to save when facing idiosyncratic labor income shocks with incomplete asset markets.

level of forecasted N-year ahead public debt, after controlling for a host of macroeconomic conditions relevant to interest rate determination, so that short-term shocks and fluctuations have little or no impact on the fiscal measures. Gamber and Seliski (2019) replicate the methodology from these seminal papers using more recent data and find evidence of a structural break after in estimates after the Great Recession. Relative to these papers, we make several methodological improvements that address estimation bias from non-stationarity in the evolution of public debt, and endogeneity in the public debt measure which includes a short-term interest rate component from the payment of preexisting debt obligations and can therefore the interest rate may respond macroeconomic shocks that are unrelated to the evolution of public debt.

In our empirical work we study heterogeneity in the interest rate elasticity due to wealth inequality, and then quantitatively examine model mechanisms that generate the relationship. Several recent papers have studied the relationship between inequality and public debt. Peterman and Sager (2022) calibrate a quantitative life cycle model with incomplete markets to match wealth and income distributions by age, and find that public debt lower welfare due to the variation of consumption by age and due to public debt's ability to widen wealth inequality. Mian et al. (2020) document that wealthy households generate a large stockpile of savings when top income shares increase, as these households have a higher propensity to save out of lifetime income, while Mian et al. (2021) develop a model in which households at the top of the wealth distribution eventually exert such a strong effect on the aggregate interest rate that households at the bottom of the wealth distribution begin to borrow at a higher rate and aggregate demand decreases. Finally, Fagereng et al. (2019) show in Norwegian data that saving rates increase with wealth when including capital gains in wealth measures, because wealthy households enjoy a flow of income from these assets.

Our paper is also related to a long-standing literature on characterizing empirical and optimal fiscal policy rules and, particularly, whether fiscal policy can be understood as a stationary process by which the government stabilizes debt-to-GDP over sufficiently long periods of time. Barro (1979) shows that even when Ricardian equivalence holds, public debt may increase the efficiency of the economy by smoothing distortionary tax rates over time, and that the public debt-to-GDP ratio is not mean reverting over time. In contrast, using data from 1916-1995, Bohn (1998) shows that the primary surplus increases with the debt-to-GDP ratio, and that the debt-to-GDP ratio is indeed stationary. Campbell et al. (2023) have recently updated these results and find that neither the primary surplus-to-GDP nor the public debt-to-GDP ratios are stationary, while the tax revenues are indeed stationary. In line with these previ-

ous papers, we take as given that public debt and its components are non-stationary and construct measures accordingly. We also find that tax revenues empirically explain changes in the interest rate once one deals with the endogeneity of net interest payments.

Discussion of Mechanisms. Before proceeding to develop a quantitative model in Section 2, it is useful to conceptualize the mechanisms by which public debt affects the interest rate. To start, we focus on the standard prediction that public debt crowds out the stock of capital held by private firms. The reasoning for this mechanism follows from the asset market clearing condition $A_t = K_t + B_t$, that equates the stock of household savings A_t with the demand for those loanable funds by both firms in the form of private capital K_t and the government in the firm of public debt B_t. Holding household savings constant, if the government increases its demand for funds by dB units then there are mechanically dB fewer units available to private firms as private capital. Thus an increase in public debt to $B_t + dB$ reduces the economy's amount of private capital to $K_t - dB$. Such a reallocation of private capital away from firms increases the marginal productivity of an additional unit of capital and thus the firms are willing to pay a higher interest rate, r_t, to households,

$$\frac{dr_t}{dB_t} = -\frac{dr_t}{dK_t} > 0 \, .$$

This mechanism illustrates the *direct crowd-out channel* by which public debt exerts upward pressure on the interest rate.

There are, however, significant simplifications employed by this stylized illustration of private capital crowd-out. First, the illustration abstracts from the equilibrium response of household savings decisions to higher public debt. Representing the amount of private capital that is crowded out by an increase in public debt by dK_t/dB_t, we can rearrange the asset market clearing condition $K_t = A_t - B_t$ to derive the change in the interest rate from an increase in public debt when household savings is allowed to respond,

$$\frac{dr_t}{dB_t} = \underbrace{\left(\frac{dA_t}{dB_t} - 1\right)}_{\equiv dK_t/dB_t} \cdot \frac{dr_t}{dK_t} \, .$$

The expression $dK_t/dB_t = dA_t/dB_t - 1$ indicates that the change in private capital is derived from the response of household savings to increased public debt dA_t/dB_t, in addition to the standard direct crowd-out channel represented by $-dB_t/dB_t = -1$.

Clearly, we obtain an interest rate response corresponding to the direct crowd-out case when savings does not respond to public debt, e.g., if $dA_t/dB_t = 0$. However, if household savings were to increase in response to increased public debt, $dA_t/dB_t > 0$, as would naturally be the response to a higher interest rate arising from the direct crowd-out channel, the increase in loanable funds to firms mitigates the crowd-out of private capital and thus the interest rate may increase by less. We call this the *savings channel* by which public debt exerts downward pressure on the interest rate.

Thus, the overall response of interest rates to increased public debt depends on the degree of diminishing marginal returns to an additional unit of capital and the magnitude of the savings channel. Generally, for a comparative static on long-run outcomes, the savings channel will be bounded between zero and one, e.g., $0 < dA_t/dB_t < dB_t/dB_t = 1$.

Second, the illustration is a comparative static on long-run outcomes that abstracts from the economy's gradual transition to a new steady state with higher public debt. Indeed, the transition dynamics of the interest rate could critically depend on the type of fiscal policy and whether (and how) that policy directly impacts households' savings decisions, since the fiscal policy can have a substantial effect on the savings channel over the transition. For example, an increase in public debt that finances lump sum transfers to households will result in a significant savings channel as households save a large portion of the transfer to afford a higher level of lifetime consumption. By boosting household savings, the lump sum transfer exerts downward pressure on the interest rate and slows its growth toward its new steady state value. In contrast, an increase in public debt that finances government expenditures does not directly boost household savings. Instead, households respond to a higher interest rate caused by the direct crowd-out channel by gradually increasing their savings, which generates a weaker savings effect as the interest rate slowly declines toward its new steady state value.

We will quantify these mechanisms after developing and calibrating our model that introduces a large set of ingredients to inform savings behavior.

Layout of Paper. The remainder of this paper is organized as follows. Section 2 details our model and characterizes the economic environment. Section 3 quantifies the model and discusses policy counterfactuals characterizing the relationship between government debt and interest rates. After finding that the dynamic effects of government debt on interest rates are specific to the type of policy shocks, section 4 updates the canonical empirical estimates of this relationship and compares them to our model results. Section 5 concludes.

2 Economic Environment

We will consider the economic effects of a change in public debt within the quantitative life cycle model with incomplete markets.[7] We define the stationary recursive competitive equilibrium around a balanced growth path. For ease of explication, we present the detrended stationary recursive competitive equilibrium and suppress time-dependence in our notation. However, we make explicit any primitives that require additional assumptions for the model to be consistent with balanced growth.

Within this environment, we will measure how the interest rate varies with different fiscal policies that each create additional government debt, and quantify the mechanisms by which households respond to these policy changes. We use the model to conduct three different policy experiments that increase the government's debt by the same amount: a temporary one period increase in the deficit from an increase in government purchases, a decrease in tax revenue, and an increase in lump sum transfers. Each policy experiment induces a transition to a steady state with a new higher level of debt but a different interest rate elasticity.

2.1 Production

We assume there exists a large number of firms that sell a single consumption good in a perfectly competitive product market, purchase inputs from perfectly competitive factor markets, and each operate an identical constant returns to scale production technology, $Y = ZF(K, L)$. These assumptions on primitives admit a representative firm that chooses capital (K) and labor (L) inputs in order to maximize profits, given an interest rate r, a wage rate w, a level of total factor productivity Z, and capital depreciation rate $\delta \in (0, 1)$. Following Aiyagari and McGrattan (1998), we assume that total factor productivity grows over time at rate $g_z > 0$ which generates output growth at rate denoted by $g_y > 0$.

2.2 Households

Demographics: Time is discrete and each model period represents a year. The household is the unit of measurement within our model and each period the economy is inhabited by J overlapping generations of households.[8] Each period a new cohort

[7] This model was initially developed in Peterman and Sager (2022) for studying the optimal level of public debt.

[8] Consistent with most of the literature, we model household decisions as joint and assume away intrahousehold frictions that can distort allocations of consumption, hours and savings.

is born and the size of each successive newly born cohort grows at a constant rate $g_n > 0$. All households live for a maximum of J periods, where age is indexed by $j = 1, \ldots, J$. All living households face mortality risk such that, conditional on living to age j, ψ_j is the probability of an age-j household living to age $j+1$. Accordingly, the terminal-age survival probability is $\psi_J = 0$.

Preferences: Households derive utility from lifetime paths of consumption, labor hours, retirement status and assets, denoted $\{c_j, h_j, d_j, a_{j+1}\}_{j=1}^{J}$, according to the following preferences:

$$\mathbb{E}_1 \sum_{j=1}^{J} \beta^{j-1} \left(\prod_{i=1}^{j-1} \psi_i \right) \Big[u_j(c_j) - v(h_j, d_j) + \beta(1-\psi_j)\phi(a_{j+1}) \Big]$$

where β is the time discount factor, $u_j(c)$ and $v(h,d)$ are instantaneous utility functions over consumption, labor hours and retirement status, respectively, satisfying standard conditions.[9] Households derive instantaneous utility $\phi(a_{j+1})$ from bequeathing their assets, should they receive a mortality shock and die. Expectations are taken with respect to the stochastic processes governing labor productivity.

Between the ages $\underline{J}_{ret}$ and $\bar{J}_{ret}$, a household makes the irreversible decision whether to retire, and once retired an agent no longer has the option to work. Any household that has not retired by age $\bar{J}_{ret}$ is forced to retire. We denote the age at which a household chooses to retire by J_{ret} and define $d_j \equiv \mathbb{1}(j < J_{ret})$ to be an indicator variable that equals one when a household chooses to continue working and zero upon retirement.

Labor Earnings: Non-retired households are endowed with one unit of time per period, which they split between leisure and market labor. During each period of working life, an household's labor earnings are we_jh_j, where w is the wage rate per efficiency unit of labor, e_j is the household's idiosyncratic labor productivity drawn at age j, and h_j is the share of the time endowment that the household chooses to work at age j.

Following Kaplan (2012), we assume that labor productivity shocks can be decomposed into four sources:

$$\log(e_j) = \kappa + \theta_j + \nu_j + \epsilon_j$$

[9] The instantaneous utility over consumption accounts for changes in household size through an adult equivalent normalization and therefore varies with age, which we detail in Section 3. Furthermore, this definition of preferences embeds the assumption that the disutility of labor, $v(h,d)$, and utility from bequests, $\phi(a)$, grow over time. In particular, along a balanced growth path they will grow at the same rate as the utility over individual consumption (for more detail see Appendix B).

where (i) $\kappa \overset{iid}{\sim} \mathcal{N}(0, \sigma_\kappa^2)$ is an individual-specific fixed effect that is drawn once when the household enters the economy and remains fixed, (ii) $\{\theta_j\}_{j=1}^J$ is an age-specific fixed effect that evolves over the life cycle in a predetermined manner, (iii) ν_j is a persistent shock that follows a first-order Markov process, and (iv) $\epsilon_j \overset{iid}{\sim} \mathcal{N}(0, \sigma_\epsilon^2)$ is a transitory shock that is drawn each period.

For notational compactness, we denote the relevant state as a vector $\varepsilon_j = (\kappa, \theta_j, \nu_j, \epsilon_j)$ that contains each element necessary for computing contemporaneous labor earnings, $e_j \equiv e(\varepsilon_j)$, and forming expectations about future labor earnings. Denote the Markov process governing the process for ε by $\pi_j(\varepsilon_{j+1}|\varepsilon_j)$ for each ε_j, ε_{j+1} and for each $j = 1, \ldots, \bar{J}_{ret}$.

Assets and Bequests: Households have access to a single asset, a non-contingent one-period bond denoted a_j with a market determined rate of return of r. Households are endowed with zero initial wealth, such that $a_1 = 0$ for each household. Working households may take on a net debt position, in which case they are subject to a borrowing constraint that requires their net assets be bounded below by $\underline{a} \in \mathbb{R}$ and face an interest rate on repayment of r/ψ_j at each age. For notational convenience we define $\tilde{r}_j(a)$ as the interest rate that takes on a value of r when $a \geq 0$ and r/ψ_j when $\underline{a} \leq a < 0$.

Households may hold assets when they die, which are redistributed to living households as bequests. During their lifetimes, households receive a flow of bequests at each age (including upon entering the economy). This flow of bequests varies with the individual-specific fixed effect component of labor productivity and age according to the function $b_j(\kappa)$.

Medical Expenditures: Given the potential for households' late-life expenses to affect lifetime savings decisions, we incorporate out-of-pocket medical expenditures into the model. Households face medical expenses that deterministically as a function of age, denoted μ_j, beginning at age 70 and face zero medical expenses prior to age 70. We focus on late-life medical expenditures because previous work, notably DeNardi, French, and Jones (2010), finds that average out-of-pocket expenditures rise rapidly after age 70 and begin to constitute a notable share of household expenditures.[10]

[10]For example, DeNardi et al. (2016) find that average out-of-pocket medical expenditures rise from $1000 at age 75 to $17,700 by age 100.

2.3 Government

The government (i) consumes resources G, (ii) collects Social Security taxes and distributes Social Security payments to retired households, (iii) distributes bequests across living households, (iv) distributes transfers to cover a portion of medical expenses and (v) collects income taxes from each household.

Social Security: The model's Social Security Program is financed with a payroll tax, τ_{ss}, and taxable income is capped per-period by $\bar{m}$. Therefore, the consumer pays a payroll tax given by: $\tau_{ss} \min\{weh, \bar{m}\}$.

Social security payments are computed using the averaged indexed monthly earnings (AIME) that summarizes a household's lifetime labor earnings. Following Huggett and Parra (2010) and Kitao (2014), the AIME is denoted by $\{m_j\}_{j=1}^{J}$, has an initial value $m_1 = 0$ and evolves as follows:

$$m_{j+1} = \left\{ \begin{array}{ll} \frac{1}{j}\left(\min\{we_j h_j, \bar{m}\} + (j-1)m_j\right) & \text{for } j \leq 35 \\ \max\left\{m_j, \frac{1}{j}\left(\min\{we_j h_j, \bar{m}\} + (j-1)m_j\right)\right\} & \text{for } j \in (35, J_{ret}) \\ m_j & \text{for } j \geq J_{ret} \end{array} \right\}$$

The AIME is a state variable for determining future benefits. Benefits are derived from the AIME in two steps: a base payment is determined, and then the base payment is adjusted according to a household's retirement age. The base payment, denoted by $b^{ss}_{base}(m_{J_{ret}})$, is computed as a piecewise-linear function over the household's average labor earnings at retirement $m_{J_{ret}}$:

$$b^{ss}_{base}(m_{J_{ret}}) = \left\{ \begin{array}{lll} \tau_{r1} m_{J_{ret}} & \text{for} & m_{J_{ret}} \in [0, b^{ss}_1) \\ \tau_{r1} b^{ss}_1 + \tau_{r2}(m_{J_{ret}} - b^{ss}_1) & \text{for} & m_{J_{ret}} \in [b^{ss}_1, b^{ss}_2) \\ \tau_{r1} b^{ss}_1 + \tau_{r2} b^{ss}_2 + \tau_{r3}(m_{J_{ret}} - b^{ss}_1 - b^{ss}_2) & \text{for} & m_{J_{ret}} \in [b^{ss}_2, b^{ss}_3) \\ \tau_{r1} b^{ss}_1 + \tau_{r2} b^{ss}_2 + \tau_{r3} b^{ss}_3 & \text{for} & m_{J_{ret}} \geq b^{ss}_3 \end{array} \right\}$$

The base payment is adjusted according to a household's retirement age to penalize early retirement and credit delayed retirement. The adjustment is given by:

$$b_{ss}(m_{J_{ret}}) = \left\{ \begin{array}{lll} (1 - D_1(J_{nra} - J_{ret}))b^{ss}_{base}(m_{J_{ret}}) & \text{for} & \underline{J}_{ret} \leq J_{ret} < J_{nra} \\ (1 + D_2(J_{ret} - J_{nra}))b^{ss}_{base}(m_{J_{ret}}) & \text{for} & J_{nra} \leq J_{ret} \leq \bar{J}_{ret} \end{array} \right\}$$

where $D_i(\cdot)$ are functions governing the benefits penalty or credit, and J_{nra} is the "normal retirement age".

Income Taxation: Taxable income is defined as the sum of labor income and capital income from assets and bequests, net of social security contributions from an employer which are considered half of the total contribution:

$$y_j(h,a,\varepsilon,d) \equiv \begin{cases} we(\varepsilon)h + \tilde{r}_j(a)a + rb_j(\varepsilon) - \frac{\tau_{ss}}{2}\min\{we(\varepsilon)h, \bar{m}\} & \text{if } d=1 \\ r(a+b_j(\varepsilon)) & \text{if } d=0 \end{cases}$$

The government taxes each household's taxable income according to an increasing and concave function, $\Upsilon(y_j(h,a,e,d))$.

Medical Transfers: In the spirit of U.S. Medicaid program and other means-tested public assistance programs, low-income households are given transfers in order to allow them to pay for medical expenses and still have resources available for consumption. The government provides transfers to retired households after the age $\bar{J}_{ret}$ that guarantee a minimum consumption level, denoted $\underline{c} > 0$, after paying medical expenditures and taxes. The transfers are given by,

$$Tr_j(a,m,\varepsilon) = \max\left\{0,\ \underline{c} - \left[b_{ss}(m) + r(a+b_j(\varepsilon)) - \Upsilon\left(r(a+b_j(\varepsilon))\right) - \mu_j\right]\right\}$$

where households do not receive a transfer if their after-tax income net of medical expenses allows for consumption in excess of the policy's minimum guarantee $(\underline{c})$.[11]

Public Savings and Budget Balance: Each period, the government has a debt balance B and saves or borrows (denoted B') at the market interest rate r. If the government borrows, then $B' < 0$ and the government repays rB' next period. If the government saves, then $B' > 0$ and the government collects asset income rB' next period. The resulting government budget constraint is:

$$G + Tr + (1+g_y)B' - B = rB + R \tag{1}$$

where R is aggregate revenues from income taxation and G is the level of government expenditures that are independent from aggregate medical transfers, Tr.[12] The

[11] We only model these transfers for non-working households because the consumption floor is sufficiently low to be effectively irrelevant for working households. Working households only need to spend a small fraction of their time endowment to generate enough income to relax the consumption floor and, thus, there would be no medical related transfers to working households even if explicitly allowed.

[12] We assume government expenditures exclusive of aggregate medical transfers, G, are unproductive.

model's Social Security system is self-financing and therefore does not appear in the governmental budget constraint.

2.4 Consumer's Problem

The household's state is $(a, \varepsilon, m, d_{-1})$ and consists of asset holdings a, labor productivity shocks $\varepsilon \equiv (\kappa, \theta, \nu, \epsilon)$, Social Security contribution (AIME) variable m, and retirement status d_{-1}.

Prior to retirement age, the age-j household's recursive problem is:

$$\begin{aligned} V_j(a, \varepsilon, m, 1) = \max_{c, a', h} \; & \big[u_j(c) - v(h, 1)\big] + [\beta(1+g_w)^{1-\sigma}]\psi_j \sum_{\varepsilon'} \pi_j(\varepsilon'|\varepsilon) V_{j+1}(a', \varepsilon', m', 1) \\ & + [\beta(1+g_w)^{1-\sigma}](1-\psi_j)\phi(a') \end{aligned} \tag{2}$$

subject to

$$\begin{aligned} c + (1+g_w)a' &\leq we(\varepsilon)h + \big(1 + \tilde{r}_j(a)\big)a + (1+r)b_j(\varepsilon) - \tau_{ss}\min\{we(\varepsilon)h, \bar{m}\} - \Upsilon(y_j(h, a, \varepsilon, 1)) \\ a' &\geq \underline{a} \end{aligned}$$

where g_w is the steady state rate of wage growth and σ is the coefficient of relative risk aversion (see Section 3 for the underlying utility specification).

A working household chooses whether to retire (indicated by $d = 0$) or not ($d = 1$) between ages $\underline{J}_{ret}$ and $\bar{J}_{ret}$ or faces mandatory retirement after age $\bar{J}_{ret}$. The retired age-j household's stationary recursive problem is:

$$\begin{aligned} V_j(a, \varepsilon, m, 0) = \max_{c, a'} \; & u_j(c) + [\beta(1+g_w)^{1-\sigma}]\psi_j V_{j+1}(a', \varepsilon, m, 0) \\ & + [\beta(1+g_w)^{1-\sigma}](1-\psi_j)\phi(a') \end{aligned} \tag{3}$$

subject to

$$\begin{aligned} c + (1+g_w)a' &\leq (1+r)(a + b_j(\varepsilon)) + b_{ss}(m) - \Upsilon\left(r(a + b_j(\varepsilon))\right) + \big(Tr_j(a, m, \varepsilon) - \mu_j\big) \\ a' &\geq \underline{a} \end{aligned}$$

where the post-retirement state ε simply records productivity type κ for determining bequest inflows. The construction of the stationary Bellman equation is presented in

Two recent papers, Röhrs and Winter (2017) and Chatterjee, Gibson, and Rioja (2017) have relaxed the standard Ramsey assumption that government expenditures are unproductive. Both papers show that public savings is optimal with productive government expenditures, intuitively because there is an additional benefit to aggregate output.

Appendix B.

2.5 Stationary Recursive Competitive Equilibrium

We study a stationary equilibrium along a balanced growth path in which all aggregate variables grow at the same rate as output, g_y.

Households are heterogeneous with respect to their age $j \in \mathbf{J} \equiv \{1, \ldots, J\}$, wealth $a \in \mathbf{A}$, labor productivity $\varepsilon \in \mathbf{E}$, average lifetime earnings $m \in \mathbf{X}$, and retirement status $d \in \mathbf{D} \equiv \{0, 1\}$. Let $\mathbf{S} \equiv \mathbf{A} \times \mathbf{E} \times \mathbf{X} \times \mathbf{D}$ be the state space and $\mathcal{B}(\mathbf{S})$ be the Borel σ-algebra on $\mathbf{S}$. Let $\mathbf{M}$ be the set of probability measures on $(\mathbf{S}, \mathcal{B}(\mathbf{S}))$. Then $(\mathbf{S}, \mathcal{B}(\mathbf{S}), \lambda_j)$ is a probability space in which $\lambda_j(S) \in \mathbf{M}$ is a probability measure defined on subsets of the state space, $S \in \mathcal{B}(\mathbf{S})$, that describes the distribution of individual states across age-j households. Denote the fraction of the population that is age $j \in \mathbf{J}$ by ω_j. For each set $S \in \mathcal{B}(\mathbf{S})$, $\omega_j \lambda_j(S)$ is the fraction of age $j \in \mathbf{J}$ and type $S \in \mathbf{S}$ households in the economy. We can now define a stationary recursive competitive equilibrium of the economy.

Definition (Equilibrium): Given a government policy (G, B, B'), a *stationary recursive competitive equilibrium* is (i) an allocation for consumers described by policy functions $\{c_j, a'_j, h_j, d_j\}_{j=1}^J$ and consumer value function $\{V_j\}_{j=1}^J$, (ii) an allocation for the representative firm (K, L), (iii) prices (w, r), (iv) bequests $b_j(\kappa)$, (v) a government policy $(\Upsilon, \tau_{ss}, b_{ss}, \underline{c})$, and (vi) distributions over households' state vector at each age $\{\lambda_j\}_{j=1}^J$ that satisfy:

(a) Given prices, policies and bequests, $V_j(a, \varepsilon, m, d_{-1})$ solves the Bellman equation (2) and (3) with associated policy functions $c_j(a, \varepsilon, m, d_{-1})$, $a'_j(a, \varepsilon, m, d_{-1})$, $h_j(a, \varepsilon, m, d_{-1})$ and $d_j(a, \varepsilon, m, d_{-1})$.

(b) Given prices (w, r), the representative firm's allocation minimizes cost, $r = ZF_K(K, L) - \delta$ and $w = ZF_L(K, L)$.

(c) Total bequests from households of type-κ who die at the end of this period are distributed across next period's living households of type-κ according to the function $b_j(\kappa)$. Then the following condition must hold for each κ,

$$(1 + g_n)\pi(\kappa) \sum_{j=1}^{J} \omega_j b_j(\kappa) = \sum_{j=1}^{J} \omega_j (1 - \psi_j) \int a'_j(a, \varepsilon, m, d_{-1}|\kappa) \mathbf{d}\lambda_j(a, \varepsilon, m, d_{-1}|\kappa),$$

such that living households' savings equal next period's wealth net of bequests,

$$\sum_{j=1}^{J} \omega_j \psi_j \int a'_j(a,\varepsilon,m,d_{-1}|\kappa)\mathbf{d}\lambda_j(a,\varepsilon,m,d_{-1}|\kappa) = (1+g_n)\sum_{j=1}^{J}\omega_j \int a\mathbf{d}\lambda_j(a,\varepsilon,m,d_{-1}|\kappa)$$

where $\lambda(a,\varepsilon,m,d_{-1}|\kappa)$ denotes the mass over (a,ε,m,d_{-1}) when holding the value of κ fixed, such that $\pi(\kappa) = \int \mathbf{d}\lambda(a,\varepsilon,m,d_{-1}|\kappa)$ is the measure of type-κ households.

(d) Government policies satisfy budget balance in equation (1), where aggregate income tax revenue is given by:

$$R \equiv \sum_{j=1}^{J} \omega_j \int \Upsilon\left(y_j\left(h_j(a,\varepsilon,m,d_{-1}),a,\varepsilon,d_j(a,\varepsilon,m,d_{-1})\right)\right)\ \mathbf{d}\lambda_j(a,\varepsilon,m,d_{-1}), \tag{4}$$

and aggregate medical expenditure transfers are given by:

$$Tr \equiv \sum_{j=\bar{J}_{ret}+1}^{J} \omega_j \int Tr_j(a,m,\varepsilon;\underline{c})\ \mathbf{d}\lambda_j(a,\varepsilon,m,d_{-1}).$$

(e) Social security is self-financing:

$$\sum_{j=1}^{J} \omega_j \int d_j(a,\varepsilon,m,d_{-1})\tau_{ss}\min\{we(\varepsilon)h_j(a,\varepsilon,m,d_{-1}),\bar{m}\}\mathbf{d}\lambda_j(a,\varepsilon,m,d_{-1})$$

$$= \sum_{j=1}^{J} \omega_j \int (1-d_j(a,\varepsilon,m,d_{-1}))b_{ss}(m)\mathbf{d}\lambda_j(a,\varepsilon,m,d_{-1}). \tag{5}$$

(f) Given policies and allocations, prices clear asset and labor markets:

$$K - B = \sum_{j=1}^{J} \omega_j \int \left(a + b_j(\kappa)\right)\ \mathbf{d}\lambda_j(a,\varepsilon,m,d_{-1}) \tag{6}$$

$$L = \sum_{j=1}^{J} \omega_j \int d_j(a,\varepsilon,m,d_{-1})e(\varepsilon)h_j(a,\varepsilon,m,d_{-1})\ \mathbf{d}\lambda_j(a,\varepsilon,m,d_{-1}) \tag{7}$$

and the allocation satisfies the resource constraint (guaranteed by Walras' Law):

$$C + (1+g_y)K' + G + Tr = ZF(K,L) + (1-\delta)K \tag{8}$$

where

$$C = \sum_{j=1}^{J} \omega_j \int c_j(a, \varepsilon, m, d_{-1}) \mathbf{d}\lambda_j(a, \varepsilon, m, d_{-1}).$$

(g) Given consumer policy functions, distributions across age j households $\{\lambda_j\}_{j=1}^{J}$ are given recursively from the law of motion $T_j^* : \mathbf{M} \to \mathbf{M}$ for all $j \in \mathbf{J}$ such that T_j^* is given by:

$$\lambda_{j+1}(\mathcal{A} \times \mathcal{E} \times \mathcal{X} \times \mathcal{D}) = \sum_{d_{-1} \in \{0,1\}} \int_{A \times E \times X} Q_j\left((a, \varepsilon, m, d_{-1}), \mathcal{A} \times \mathcal{E} \times \mathcal{X} \times \mathcal{D}\right) \mathbf{d}\lambda_j$$

where $\mathcal{S} \equiv \mathcal{A} \times \mathcal{E} \times \mathcal{X} \times \mathcal{D} \subset \mathbf{S}$, and $Q_j : \mathbf{S} \times \mathcal{B}(\mathbf{S}) \to [0, 1]$ is a transition function on $(\mathbf{S}, \mathcal{B}(\mathbf{S}))$ that gives the probability that an age-j household with current state $\mathbf{s} \equiv (a, \varepsilon, m, d_{-1})$ transits to the set $\mathcal{S} \subset \mathbf{S}$ at age $j + 1$. The transition function is given by:

$$Q_j\left((a, \varepsilon, m, d_{-1}), \mathcal{S}\right) = \left\{ \begin{array}{cc} \psi_j \cdot \pi_j(\mathcal{E}|\varepsilon)^{d_{-1}} & \text{if } a_j'(\mathbf{s}) \in \mathcal{A},\ m_j'(\mathbf{s}) \in \mathcal{X},\ d_j(\mathbf{s}) \in \mathcal{D} \\ 0 & \text{otherwise} \end{array} \right\}$$

where households that continue working and transition to set $\mathcal{E}$ choose $d_j(\mathbf{s}) = 1$, while households that transition from working life to retirement choose $d_j(\mathbf{s}) = 0$. For $j = 1$, the distribution λ_j reflects the invariant distribution $\pi_{ss}(\varepsilon)$ of initial labor productivity over $\varepsilon = (\kappa, \theta_1, 0, \epsilon_1)$.

(h) Aggregate capital, governmental debt, prices and the distribution over consumers are stationary, such that $K' = K$, $B' = B$, $w' = w$, $r' = r$, and $\lambda_j' = \lambda_j$ for all $j \in \mathbf{J}$.

3 Calibration

3.1 Parameters and Functional Specifications

In this section we calibrate the life cycle model along the balanced growth path. There are two subsets of parameters. One subset is set directly from empirical estimates, while the other subset is chosen so that the model matches a number of empirical moments. Table 1 summarizes the target, source and value for each parameter, and Table 2 evaluates model fit by comparing model generated moments to empirical moments.

Production: We assume that the aggregate production function is Cobb-Douglas of the form $F(K, L) = K^{\alpha} L^{1-\alpha}$ where $\alpha = 0.36$ is the income share accruing to capital and total factor productivity is normalized to one, $Z = 1$. The depreciation rate is set to $\delta = 0.0833$ which allows the model to match the empirically observed investment-to-output ratio of 24.2%.

Demographics: Households enter the economy at age 20 (or model age $j = 1$) and we set conditional survival probabilities $\{\psi_j\}_{j=1}^{J-1}$ to match survival rates for households in the data from Bell and Miller (2002). We impose that surviving households exogenously die after age 100 (or model age $J = 81$) by setting the terminal survival probability $\psi_J = 0$. Household mortality is defined as either both members of a married household dying or the sole remaining adult of a household dying. Household mortality rates account for demographic changes within the household and variation in individuals' mortality rates by age and sex. We set the population growth rate to $g_n = 0.011$ to match annual population growth in the US.

Preferences: The utility function is separable in the utility over consumption, $u_j(c)$, labor hours and retirement, $v(h, d)$, and bequests, $\phi(a')$. We parameterize the utility specification as,

$$u_j(c) = \frac{(c/n_j)^{1-\sigma}}{1-\sigma}$$

$$v(h, d) = \chi_1 \frac{h^{1+\frac{1}{\gamma}}}{1+\frac{1}{\gamma}} + d\chi_2$$

$$\phi(a') = \chi_b \frac{(\chi_a + a')^{1-\sigma}}{1-\sigma} .$$

Utility over consumption is age-dependent in order to capture how a changing average household size over the life cycle affects household consumption decisions. Accordingly, we use the adult equivalent scale, n_j, to adjust consumption in the utility function. Following Hur (2018), we compute the adult equivalent scale at each age (of the head of household) to convert households of varying sizes into a standardized measure,

$$n_j \equiv \left[\omega_j^{single} \cdot 1\right] + \left[(1 - \omega_j^{single}) \cdot 1.5\right] + (1/3) n_j^c$$

where ω_j^{single} is the fraction of single-adult households with an age-j head of household, and n_j^c is the average number of children in a household with an age-j head of household.

Table 1: Calibration Targets and Parameters for Baseline Economy.

Description	Parameter	Value	Target or Source
Demographics			
Maximum Age	J	81 (100)	By Assumption
Min/Max Retirement Age	$\underline{J}_{ret}, \bar{J}_{ret}$	43, 51 (62, 70)	Social Security Program
Population Growth	g_n	1.1%	Conesa et al (2009)
Survival Rate	$\{\psi_j\}_{j=1}^{J}$		Bell and Miller (2002)
Preferences and Borrowing			
Coefficient of RRA	σ	2.0	Conesa et al (2009)
Frisch Elasticity	γ	0.5	Kaplan (2012)
Coefficient of Labor Disutility	χ_1	159.8	Avg. Hours Worked = 0.2687
Fixed Utility Cost of Labor	χ_2	2.41	70% retire by NRA
Discount Factor	β	1.005	Capital/Output = 2.9
Borrowing Limit	$\underline{a}$	-0.50	Avg. Borrowing/Avg. Wealth = -0.0483
Bequests			
Coefficient on Bequest Utility	χ_b	30	Avg. Bequest/ Avg. Wealth = 0.0088
Non-Homotheticity of Bequests	χ_a	4.85	Bequests, 90th pct/Avg. Labor Earnings = 10.7
Bequest Age-Profile	$\{b_j(\kappa)\}_{j=1}^{J}$		Feiveson and Sabelhaus (2019)
Production Technology			
Capital Share	α	0.36	NIPA
Capital Depreciation Rate	δ	0.0833	Investment/Output = 0.242
Output Growth	g_y	1.85%	NIPA
Labor Productivity			
Persistent Shock, autocorrelation	ρ	0.960	PSID
Persistent Shock, variance	σ_ν^2	0.014	PSID
Permanent Shock, variance	σ_κ^2	0.054	PSID
Transitory Shock, variance	σ_ϵ^2	0.086	PSID
Mean Earnings, Age Profile	$\{\theta\}_{j=1}^{\bar{J}_{ret}}$		PSID
High Labor Productivity Shock	ν_{max}	8.8	Top 40% Wealth = 94.6%
High Labor Prod. Persistence	p_S	99.4%	Top 20% Labor Share = 63.5%
High Labor Prod. Probability	π_S	0.5%	
Government Budget			
Government Consumption	$(G+Tr)/Y$	0.155	NIPA Average 1998-2007
Government Savings	B/Y	-0.667	NIPA Average 1998-2007
Marginal Income Tax	τ_0	0.258	Gouveia and Strauss (1994)
Income Tax Progressivity	τ_1	0.786	Gouveia and Strauss (1994)
Income Tax Progressivity	τ_2	3.324	Balanced Budget
Medical Expenses			
Consumption Floor (% of wL)	$\underline{c}$	15%	Kopecky and Koreshkova (2014)
Average Household Expenses	$\{\mu_j\}_{j=\bar{J}_{ret}}^{J}$		HRS-AHEAD
Social Security			
Payroll Tax	τ_{ss}	0.151	Balanced Budget
SS Replacement Rates	$\{\tau_{ri}\}_{i=1}^{3}$	(0.9,0.32,0.15)	Social Security Program
SS Replacement Bend Points	$\{b_i^{ss}\}_{i=1}^{3}$	(0.21,1.29,2.42)	Social Security Program
SS Early Retirement Penalty	$\{D_i\}_{i=1}^{2}$	See Text	Social Security Program

Utility over consumption is a CRRA specification with a coefficient of relative risk aversion $\sigma = 2$, which is consistent with Conesa et al. (2009) and Aiyagari and McGrattan (1998). Disutility over labor exhibits a constant intensive margin Frisch elasticity. We choose $\gamma = 0.5$ such that the Frisch elasticity consistent with the majority of the related literature as well as the estimates in Kaplan (2012).[13]

We calibrate the labor disutility parameter χ_1 so that the cross sectional average of hours is 0.2687 of the time endowment, which we find in the PSID.[14] Finally, χ_2 is the fixed utility cost of not being retired, which generates an active extensive margin by introducing a non-convexity in the utility function. We choose χ_2 to match the empirical observation that seventy percent of the population has retired by age 66.

We employ a standard utility function over bequests (c.f., DeNardi (2004)), where χ_b is a coefficient that determines the level of utility over bequests, and χ_a is a non-homothetic term that measures the extent to which bequests are luxury goods. We choose χ_b to match the bequest-to-wealth ratio of 0.0088 (see Gale and Scholz (1994)), and choose χ_a to match the 90th percentile of the bequest distribution for households over the age of 70, normalized by average labor income, of 10.7.

Bequests: The total level of bequests to households is determined by the total amount of wealth held by households who died in the previous period. Bequests are redistributed to households with the same individual-specific labor productivity type, κ. For each type, these bequests are allocated to living households to match shares of bequests received by age in the Survey of Consumer Finances, according to the function $b_j(\kappa)$.

Liquidity Constraints: We choose the lower bound on assets $\underline{a} \leq 0$ so that the model matches the level of net wealth conditional on borrowing, relative to the total level of private wealth. This ratio is -4.83% in the 2007 Survey of Consumer Finances for ages 20 to 60.

Labor Productivity Process: The parameters for the labor productivity process are jointly estimated from the PSID using the procedure described in Kaplan (2012). However, the parameters are determined on a household-based concept rather than an individual-based concept.[15]

[13] Although Peterman (2016) finds that estimates of the Frisch elasticity tend to be larger for non-primary earners, the Frisch elasticity of 0.5 still falls within the range of these estimates. Moreover, we have found that changing the Frisch elasticity within this range did not materially change the main results herein.

[14] To compute hours worked in the PSID, we normalize a households' annual hours worked (by 40 hours per week, 52 weeks per year) and number of potential earners.

[15] We thank Aladangady, Bi, and Peterman (2020) for sharing their estimates from preliminary work.

A well known problem with a log-normal income process is that it cannot generate the degree of wealth and labor income inequality we observe in the data. To match the wealth and labor income distributions, we follow Castañeda, Díaz-Giménez, and Ríos-Rull (2003) and Kindermann and Krueger (2020), by modeling the persistent component of the labor productivity process in two parts. First, we include a standard persistent component of labor productivity which follows a standard first-order autoregressive process given by $\nu_{j+1} = \rho\nu_j + \eta_{j+1}$ with $\eta \overset{iid}{\sim} \mathcal{N}(0, \sigma_\nu^2)$ and $\eta_1 = 0$, which we estimate using the PSID. Second, we include an extremely high labor productivity state.[16] We refer to this additional high labor productivity state as a *superstar shock* and we set the probability of receiving the superstar shock to 0.5%, which implies that 20% of households that retire at the normal retirement age experienced being a superstar at some point during their working lifetime. We choose the value of the superstar shock (ν_{max}) and probability of remaining a superstar (p_S) so that the top 40% of the population holds 94.6% of total wealth and the top 20% of the population receives 53.5% of total labor income (see Krueger, Mitman, and Perri (2016)).[17] The value of ν_{max} is 8.8, which implies that a superstar earns approximately 9 times more than the average non-superstar, and the value of p_S is 0.994, which implies that the superstar state is extremely persistent.

Government Debt and Income Taxation: Consistent with Aiyagari and McGrattan (1998) we set government debt equal to two-thirds of output. We set the sum of government consumption and aggregate medial transfers to 15.5 percent of output. This ratio corresponds to the average of government expenditures to GDP from 1998 through 2007.[18]

We model the income tax using a standard functional form from Gouveia and Strauss (1994),

$$\Upsilon(y) = \tau_0 \left(y - \left(y^{-\tau_1} + \tau_2 \right)^{-\frac{1}{\tau_1}} \right) .$$

We parameterize the function with the authors' estimates of $\tau_0 = 0.258$ and $\tau_1 = 0.768$,

[16] Preference heterogeneity is an alternate way to introduce a skewed wealth. However, there are two downsides to using preference heterogeneity. First, in a model similar to ours that excludes altruism, Hendricks (2007) demonstrates that matching the wealth distribution requires including a large mass of both patient and impatient agents with a considerably larger gap in patience between these groups than is consistent with empirical estimates. Second, it is unclear what discount rate should be used to measure social welfare.

[17] We assume that no household enters the economy as a superstar. We also assume that upon exiting the superstar state, households transition to the median persistent labor productivity state. Finally, superstars are not subject to the transitory labor productivity shock.

[18] We exclude government expenditures on Social Security since they are explicitly included in our model.

and calibrate τ_2 to ensure the government budget constraint is satisfied.

Medical Expenditures and Transfers: Following DeNardi, French, and Jones (2010) and Kopecky and Koreshkova (2014), the consumption floor $\underline{c}$ is a stylized representation of the US Medicaid program and other means-tested public assistance programs. We set $\underline{c}$ to 15 percent of average labor earnings for all households until the normal retirement age (of 66 years old), which is the midpoint of estimates in Kopecky and Koreshkova (2014).[19] Finally, we construct the age-dependent medical expenditures, $\{\mu_j\}$, to match average household out-of-pocket medical expenditures in the HRS-AHEAD dataset while accounting for household composition and survivorship.

Social Security: Consistent with the minimum and maximum retirement ages in the U.S. Social Security system, we set the interval in which households can retire to the ages 62 and 70, and we set the normal retirement age to 66. The early retirement penalty and delayed retirement credits are set in accordance with the Social Security program. In particular, if households retire up to three years before the normal retirement age, then households' benefits are reduced by 6.7 percent for each year they retire early. If they choose to retire four or five years before the normal retirement age, then benefits are reduced by an additional 5 percent for these years. If households choose to delay retirement past normal retirement age, then their benefits are increased by 8 percent for each year they delay. The marginal replacement rates in the progressive Social Security payment schedule ($\tau_{r1}, \tau_{r2}, \tau_{r3}$) are also set at their actual respective values of 0.9, 0.32 and 0.15. Following Huggett and Parra (2010), the bend points where the marginal replacement rates change (b_1^{ss}, b_2^{ss}, b_3^{ss}) and the maximum earnings ($\bar{m}$) are set equal to the actual multiples of mean earnings used in the U.S. Social Security system so that b_1^{ss}, b_2^{ss} and $b_3^{ss} = \bar{m}$ occur at 0.21, 1.29 and 2.42 times average earnings in the economy. We set the payroll tax rate, τ_{ss} such that the program's budget is balanced. In our baseline model the payroll tax rate is 15.1 percent.

3.2 Baseline Economy

The model matches the set of targeted moments very closely, as shown in Table 2. Next we examine the fit of the model against a set of moments that we did not target in calibration. To do so, we construct empirical counterparts to objects in the model to be consistent with a household concept, using the Consumer Expenditure Survey for

[19] Using a different target (a 14% takeup rate in public transfers of this type), Kopecky and Koreshkova (2014) calibrate a very similar consumption floor of 16.5% of labor earnings. However, measuring the consumption floor is generally difficult, in part, because of the heterogeneity in programs available for different households.

Table 2: Target Moments, Model and Data

Moments	Data	Model
Capital-to-Output Ratio	2.900	2.903
Retired by age 66 (%)	0.700	0.700
Hours Worked up to age 66	0.269	0.269
Bequests-to-Wealth Ratio	0.009	0.010
Bequests, 90th pct / Avg. Labor Earnings	10.695	10.677
Top 40% Wealth Share	0.946	0.946
Top 20% Labor Income Share	0.635	0.635
Borrowers' Debt-to-Wealth	-0.048	-0.048

household consumption expenditures, the Panel Study of Income Dynamics (PSID) for household labor income and hours, and the Survey of Consumer Finances (SCF) for wealth.

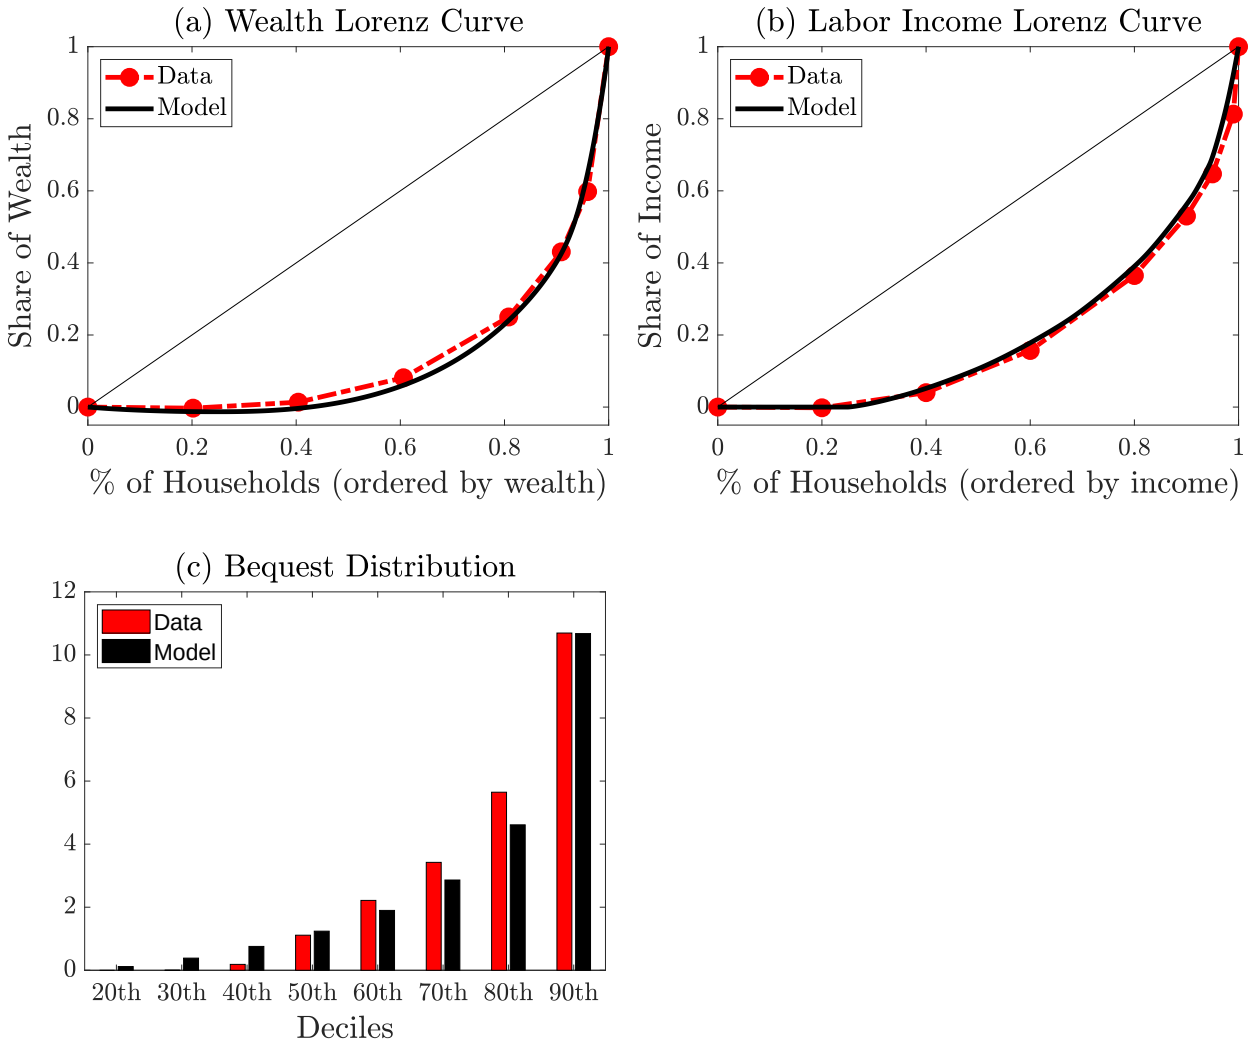
First, Figure 1 examines the aggregate distributions of wealth and labor, as well as the distribution of bequests from households over the age of 70. Although the model is calibrated to match a sparse set of distributional moments (the share of wealth above and below the 60th percentile, the share of labor income to the top 20%, the average bequest and the share of bequests left by households at the 90th percentile of the bequest distribution), it characterizes the whole distribution in each case. Panel (a) shows that the model's wealth Lorenz curve matches the data quite well over the entire wealth distribution;[20] panel (b) shows that the model closely matches the Lorenz curve for labor income; and panel (c) shows that the model matches the distribution of bequests quite well.[21]

Next, because the evolution of households' allocations along the life cycle will be a key determinant of the effects of fiscal policy, we check the model's fit against the life cycle profiles of average consumption, savings, labor income, and hours by age in Figure 2. In panel (a), we find that consumption in the model matches that in the data

[20] A well known feature of this class of models is a tendency to underpredict the amount of concentration at the very top of the wealth distribution. Because the model was calibrated to capture wealth inequality in the top 40% of the distribution and we did not construct the model to capture the extreme skewness in the top 1% of the wealth distribution, we exclude the top 1% from both the model and the data in the depicted comparison. We calculate wealth as the households' net worth from the 2016 Survey of Consumer Finances.

[21] We calculate bequests from the Health and Retirement Survey (HRS) which samples households over the age of 70. In order to focus on bequests that are not transferred within the household, we exclude estates in which there were assets transferred to the spouse.

Figure 1: Wealth and Bequest Distributions, Model and Data

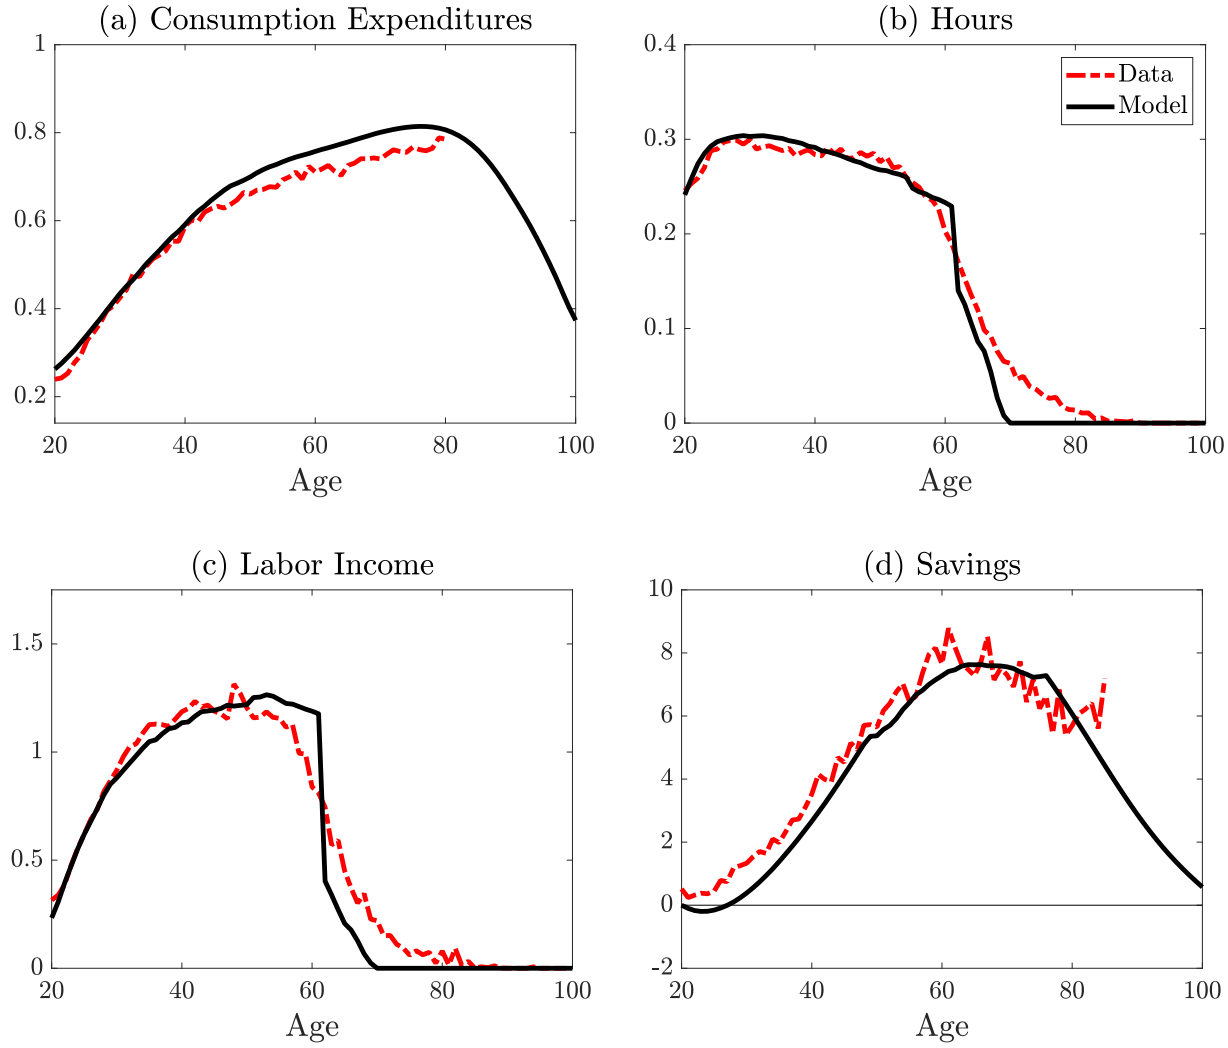


over the available range of ages, both in terms of contour and level.[22] Likewise, from panels (b) and (c), average household hours and labor income in the model are very close to those in the data at each age.[23] Lastly, panel (d) demonstrates that the model

[22] We follow Aguiar and Hurst (2013) in measuring the relevant empirical counterpart to the model's consumption expenditures as total household consumption expenditures including primary and owner equivalent rents, and excluding durable goods expenditures, work related expenses and education. We further exclude medical expenditures as these are explicitly modeled for post-retirement households. We measure consumption expenditures for head of households of ages 20-80, and estimate age profiles with cohort fixed effects, normalized year fixed effects, and controls for education and sex of the head of household. Consistent with Aguiar and Hurst (2013), we do not include later ages due to a small sample in the Consumer Expenditure Survey. We express consumption expenditures as a fraction of aggregate labor income in the model and in the PSID.

[23] We measure hours worked and labor income from the PSID. Household hours are calculated as the annual hours worked by the head of household and a spouse (if any), normalized by total annual

Figure 2: Life Cycle Profiles, Model and Data



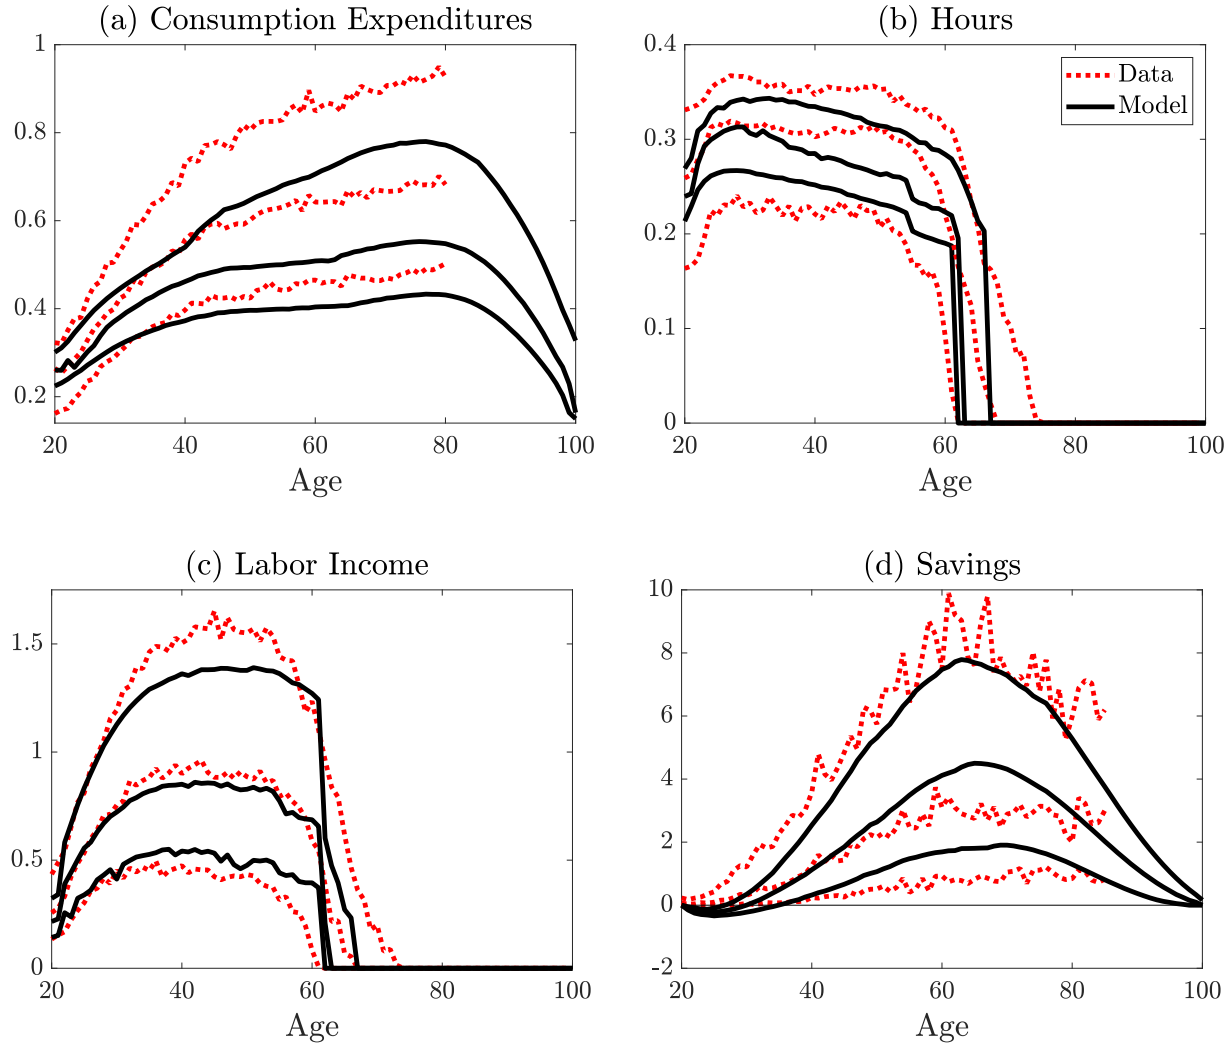
Notes: Solid lines depict the model's life cycle profile under the baseline public debt policy (67% of output). Dashed lines depict empirical life cycle profiles. Hours profiles are normalizd by total annual hours. All other model generated life cycle profiles are normalized by the model's aggregate labor income and empirical life cycle profiles are normalized by PSID aggregate labor income.

generates a very good match to the empirical life cycle profile for average savings.

Finally, in Figure 3 we assess how well the model characterizes the dispersion in consumption, savings, labor income and hours over the life cycle by comparing the 25th, 50th, and 75th percentiles at each age in the model and data. Although the model does not explicitly target the dispersion in household allocations by age, it

hours (40 hours per week times 52 weeks per year) and the number of earners within the household (1 or 2). To be consistent with the PSID which topcodes labor income reports, panel (c) compares the data to the bottom 99% of labor income observations in the model.

Figure 3: Distribution of Life Cycle Profiles, Model and Data



Notes: Red dashed lines depict data percentiles (25th, 50th, 75th) that have been normalized by PSID aggregate labor income. Black solid lines depict model percentiles (25th, 50th, 75th) that have been normalized by model generated aggregate labor income. The hours profile was normalized by total annual hours.

generates a reasonably close fit to its empirical counterparts, particularly for hours, labor income and savings. Although model underpredicts the levels of consumption at these percentiles, which implies a more skewed consumption distribution in the model, it still captures well the rise in consumption dispersion relative to the median over the life cycle.

Overall, the model describes the data well. We find that the model not only has a tight in-sample fit, but also matches a large number of untargeted moments that characterize life cycle averages, aggregate inequality and how inequality evolves over

Table 3: Interest Rate Elasticity With Respect to Public Debt-to-Output

	Yields		Forward Rates		New
Counterfactual	**10yr**	**30yr**	**5yr/10yr**	**10yr/10yr**	**Steady State**
Higher Gov. Spending	2.06	1.82	1.98	1.80	1.40
Lump-Sum Transfer	0.33	0.88	0.80	1.04	1.40
Tax Reduction	-0.03	0.49	-0.01	0.48	1.40
Persistent Tax Reduction	-0.41	-0.51	-1.19	-1.05	1.41

the life cycle.

3.3 Policy Counterfactuals

Given a model that is calibrated to recent U.S. data on households and macroeconomic aggregates, we now measure the response of interest rates at varying maturities (horizons) to changes in fiscal policies that add to the public debt. First we consider a series of unanticipated one-time shocks, and then we turn to an illustrative example of a more persistent shock – a ten year tax reduction that is calibrated to mimic annual effect of the 2018 Tax Cut and Jobs Act (TCJA) – in order to quantify mechanisms corresponding to more plausible magnitudes and duration of tax cuts.

One-time Public Debt Shocks. First, we consider the implications of (i) a one time increase in government spending, G, (ii) a one-period decrease in the linear component of total income taxes, τ_0 of the tax function $\mathrm{Y}(y_j(h,a,e,d))$, and (iii) a one-period increase in lump-sum transfers that are equally distributed across households, Tr. All three experiments assume fiscal shocks that generate a 3.3 percent increase in public debt relative to the initial level of output, so that each economy converges to the same steady state with 3.3 percent higher debt.

In the new steady state equilibrium, the interest rate is 4.6 basis points higher which implies a long-term elasticity of interest rates with respect to government debt of 1.4 basis points, which is within the range of our empirical estimates. However, while past research argues that the elasticity estimated using the 5-year forward, 10-year interest rate is a good approximation to the long-term elasticity, we find that there is a large amount of variation in the elasticity of this rate over transition to the new steady state. Panel (d) of Figure 4 shows significant variation in the one period spot interest rate transition paths across the policy experiments, while Table 3 reports the large dispersion in the model's debt elasticities on the 5-year forward 10-year interest

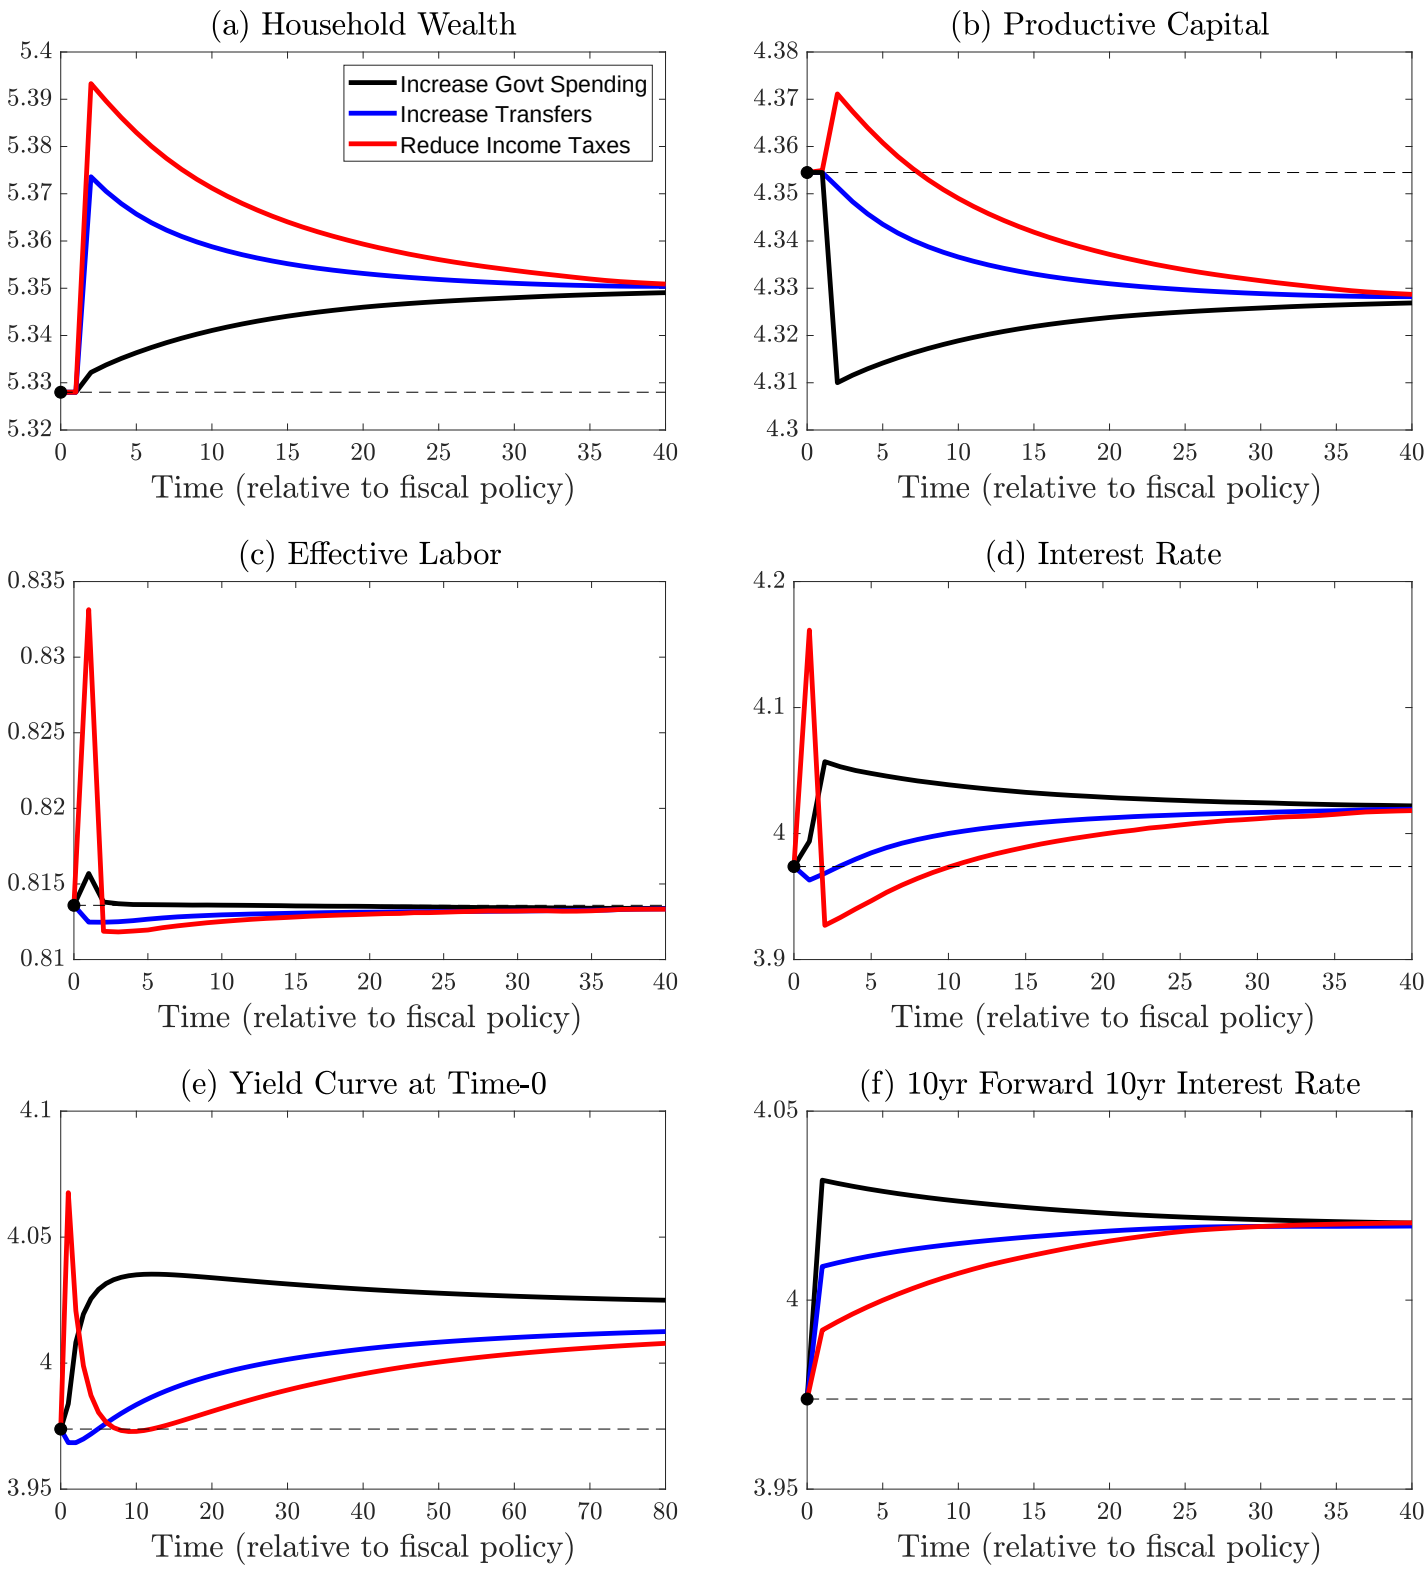


Figure 4: **Transition Paths for One-Time Debt Increases**

rate at the time of the shock: 2 basis points when the debt is generated from government purchases, 0.8 basis points when it is generated from a lump-sum transfer, and -0.1 basis points when it is generated from a tax cut. Moreover, significant dispersion in elasticities persists at longer horizons for interest rates (e.g., for the 10-year forward 10-year rate along with the 10 and 30 year yields).

The variation in the elasticity with respect to the 5-year forward 10-year rates is due to the distinct transitional dynamics across the three experiments. For example, public debt-financed government expenditures do not directly transfer resources to households, and so household savings barely responds at the time of the fiscal policy shock and there is essentially a one-to-one crowd out of productive capital that induces a higher interest rate (see the black lines in panels (a), (b) and (d) of Figure 4). Over time, households respond to the higher interest rate by accumulating more savings which puts downward pressure on the interest rate. All told, the total increase in household savings is less than the increase in government debt, and thus crowd-out of productive capital persists as the interest rate settles at a higher level than in the initial steady state. Because the amount of productivity capital crowd-out between years 10 to 20 is higher than in the new steady state, the estimated elasticity on the 5-year forward 10-year rate is 2 basis points, higher than the steady state elasticity of 1.4 basis points. Similarly, the the 10-year forward 10-year rate, 10 and 30-year yield elasticities are larger than the elasticity on the steady state rate.

In contrast, when the higher debt finances a one period cut in taxes initially there is a *crowd in* of productive capital due to two forces (see the large increase in the red line of panel (a) of Figure 4). First, households directly receive all of the resources that generate the debt, and since the tax cut is only temporary households initially save these additional resources to fund consumption in future periods. Moreover, since households experience a one-period income tax reduction, they also substantially increase their effective labor supply (see panel (c) of Figure 4) and save this additional income for future consumption. In total, these two effects lead the productive capital stock to be larger than that in the initial steady state, thus pushing down the equilibrium interest rate. Since it takes many periods for households to spend their elevated level of savings, the interest rate 10 to 20 years after the tax cut is lower than its initial steady state value, leading to a *negative* elasticity on the 5-year forward, 5-year interest rate and 10-year yield of essentially zero and a smaller elasticity on the 10-year forward, 10-year rate and 30-year yield relative to the steady state elasticity (see Table 3). When the public debt finances a one-time transfer, the transition dynamics are somewhat similar to those with a tax cut, but without the strong labor response, thus yielding a higher elasticity compare to the tax cut but smaller than the elasticity with

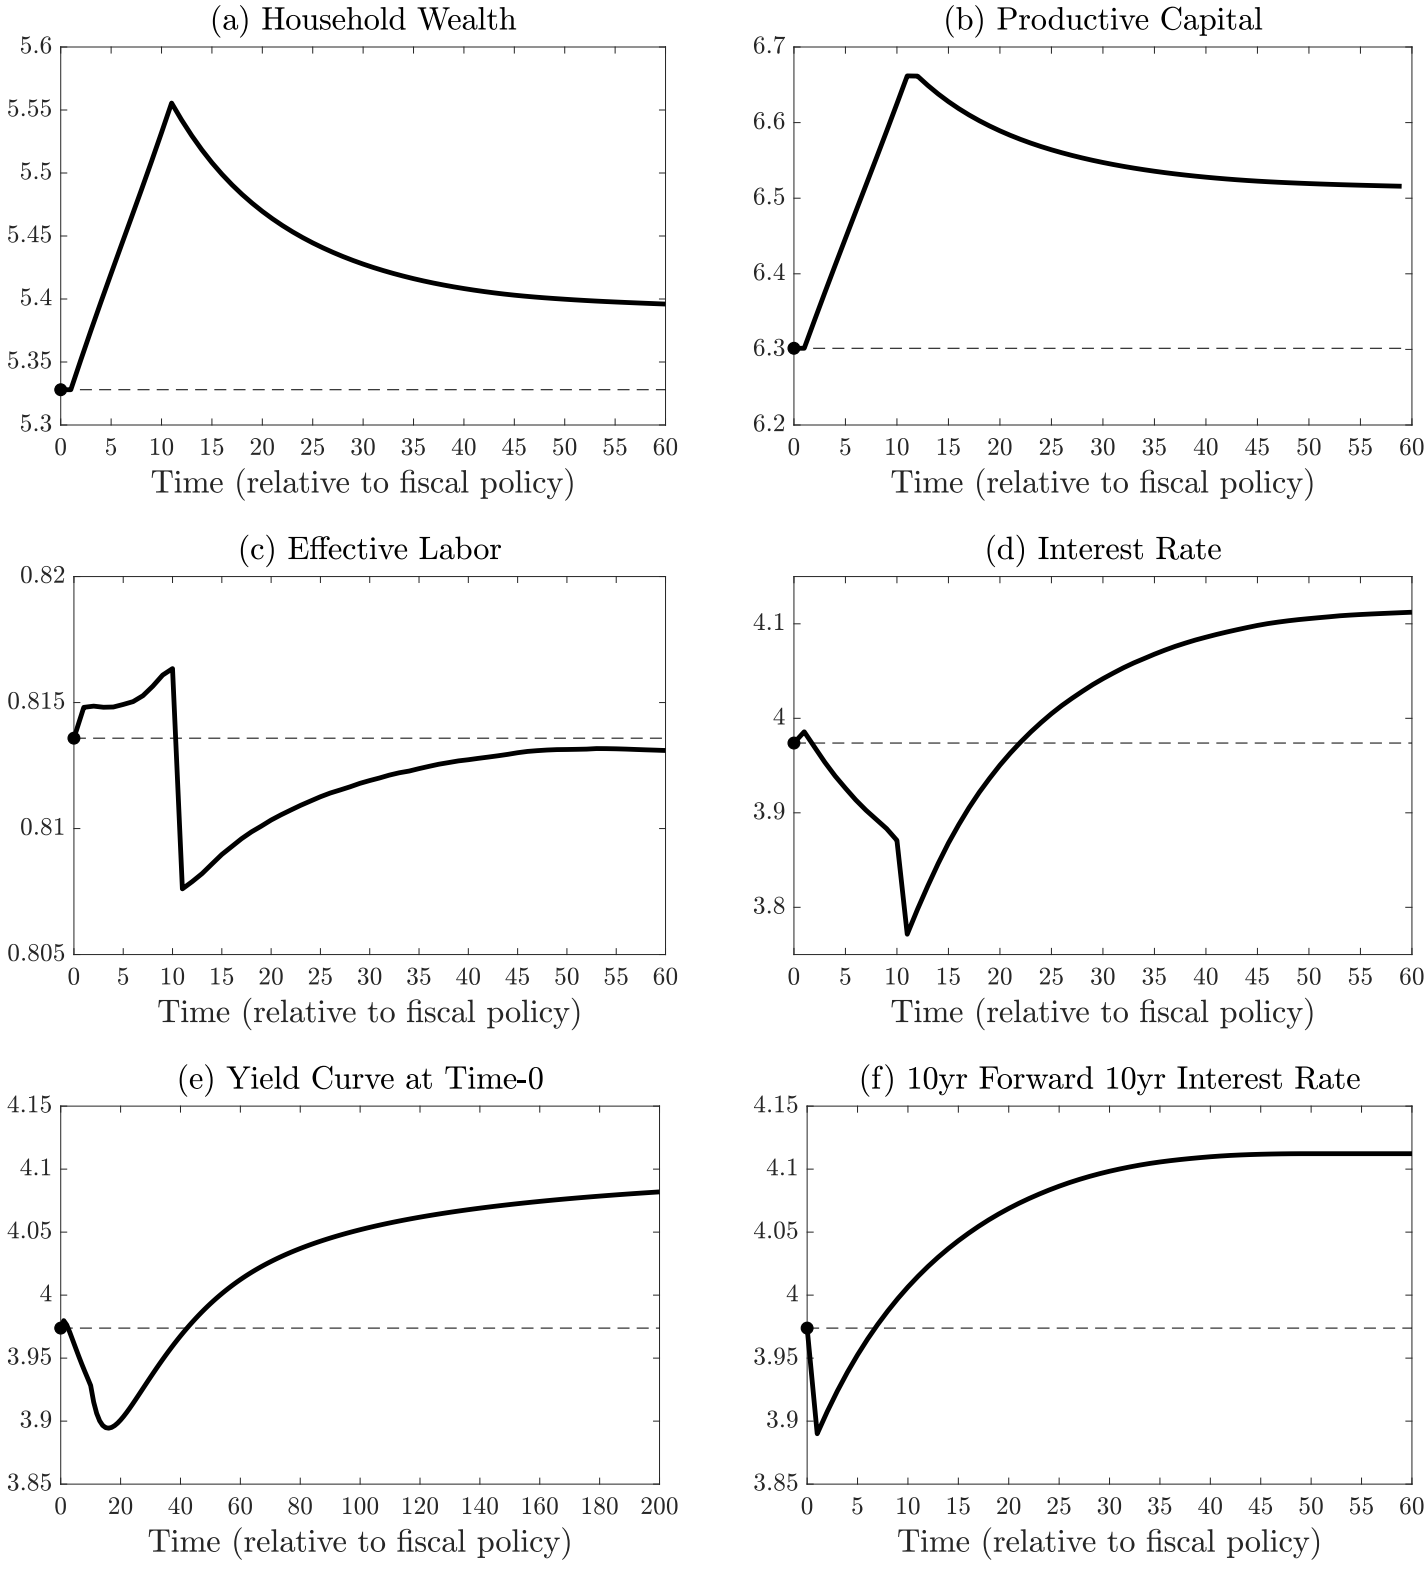


Figure 5: **Transition Paths for Persistent Tax Reduction**

the government spending shock.

Taken as a whole, the dispersion in transitional responses to fiscal shocks indicate that the "interest rate elasticity with respect to public debt" is not a well-defined object. The long transitional paths after each type of fiscal shock makes standard steady state comparisons misleading for policy analysis within standard planning periods (e.g. 10 to 20 years). Importantly, fiscal policies that transfer resources to households lead to quite different savings and interest rate dynamics than fiscal policies that do not direct resources directly to the households which can lead to the opposite effect on yields at the time of the shock. Thus, these results indicate that it is incorrect to consider one interest rate elasticity with respect to homogeneous debt concept over a 10-to-20 year horizon, but instead one must consider the source of the debt.

Persistent Tax Cut Experiment. Next we consider a 10-year income tax cut equal to 1 percent of GDP per year, after which time taxes revert to their original rate. This example is motivated by the reduction in taxes from the Tax Cut and Jobs Act (TCJA) which permanently reduced the corporate tax rate and temporarily reduced personal income taxes.[24] In order to capture the differential effects of the tax cut across the income distribution we match the revenue reductions as a share of output across the income distribution to estimates in Gale et al. (2019). Although most of the provisions that affected personal taxes were only enacted for seven years, we allow the tax cut to persist for ten years, for illustrative purposes, after which taxes revert back and the economy transitions to the new steady state with a debt-to-GDP ratio that is 10 percentage points higher.

In Figure 5, we observe that the interest rate falls over the ten years of tax cuts and productive capital increases as households work more hours and accumulate savings at a rapid pace in anticipation of future taxes returning to their higher rate. After the tax cut ends, households hold a higher level of savings that they spend down in order to afford a higher level of lifetime consumption. Accordingly, the interest rises thereafter in lockstep with household savings drawing down and with declining resources for productive capital.

Relative to the previous one-time tax reduction counterfactual, this experiments measures the impact over a policy that persists for a longer period. In particular, households accumulate notably more savings over the ten year window of the tax cut along with increasing their labor. Thus, the spot interest rate falls at the onset of the tax change and continues to decline during the 10 years in which the tax rate is lower.

[24] Although there have been numerous estimates of the size of the TCJA which have evolved over time, generally the range of estimates by the Joint Committee on Taxation and Congressional Budget Office have been centered around 1 percentage point of GDP.

After the tax cut expires the spot interest rate begins to rise slowly, but continues to be below the initial steady state value for 20 years. Consequently, we find that the elasticity on the forward rates and yields are *negative*.

4 Estimating the Interest Rate Elasticity of Debt

Next, we relate the model results to empirical estimates of the response of the real interest rate to a change in public debt. Conceptually there are several threats to identification from naively regressing interest rates on measures of public debt, most importantly that the short-term the interest rate varies with macroeconomic conditions and shocks that are unrelated to variation in public debt. In order to purge our estimates of confounding shocks, we follow Laubach (2009) and Engen and Hubbard (2004) in using forecasted values of debt and medium-term forward interest rates, as well as including a large set of controls for other macroeconomic shocks that simultaneously affect debt and interest rates. Importantly, we update the standard approach by identifying elasticities from *forecast revisions* instead of forecasted levels of debt. One advantage of this strategy is that the evolution of government debt is plausibility non-stationary (c.f. Barro (1979), Bohn (1998), and Campbell et al. (2023)), while forecast revisions are stationary and thus estimates exhibit greater coefficient stability. Second, constructing forecast revisions allows us to isolate different spending sources that explain the change of public debt and test whether elasticity estimates are sensitive to these sources.

4.1 Methodology, Data and Estimates

Framework and Previous Estimation. The estimation framework proceeds from a standard government budget constraint. The government spends G_t and receives T_t in taxes from households each period, to form the primary surplus of $T_t - G_t$. The government chooses its level of new borrowing, D_t, and pays interest on existing debt at the prevailing interest rate, $r_t D_{t-1}$. We define debt as negative $D_t < 0$, but the government can also save by choosing $-D_t > 0$. Thus the government's budget constraint is,

$$D_t - D_{t-1} = (T_t - G_t) + r_t D_{t-1} \ .$$

Accordingly, j-period ahead expectations of public debt are obtained by cumulating the government budget constraint,

$$\mathbb{E}_t[D_{t+j}] = D_{t-1} + \mathbb{E}_t\Big(\sum_{j=0}^{J}(T_{t+j} - G_{t+j})\Big) + \mathbb{E}_t\Big(\sum_{j=0}^{J} r_{t+j} D_{t-1+j}\Big) \ .$$

Laubach (2009) and Engen and Hubbard (2004) estimate the impact of the *level of* public debt relative to GDP, denoted Y_t, on the j-period ahead interest rate with the following specification,

$$r_{t+j} = \alpha + \beta \mathbb{E}_t[D_{t+j}/Y_{t+j}] + \gamma X_t + \varepsilon_t \tag{9}$$

where X_t is a set of time-varying controls for macroeconomic conditions that might confound the effect of debt-to-GDP on interest rates.

For this specification, the interest rate data we use in estimation consists of updated yield curve estimates from Gürkaynak et al. (2007) to derive the 5-year ahead forward interest rate on a 10-year treasury bond. For the debt measure, we either use the CBO's 5-year ahead forecast for total public debt-to-GDP. Additionally, we control for expected total PCE inflation to effectively estimate the effect on the real interest, projected GDP growth to adjust the interest rate for growth trends as is predicted by a standard Euler equation, and the dividend yield to control for the effect of equity returns on the interest rate. We further control for the stock of US Treasuries held by the Federal Reserve and the stock of US Treasuries held by foreign governments to control for composition of government debt.[25] Our choice of control variables closely follows Laubach (2009) and Gamber and Seliski (2019), and thus our results will be comparable.

Column (1) of Table 4 shows that a 1 percent increase in total debt-to-GDP leads to an increase in the 5-year ahead, 10-year interest rate of 2.4 basis points. This estimate falls within the range of those found in Gamber and Seliski (2019), which uses a similar sample period of 1987 to 2017 as well as the same controls variables. However, column (2) shows that when we extend the sample endpoint from 2017 to 2020 the estimates become negative and small, which we find is robust to several specifications. The sensitivity of the estimate to the sample period is plausibly due to a unit root in the time series for both public debt and interest rates. Both panels of Figure 6 show that interest rates and debt have a secular trend with multiple regimes.

[25] Including the composition of debt across domestic and foreign sources helps alleviate concerns that the elasticity is driven by a global savings glut during this time period, as discussed in Bernanke (2005).

Table 4: Coefficient Stability Across Specifications

	(1)	(2)	(3)	(4)
5yr Total Debt	0.024^{***}	-0.005	0.050^{***}	0.048^{***}
	(0.007)	(0.009)	(0.019)	(0.017)
Dependent Variable	$r_t^{5,10}$	$r_t^{5,10}$	$\Delta r_t^{5,10}$	$\Delta r_t^{5,10}$
Controls	×	×	×	×
Output Gap	t+1-t+5	t+1-t+5	t+1-t+5	t+1-t+5
GDP Growth	t+5	t+5	t+5	t+5
Aging				
Sample period	1987-2017	1987-2020	1987-2017	1987-2020
N	89	95	89	95

*Notes: The dependent variable is the 5-year ahead 10-year forward interest rate. Standard errors in parentheses. * denotes significance at the 10% level, ** at the 5% level, and *** at the 1% level.*

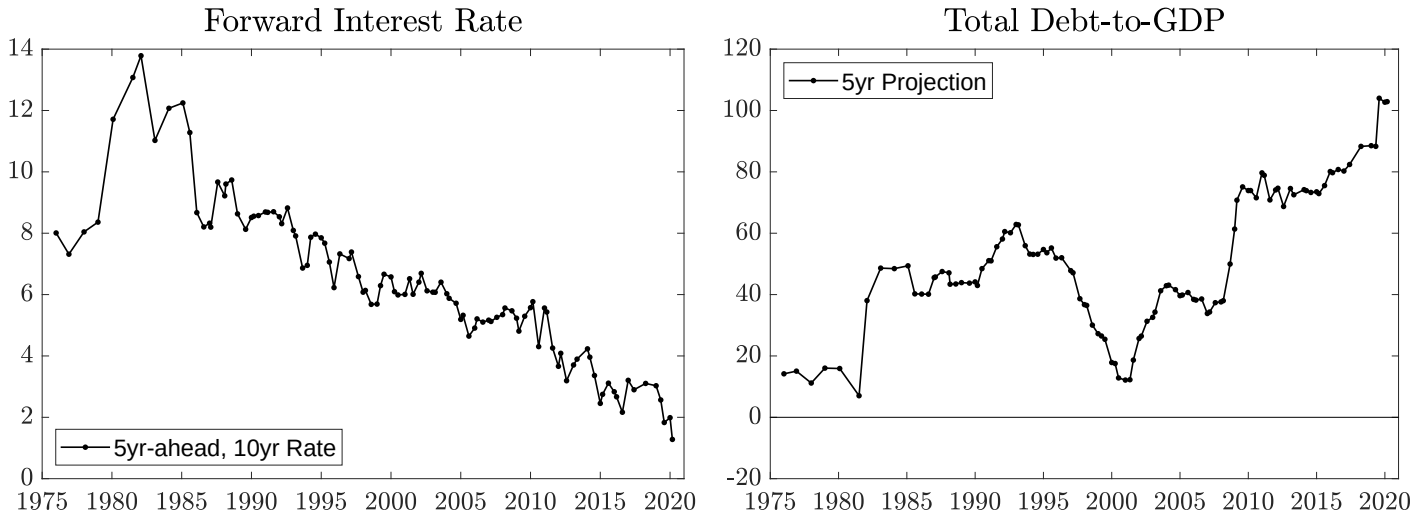


Figure 6: **Interest Rates and Debt**

Figure 7 plots interest rates against expected public debt-to-GDP from 1987-2020 and breaks out three separate regimes: a positively relationship between 1987-1999, an essentially flat relationship between 2000-2009, and a negative relationship between 2010-2020. Importantly, we find that estimates fall considerably when the sample period is extended to include additional post-2010 data, which features a negative trend in interest rates and an upward trend in debt.

Baseline Specification. Our strategy for simultaneously introducing parameter stability over sample periods and purging threats to estimation contained in short-

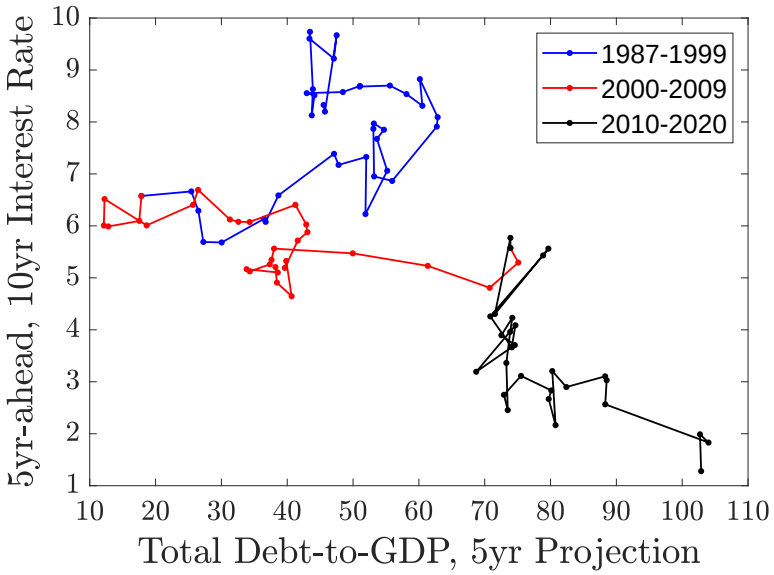


Figure 7: **Interest Rates and Debt Regimes**

term macroeconomic shocks is to consider a difference-in-difference style estimator that leverages information contained in successive forecasts. In particular, the j-period ahead forecast error is the difference between the expected debt and the actual realization of debt,

$$\eta_{t,t+j} = \mathbb{E}_t[D_{t+j}/Y_{t+j}] - D_{t+j}/Y_{t+j}$$

and the forecast revision is therefore the difference in forecast errors for date $t+j$ as of date $t-1$ and as of date t,

$$\Delta\eta_{t,t+j} \equiv \eta_{t,t+j} - \eta_{t-1,t+j} = \mathbb{E}_t[D_{t+j}/Y_{t+j}] - \mathbb{E}_{t-1}[D_{t+j}/Y_{t+j}] \ .$$

The primary distinction between the forecast at time $t-1$ and time t is that the date t forecast contains additional information, and can thus be thought of as being the result of a treatment of additional information relative to the $t-1$ forecast. Accordingly, we estimate the impact of the unexpected component of debt on the change in the j-period ahead interest rate, which gives the response of interest rates to news about debt, with the following specification,

$$\Delta r_{t+j} = \alpha + \beta\Delta\eta_{t,t+j} + \gamma\Delta X_t + \varepsilon_t \tag{10}$$

where ΔX_t is the first difference of the time-varying controls.

In addition to the control variables used in the levels specification in equation (9), we also control for any residual impact of the business cycle with the expected out-

put gap from the CBO, to help absorb the effect of persistent business cycle shocks on public debt forecasts. Moreover, we control for supply side determinants of trend interest rates with a measure of population aging, specifically the ratio of the population over the age of 65 relative to that between ages 20 and 64.[26] Finally, when the control variable is a forecasted value from the CBO (e.g., GDP growth and the output gap), we use forecast revisions as controls in this specification.

Before proceeding to introduce further aspects of the identification strategy, we assess the stability of the estimates derived from equation (10). Columns (3) and (4) of Table 4 show that the updated specification yields estimates that are not sensitive to sample period, unlike the estimates in the specification in columns (1) and (2). Moreover, trends in the level data apparently attenuated the magnitude of the coefficient in half.

Endogeneity from Net Interest Payments. Next, given a strategy for identifying the effect of public debt on interest rates from the news contained in forecast revisions, we remain concerned that the way government debt is measured can threaten identification. In particular, one component of government debt is the net interest payment on existing debt, which contains interest rates from previous debt obligations. Thus, the component of net interest payments introduce a concern over reverse causality since those payments are a function of interest rates, and those interest rates may furthermore reflect a host of other shocks that are unrelated to the formation of the primary debt.

We address this concern by decomposing revisions to the total debt forecast into forecast revisions over the primary deficit and net interest payments. We do so by taking CBO deficit forecasts for the primary debt and net interest payments, which can then aggregate to debt forecast revisions over the projection horizon. We can subsequently study the primary deficit absent the plausibly endogenous component of net interest payments. This is a notable advantage of estimating in differences, as we do not need the *level* of debt due to the primary deficit or net interest payments but only the change in those components. This is precisely why prior work has not been able to separate the primary deficit from the net interest rate, empirical specifications in levels provide no way of separating debt projections into components.

In columns (1) and (2) of Table 5, we demonstrate that the potential bias from including net interest payments. We find that the estimate in column (1), which identifies the elasticity from total debt, is larger than the estimate in column (2) which excludes net interest payments. Thus we find that excluding interest payments from

[26] Controlling for population aging is consistent with work such as Carvalho et al. (2016) and Acemoglu and Restrepo (2017).

Table 5: Specifications in First Differences

	(1)	(2)	(3)	(4)
5yr Total Debt	0.048***			
	(0.017)			
5yr Primary Debt		0.041**	0.042*	
		(0.021)	(0.022)	
10yr Primary Debt				0.011
				(0.012)
Dependent Variable	$\Delta r_t^{5,10}$	$\Delta r_t^{5,10}$	$\Delta r_t^{5,10}$	$\Delta r_t^{10,10}$
Controls	×	×	×	×
Output Gaps	t+1-t+5	t+1-t+5	t+1-t+5	t+1-t+10
GDP Growth	t+5	t+5	t+5	t+10
Aging	×	×	×	×
Sample period	1987-2020	1987-2020	1996-2020	1996-2020
N	95	95	68	68

*Notes: The dependent variable is the 5-year ahead 10-year forward interest rate. Standard errors in parentheses. * denotes significance at the 10% level, ** at the 5% level, and *** at the 1% level.*

the debt measure leads to somewhat smaller estimates, by around 7 basis points.

Forecast Horizon and Debt Maturity. The prior literature estimated effect of 5-year forward, 10-year interest rates on 5-year ahead public debt forecasts, and treated the 5-year horizon as a sufficiently close proxy to the long-run in which contemporaneous shocks could no longer contaminate the measures of interest rates and debt. In part, the choice of a 5-year ahead horizon was due to a data limitation, as the CBO only started forecasting at longer horizons in 1996.[27] We now turn to estimating our specification at a 10-year horizon, the longest horizon contained in post-1996 CBO projections, in order to obtain the elasticity of the 10-year ahead, forward interest rate on a 10-year bond with respect to the 10-year ahead forecast revision on the primary debt.[28]

Columns (3) and (4) of Table 5 estimate the interest rate effect of higher public debt over the full ten year forecast horizon. Column (3) provides the specification with 5-year primary debt projections from 1996-2020 as a benchmark for comparison. In

[27] As an example, Laubach (2009) employs the sample period of 1976-2003, which did not include many observations with 10-year ahead forecasts.

[28] We also include 10-year horizons for the relevant control variables, such as the changes in the GDP growth rate and output gap.

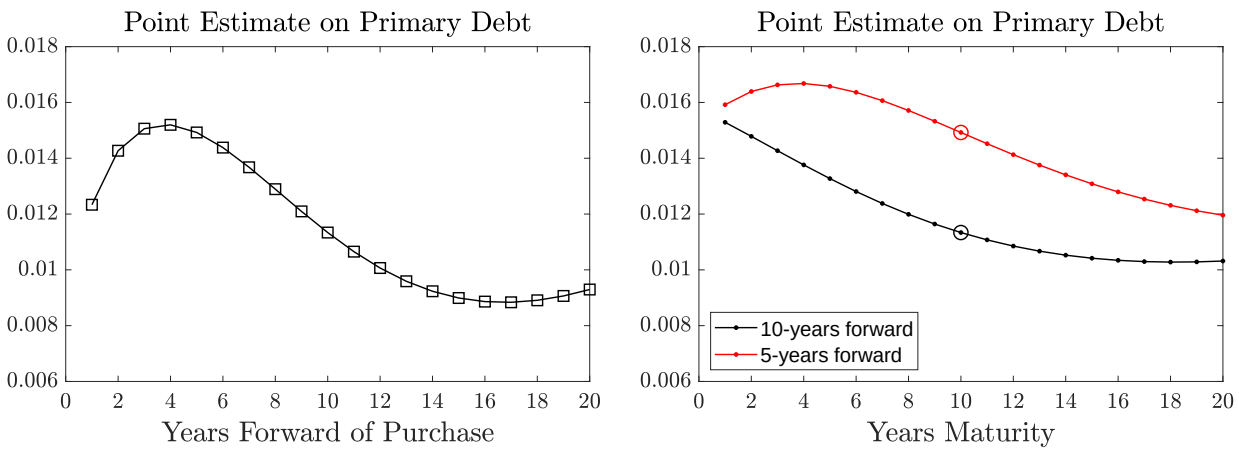


Figure 8: **Local Projection of Interest Rates on 10-year Debt Projections**

column (4), when all ten years of debt information are included in the specification, we find a point estimate of 1.1 basis points.[29] It should not be surprising that the estimate is smaller at a longer horizon for public debt projections, since the magnitude of the change in forecasted debt over a longer horizon increases (e.g., since changes in annual deficits are persistent) but the dependent variable of interest rate changes has far less variation as we increase the number of years forward at which the bond is purchased. Mechanically, the point estimate falls as an independent variable with greater variance explains a dependent variable that has little changed.

Next, we show that the estimated elasticity varies with the particular forward rates that are used for identification. To do this, we consider a set of local projections of the interest rate with respect to the number of years forward that the bond will be purchased and the number of years the bond will be held. The left panel of Figure 8 shows that the elasticity with respect to 10-year the public debt projection declines as we increase the interest rate's purchase timing from 5-years forward to 10-years forward, while the left panel of Figure 9 shows that the elasticity with respect to the 5-year public debt projection is larger in magnitude and declines more rapidly. Likewise, the right panels vary the maturity of the bond and give a similar result.

Overall, the variation in the empirical estimates along both the years forward and the years to maturity are consistent with the results of our quantitative model experiments – that the effect of increased public debt on the spot interest rate is very persistent over the transition.

Elasticity Variation by Spending Source. Next, we show that the elasticity varies with the type of fiscal policy that generates the debt, which is consistent with the

[29] In robustness exercises with respect to specification, we find the coefficient that ranges between 1 and 1.4 basis points.

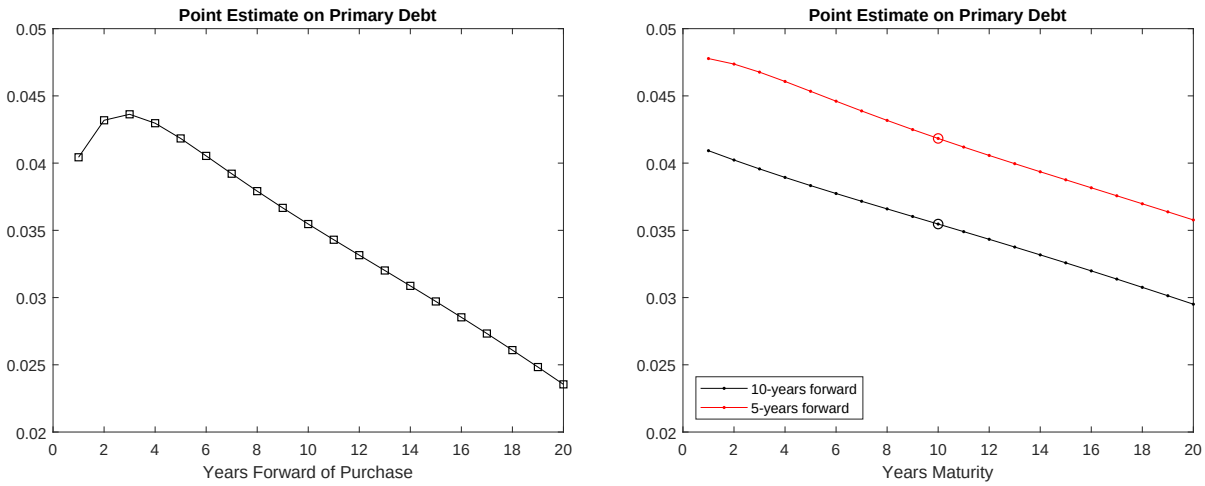


Figure 9: **Local Projection of Interest Rates on 5-year Debt Projections**

results from our quantitative model. In particular, we use the CBO's forecast decompositions to construct forecast revisions for components of the primary deficit: government revenues (mainly inflows from taxation), mandatory government outlays (such as entitlement programs), and discretionary government outlays (mainly due to changes in annual budgeting as well as emergency legislation). We reestimate the empirical specification by replacing the 10-year forecast revision in the primary deficit with the forecast revisions in all three of its components. Table 6 shows that there is considerable heterogeneity in the elasticities with respect to (negative of) revenues, discretionary outlays and mandatory outlays.

In column (2) we find that the elasticity with respect to (the negative of) revenues is similar to, though somewhat larger than, the elasticity for the primary deficit as a whole (in column (1)). Moreover, the elasticity with respect to discretionary outlays is close to 1 basis point. The elasticity with respect to mandatory outlays is negative instead of positive, raising a concern that mandatory outlays could be correlated with the business cycle despite our experimental design to purge such correlation (e.g., since they include automatic stabilizer programs such as unemployment insurance or Medicaid). In column (3) we interact each component with a dummy variable that equals one when the 10-year projected output is below the 10-year projection for potential output, and only find a significant additional effect on the elasticity with respect to mandatory spending when the output gap is negative. Notably the additional effect on mandatory spending is positive, suggesting that households consume the increases in mandatory spending during downturns but save them otherwise, or that the forward rates are correlated with the cycle. Finally, we observe that the estimates are somewhat imprecise and not consistently statistically different than zero, in

Table 6: Specifications with Debt Decomposition

	(1)	(2)	(3)
10yr Primary Debt	0.011		
	(0.011)		
10yr Revenues (neg)		0.021*	0.020*
		(0.011)	(0.012)
10yr Discretionary Outlays		0.010	0.039
		(0.025)	(0.052)
10yr Mandatory Outlays		−0.103**	−0.158***
		(0.050)	(0.050)
10yr Revenues $\times$ $\mathbf{1}[gap < 0]$			−0.017
			(0.025)
10yr Discretionary $\times$ $\mathbf{1}[gap < 0]$			−0.053
			(0.064)
10yr Mandatory $\times$ $\mathbf{1}[gap < 0]$			0.118*
			(0.070)
Dependent Variable	$\Delta r_t^{10,10}$	$\Delta r_t^{10,10}$	$\Delta r_t^{10,10}$
Controls	$\times$	$\times$	$\times$
Output Gaps	t+1-t+10	t+1-t+10	t+1-t+10
GDP Growth	t+5	t+5	t+5
Aging	$\times$	$\times$	$\times$
Sample period	1996-2020	1996-2020	1996-2020
N	68	68	68

*Notes: The dependent variable is the 5-year ahead 10-year forward interest rate. Standard errors in parentheses. * denotes significance at the 10% level, ** at the 5% level, and *** at the 1% level.*

part given the relatively short time series. Overall, the considerable range of interest rate elasticities with respect to the subsequent increase in public debt are consistent with the quantitative model experiments which found that the elasticity is policy dependent.

5 Conclusion

The U.S. federal debt-to-GDP ratio has almost doubled since the onset of the Great Recession, highlighting the importance of the relationship between public debt and long-term interest rates. In a quantitative heterogeneous agent life cycle model, we

demonstrate that it is insufficient to characterize the relationship between debt and interest rates with one elasticity. In particular, we find that the interest rate responses vary substantially with the horizon of interest rate as well as type of fiscal policy associated with the generating the debt. The type of policy not only affects the magnitude of the interest rate response but can also induce a decrease in interest rates (instead of an increase) when the policy leads to a boost to household savings.

Empirically, previous research finds that a 1 percentage point increase in public debt-to-GDP leads to a 2.5 to 5 basis point increase in the interest rate. This paper revisited these estimates of the average effect of public debt on the interest rate and found a much smaller estimate after introducing a number of methodological improvements. While these estimates characterize the average effect from the historical changes in government debt, empirically we find evidence that the relationship may vary with both the type of policy that is financed by government debt and the exact interest rate measure used to estimate the elasticity further demonstrating that the relationship between debt and interest rates cannot be characterize with one elasticity.

References

Acemoglu, D. and P. Restrepo (2017): "Secular Stagnation? The Effect of Aging on Economic Growth in the Age of Automation," *American Economic Review*, 107, 174–79.

Aguiar, M. and E. Hurst (2013): "Deconstructing Life Cycle Expenditure," *Journal of Political Economy*, 121, 437–492.

Aiyagari, S. R. and E. R. McGrattan (1998): "The optimum quantity of debt," *Journal of Monetary Economics*, 42, 447–469.

Aladangady, A., K. Bi, and W. Peterman (2020): "Decomposing the Persistence of Household and Worker Earnings Risk," Unpublished.

Barro, R. J. (1979): "On the determination of the public debt," *Journal of Political Economy*, 87, 940–971.

Bauer, M. D. and G. D. Rudebusch (2016): "Why are long-term interest rates so low?" *FRBSF Economic Letter*, 36, 1–5.

Bell, F. and M. Miller (2002): "Life Tables for the United States Social Security Area 1900 - 2100," Office of the Chief Actuary, Social Security Administration, Actuarial Study 116.

Bernanke, B. S. (2005): "The Global Saving Glut and the US Current Account Deficit," *Lecture presented at Homer Jones Lecture, St. Louis, MO.*

Bohn, H. (1998): "The behavior of US public debt and deficits," *the Quarterly Journal of Economics*, 113, 949–963.

Campbell, J. Y., C. Gao, and I. Martin (2023): "Debt and Deficits: Fiscal Analysis with Stationary Ratios," Tech. rep., National Bureau of Economic Research.

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

Castañeda, A., J. Díaz-Giménez, and J.-V. Ríos-Rull (2003): "Accounting for the U.S. Earnings and Wealth Inequality," *Journal of Political Economy*, 111, 818–857.

Chatterjee, S., J. Gibson, and F. Rioja (2017): "Optimal public debt redux," *Journal of Economic Dynamics and Control*, 83, 162–174.

Conesa, J. C., S. Kitao, and D. Krueger (2009): "Taxing Capital? Not a Bad Idea after All!" *American Economic Review*, 99, 25–48.

Del Negro, M., D. Giannone, M. P. Giannoni, and A. Tambalotti (2017): "Safety, Liquidity, and the Natural Rate of Interest," *Brookings Papers on Economic Activity*, 48, 235–316.

DeNardi, M. (2004): "Wealth Inequality and Intergenerational Links," *Review of Economic Studies*, 71, 743–768.

DeNardi, M., E. French, and J. B. Jones (2010): "Why Do the Elderly Save? The Role of Medical Expenses," *Journal of Political Economy*, 118, 39–75.

——— (2016): "Savings after retirement: A survey," *Annual Review of Economics*, 8, 177–204.

Engen, E. M. and R. G. Hubbard (2004): "Federal government debt and interest rates," *NBER macroeconomics annual*, 19, 83–138.

Fagereng, A., M. B. Holm, B. Moll, and G. Natvik (2019): "Saving behavior across the wealth distribution: The importance of capital gains," NBER Working Papers 26588, National Bureau of Economic Research.

Feiveson, L. and J. Sabelhaus (2018): "How Does Intergenerational Wealth Transmission Affect Wealth Concentration?" FEDS Notes, Board of Governors of the Federal Reserve.

——— (2019): "Lifecycle Patterns of Saving and Wealth Accumulation," Finance and Economics Discussion Series 2019-010, Board of Governors of the Federal Reserve.

Gale, W. G., H. Gelfond, A. Krupkin, M. J. Mazur, and E. J. Toder (2019): "Effects of the tax cuts and jobs act: A preliminary analysis," *National Tax Journal*, 71, 589–612.

Gale, W. G. and P. R. Orszag (2004): "Budget deficits, national saving, and interest rates," *Brookings Papers on Economic Activity*, 2004, 101–210.

Gale, W. G. and J. K. Scholz (1994): "Intergenerational Transfers and the Accumulation of Wealth," *Journal of Economic Perspectives*, 8, 145–160.

Gamber, E. and J. Seliski (2019): "The effect of government debt on interest rates," Congressional Budget Office Working Papers 2019-01, Congressional Budget Office.

Gomme, P., B. Ravikumar, and P. Rupert (2011): "The Return to Capital and the Business Cycle," *Review of Economic Dynamics*, 14, 262–78.

Gouveia, M. and R. Strauss (1994): "Effective Federal Individual Income Tax Functions: An Exploratory Empirical Analysis," *National Tax Journal*, 47, 317–339.

Gürkaynak, R. S., B. Sack, and J. H. Wright (2007): "The US Treasury yield curve: 1961 to the present," *Journal of monetary Economics*, 54, 2291–2304.

Hendricks, L. (2007): "How important is discount rate heterogeneity for wealth inequality?" *Journal of Economic Dynamics and Control*, 31, 3042–3068.

Holston, K., T. Laubach, and J. C. Williams (2017): "Measuring the Natural Rate of Interest: International Trends and Determinants," *Journal of International Economics*, 108, 59–75.

Huggett, M. and J. C. Parra (2010): "How Well Does the U.S. Social Insurance System Provide Social Insurance?" *Journal of Political Economy*, 118, pp. 76–112.

Hur, S. (2018): "The Lost Generation of the Great Recession," *Review of Economic Dynamics*, 30, 179–202.

Johannsen, B. K. and E. Mertens (2018): "A Time Series Model of Interest Rates with the Effective Lower Bound," BIS Working Papers 715, Bank for International Settlements.

Kaplan, G. (2012): "Inequality and the life cycle," *Quantitative Economics*, 3, 471–525.

Kiley, M. T. (2020): "What Can the Data Tell Us about the Equilibrium Real Interest Rate?" *International Journal of Central Banking*, 16, 181–209.

Kindermann, F. and D. Krueger (2020): "High Marginal Tax Rates on the Top 1%? Lessons from a Life Cycle Model with Idiosyncratic Income Risk," *American Economic Journal: Macroeconomics*, Forthcoming.

Kitao, S. (2014): "Sustainable Social Security: Four Options," *Review of Economic Dynamics*, 17, 756–779.

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

Krueger, D., K. Mitman, and F. Perri (2016): *Macroeconomics and Household Heterogeneity*, Elevsier, vol. 2A of *Handbook of Macroeconomics*.

Laubach, T. (2009): "New evidence on the interest rate effects of budget deficits and debt," *Journal of the European economic Association*, 7, 858–885.

Levrero, E. S. and M. Deleidi (2019): "The causal relationship between short- and long-term interest rates: an empirical assessment of the United States," MPRA Paper 93608, University Library of Munich, Germany.

Lubik, T. A. and C. Matthes (2015): "Calculating the Natural Rate of Interest: A Comparison of Two Alternative Approaches," Economic Brief 15–10, Federal Reserve Bank of Richmond.

Mian, A., L. Straub, and A. Sufi (2021): "Indebted demand," *The Quarterly Journal of Economics*, 136, 2243–2307.

Mian, A. R., L. Straub, and A. Sufi (2020): "The saving glut of the rich," Tech. rep., National Bureau of Economic Research.

Peterman, W. B. (2016): "Reconciling Micro And Macro Estimates Of The Frisch Labor Supply Elasticity," *Economic Inquiry*, 54, 100–120.

Peterman, W. B. and E. Sager (2022): "Optimal Public Debt with Life Cycle Motives," *American Economic Journal: Macroeconomics*, 14, 404–37.

Röhrs, S. and C. Winter (2017): "Reducing Government Debt in the Presence of Inequality," *Journal of Economic Dynamics and Control*, 82, 1–20.

Appendix

A Robustness of Empirical Estimates **A-2**

B Construction of the Balanced Growth Path **A-3**

B.1 Aggregate Conditions A-5

B.2 Individual Conditions A-7

B.3 Aggregation A-8

C Calibration **A-9**

A Robustness of Empirical Estimates

The existing literature has noted that the interest rate elasticity with respect to public debt is somewhat sensitive to specification. As we showed, this is partially due to specifying the empirical relationship in levels. Here we check robustness around the specification for each of the features we introduce into estimation. We omit population aging throughout, focusing on the effect of the output gap in controlling for the persistence of business cycle shocks.

Table A1: Specification in Levels

	(1)	(2)	(3)	(4)	(5)
5yr Total Debt	0.020^{***}	0.024^{***}	-0.009	-0.005	-0.005
	(0.008)	(0.007)	(0.008)	(0.009)	(0.009)
Dependent Variable	$i_t^{5,10}$	$i_t^{5,10}$	$i_t^{5,10}$	$i_t^{5,10}$	$i_t^{5,10}$
Output Gap	t+1		t+1		t+5
GDP Growth	t+5	t+5	t+5	t+5	t+5
Controls	×	×	×	×	×
Sample period	1987-2017	1987-2017	1987-2020	1987-2020	1987-2020
N	89	89	95	95	95

*Notes: The dependent variable is the 5-year ahead 10-year forward interest rate. Standard errors in parentheses. * denotes significance at the 10% level, ** at the 5% level, and *** at the 1% level.*

Table A2: Baseline Specification in First Differences

	(1)	(2)	(3)	(4)
5yr Total Debt	0.054***	0.050***	0.049***	0.049***
	(0.019)	(0.018)	(0.018)	(0.018)
Dependent Variable	$\Delta i_t^{5,10}$	$\Delta i_t^{5,10}$	$\Delta i_t^{5,10}$	$\Delta i_t^{5,10}$
Output Gap	t+1	t+1	t+5	t+1 - t+5
GDP Growth	t+5	t+5	t+5	t+5
Controls	×	×	×	×
Sample period	1987-2017	1987-2020	1987-2020	1987-2020
N	89	95	95	95

*Notes: The dependent variable is the change in 5-year ahead 10-year forward interest rate. Standard errors in parentheses. * denotes significance at the 10% level, ** at the 5% level, and *** at the 1% level.*

Table A3: Primary Deficit and Forecast Horizon

	(1)	(2)	(3)	(4)	(5)	(6)
5yr Total Debt	0.049***	0.047***				
	(0.018)	(0.019)				
5yr Primary Debt			0.041**	0.041**		
			(0.021)	(0.021)		
10yr Primary Debt					0.014	0.010
					(0.013)	(0.011)
Dependent Variable	$\Delta i_t^{5,10}$	$\Delta i_t^{5,10}$	$\Delta i_t^{5,10}$	$\Delta i_t^{5,10}$	$\Delta i_t^{5,10}$	$\Delta i_t^{10,10}$
Output Gap	t+1 - t+5	t+1 - t+5	t+1 - t+5	t+1 - t+5	t+1 - t+10	t+1 - t+10
GDP Growth	t+5	t+5	t+5	t+5	t+10	t+10
Controls	×	×	×	×	×	×
Sample period	1987-2020	1996-2020	1987-2020	1996-2020	1996-2020	1996-2020
N	95	68	95	68	68	68

*Notes: The dependent variable is the change in 5-year ahead 10-year forward interest rate. Standard errors in parentheses. * denotes significance at the 10% level, ** at the 5% level, and *** at the 1% level.*

B Construction of the Balanced Growth Path

This appendix provides a formal construction of the Balanced Growth Path for the set of economies described in Section 2. We construct the Balanced Growth Path in multiple parts. In Appendix B.1, we construct the Balanced Growth Path for aggre-

Table A4: Primary Deficit Components

	(1)	(2)
10yr Primary Debt	0.010	
	(0.011)	
10yr Revenues		0.018
		(0.011)
10yr Mandatory		−0.082
		(0.053)
10yr Discretionary		0.009
		(0.025)
Dependent Variable	$\Delta i_t^{10,10}$	$\Delta i_t^{10,10}$
Output Gap	t+1 - t+10	t+1 - t+10
GDP Growth	t+10	t+10
Controls	×	×
Sample period	1996-2020	1996-2020
N	68	68

*Notes: The dependent variable is the change in 5-year ahead 10-year forward interest rate. Standard errors in parentheses. * denotes significance at the 10% level, ** at the 5% level, and *** at the 1% level.*

gates. In Appendix B.2, we construct the Balanced Growth Path for individual agents' allocations. Finally, in Appendix B.3, we put these elements together to characterize the balanced growth path for the distribution and the aggregation of households.

B.1 Aggregate Conditions

Balanced Growth Path: A Balanced Growth Path (BGP) is a sequence

$$\{C_t, A_t, Y_t, K_t, L_t, B_t, G_t, Tr_t\}_{t=0}^{\infty}$$

such that (i) for all $t = 0, 1, \ldots$ $C_t, A_t, Y_t, K_t, B_t, G_t, Tr_t$ grow at a constant rate g_y,

$$\frac{Y_{t+1}}{Y_t} = \frac{C_{t+1}}{C_t} = \frac{A_{t+1}}{A_t} = \frac{K_{t+1}}{K_t} = \frac{B_{t+1}}{B_t} = \frac{G_{t+1}}{G_t} = \frac{Tr_{t+1}}{Tr_t} = 1 + g_y$$

(ii) per capita variables all grow at the same constant rate g_w:

$$\frac{Y_{t+1}/N_{t+1}}{Y_t/N_t} = \frac{C_{t+1}/N_{t+1}}{C_t/N_t} = \frac{A_{t+1}/N_{t+1}}{A_t/N_t} = \frac{K_{t+1}/N_{t+1}}{K_t/N_t} = \frac{B_{t+1}/N_{t+1}}{B_t/N_t} = \frac{G_{t+1}/N_{t+1}}{G_t/N_t} = \frac{Tr_{t+1}/N_{t+1}}{Tr_t/N_t} = 1 + g_w$$

and (iii) effective labor per capita is constant:

$$\frac{L_{t+1}}{N_{t+1}} = \frac{L_t}{N_t} = \frac{L_0}{N_0}$$

Denote time 0 variables without a time subscript, for example $L \equiv L_0$.

Growth Rates: Let all growth derive from TFP $g_z > 0$ and population $g_n > 0$ growth. Then on a balanced growth path we assume:

$$Z_t = (1 + g_z)^t Z$$

$$N_t = (1 + g_n)^t N$$

where Z and N are steady state values. Then, from part (iii) of the definition, growth in labor is:

$$\frac{L_{t+1}}{L_t} = \frac{L_{t+1}/N_{t+1}}{L_t/((1 + g_n)N_t)} = 1 + g_n$$

In steady state $Y = ZK^{\alpha}L^{1-\alpha}$. Let output growth be given by $g_y > 0$. Therefore the

production function gives:

$$Y_t = Z_t K_t^{\alpha} L_t^{1-\alpha} \quad \Longrightarrow \quad (1+g_y) = (1+g_z)^{\frac{1}{1-\alpha}}(1+g_n)$$

Lastly, from parts (ii) and (iii) of the Balanced Growth Path definition, we can solve for the growth of per capita variables:

$$\frac{Y_{t+1}/N_{t+1}}{Y_t/N_t} = \frac{Z_{t+1}}{Z_t}\left(\frac{K_{t+1}/N_{t+1}}{K_t/N_t}\right)^{\alpha}\left(\frac{L_{t+1}/N_{t+1}}{L_t/N_t}\right)^{1-\alpha} \Longrightarrow (1+g_w) = (1+g_z)^{\frac{1}{1-\alpha}}$$

Prices: From Euler's theorem we know:

$$Y_t = \alpha Y_t + (1-\alpha)Y_t = (r_t+\delta)K_t + w_t L_t$$

Accordingly, the wage and interest rate depend on the capital-labor ratio. Growth in the capital-labor ratio is:

$$\frac{K_{t+1}/L_{t+1}}{K_t/L_t} = (1+g_z)^{\frac{1}{1-\alpha}} = 1+g_w$$

Therefore, the growth rate for the wage is:

$$\frac{w_{t+1}}{w_t} = \frac{Z_{t+1}}{Z_t}\cdot\left(\frac{K_{t+1}/L_{t+1}}{K_t/L_t}\right)^{\alpha} = 1+g_w$$

and the growth rate for the interest rate is:

$$\frac{r_{t+1}+\delta}{r_t+\delta} = \frac{Z_{t+1}}{Z_t}\cdot\left(\frac{K_{t+1}/L_{t+1}}{K_t/L_t}\right)^{\alpha-1} = 1$$

Therefore wages grow while interest rates do not.

Equilibrium Conditions: The detrended *asset market clearing condition* is:

$$K_t = A_t + B_t \quad \Longrightarrow \quad K = A + B$$

The detrended *resource constraint* is:

$$C_t + K_{t+1} + G_t + Tr_t = Y_t + (1-\delta)K_t \quad \Longrightarrow \quad C + (g_y+\delta)K + G + Tr = Y$$

and the detrended *government budget constraint* is:

$$G_t + Tr_t + B_{t+1} - B_t = R_t + rB_t \quad \Longrightarrow \quad G + Tr = R + (r - g_y)B$$

B.2 Individual Conditions

Preferences: We assume that labor disutility and utility over bequests have a time-dependent component. Specifically, we assume labor disutility grows at the same rate as the utility over consumption, such that $v_{t+1}(h,d) = (1+g_w)^{1-\sigma} v_t(h,d)$. Therefore, total utility over consumption, hours and retirement status is:

$$u(c_t) - v_t(h_t, d_t) = \left[(1+g_w)^{1-\sigma}\right]^t \left(u(c) - v(h,d)\right).$$

Similarly, assume that the non-homothetic component of the utility over bequests grows at the same rate as the utility over consumption, such that

$$\phi(a_{t+1}) = \left[(1+g_w)^{1-\sigma}\right]^{t+1} \phi(a') \ .$$

Bequests: Because bequest inflows $b_j(\kappa)$ are generated from the households' savings from the previous period, bequests grow at the same rate as savings.

Social Security: In order for the AIME to grow at the same rate as the wage, we assume a cost of living adjustment (COLA) on Social Security taxes and payments. For social security taxes, the cap on eligible income grows at the rate of wage growth, $\bar{m}_t = (1+g_w)^t \bar{m}$. Furthermore, base payment bend points $b^{ss}_{i,t} = (1+g_w)^t b^{ss}_i$ and base payment values $\tau_{r,i,t} = (1+g_w)^t \tau_{r,i}$ for $i = 1, 2, 3$.

Medical Expenses and Policy: We assume that medical expenses are indexed to wage growth and so $\mu_t = (1+g_w)^t \mu$ at all ages. We also assume that the consumption floor achieved by medical transfers, $\underline{c}_t$, also receives a COLA so that it grows at the same rate as the wage and therefore $\underline{c}_t = (1+g_w)^t \underline{c}$. By construction, this implies that individuals' medical transfers also grow at the same rate as the wage, $Tr_t = (1+g_w)^t Tr$.

Tax Function: On a Balanced Growth Path, (c_t, a'_{t+1}, a_t) and $\tilde{y}_t$ must all grow at the same rate as the wage. Furthermore, the tax function must grow at the same rate as

the wage. Recalling the tax function, $\mathrm{Y}_t(\tilde{y}_t)$, τ_2 must grow at the same rate as $\tilde{y}_t^{-\tau_1}$. Rewrite as:

$$\mathrm{Y}_t(\tilde{y}_t) = \tau_0 \left((1+g_w)^t \tilde{y} - \left([(1+g_w)^t]^{-\tau_1} \tilde{y}^{-\tau_1} + [(1+g_w)^t]^{-\tau_1} \tau_2 \right)^{-\frac{1}{\tau_1}} \right) = (1+g_w)^t \, \mathrm{Y}(\tilde{y})$$

Individual Budget Constraint: Let the function $T(h, a, \varepsilon, m, d)$ contain income taxes, social security taxes or payments, medical expenses or transfers, and bequest inflows that a household faces. A household's time t budget constraint is:

$$c_t + a'_{t+1} \leq w_t \varepsilon_t h_t + (1+r_t) a_t - T_t(h_t, a_t, \varepsilon_t, m_t, d_t)$$

$$c + (1+g_w) a' \leq w \varepsilon h + (1+r) a - T(h, a, \varepsilon, m, d)$$

where $\{c, a', a, h, w, r, \varepsilon\}$ are stationary variables. Given that the tax function $\mathrm{Y}(\tilde{y})$ grows at rate g_w, so will the transfer function $T(h, a, \varepsilon)$ in the infinitely lived agent model. Furthermore, given that the Social Security program $\{\bar{m}, b_i^{ss}, \tau_{r,i}\}$, medical system $\{\underline{c}, Tr, \mu\}$ and bequest inflows $\{b(\kappa)\}$ grow at rate g_w, so will the transfer $T(h, a, \varepsilon, m, d)$ function in the life cycle model.

B.3 Aggregation

Distributions: For j-th cohort at time t, the measure over $(a, \varepsilon, m, d_{-1})$ is given by:

$$\begin{aligned}
\lambda_{j,t}(a_t, \varepsilon, m_t, d_{-1}) &= \lambda_{j,t-1} \left(\frac{a_t}{1+g_w}, \varepsilon, \frac{m_t}{1+g_w}, d_{-1} \right) (1+g_n) \\
&= \lambda_{j,t-i} \left(\frac{a_t}{(1+g_w)^i}, \varepsilon, \frac{m_t}{(1+g_w)^i}, d_{-1} \right) (1+g_n)^i \quad \forall \; i \leq t \\
&= \lambda_j(a, \varepsilon, m, d_{-1}) N_{t-j+1}.
\end{aligned}$$

Therefore, $\lambda_j(a, \varepsilon, m, d_{-1})$ is a stationary distribution over age j households that integrates to one.

Aggregation: Aggregate consumption in the life cycle model is constructed as follows. Define the relative size of cohorts as $\omega_1 = 1$ and:

$$\omega_{j+1} = \frac{N_{t-j}}{N_t} \cdot \prod_{i=1}^{j} \psi_i = (1+g_n)^{-j} \prod_{i=1}^{j} \psi_i = \frac{\psi_j \omega_j}{1+g_n} \quad \forall \; j = 1, \ldots, J-1$$

Let $C_{j,t}$ be aggregate consumption per age-j household, which is derived from the age-j household's allocation:

$$C_{j,t} = \int (1+g_w)^t c_j(a,\varepsilon,m,d_{-1})\mathbf{d}\lambda_j = (1+g_w)^t \int c_j(a,\varepsilon,m,d_{-1})\mathbf{d}\lambda_j = (1+g_w)^t C_j$$

where C_j is the stationary aggregate consumption per age-j household. Accordingly, aggregate consumption is:

$$\begin{aligned} C_t &= N_t\Big(C_{1,t} + \psi_1(1+g_n)^{-1}C_{2,t} + \cdots + \left(\prod_{i=1}^{J-1}\psi_i\right)(1+g_n)^{-(J-1)}C_{J,t}\Big) \\ &= (1+g_w)^t N_t \sum_{j=1}^{J}\omega_j C_j \\ &= (1+g_y)^t C \end{aligned}$$

where C is the stationary level of aggregate consumption and where we have normalized $N = 1$.

We can similarly construct the remaining aggregates $\{A, K, Y, B, G\}$ on the balanced growth path. Notably, however, labor per capita does not grow. Aggregate labor per capita is constructed as:

$$L_t = N_t \sum_{j=1}^{J}\omega_j L_j \quad \Longrightarrow \quad L = \frac{L_t}{N_t} = \sum_{j=1}^{J}\omega_j \int d_j(a,\varepsilon,m,d_{-1})\varepsilon h_j(a,\varepsilon,m,d_{-1})\mathbf{d}\lambda_j$$

which is the sum over ages of aggregate labor per age-j household.

C Calibration

In this appendix, we detail the construction of empirical measures that we feed into the model.

Demographics: To measure household survival rates from mortality tables, we account for demographic changes such as variation in the size of the household and mortality rates by age and sex, and we define household mortality as when either both members of a married household die or when the sole remaining adult of a

household dies. Accordingly, we construct the household-level mortality rate as,

$$\psi_j \equiv \omega_j^{single}\left[\omega_j^{male}\psi_j^{male} + (1-\omega_j^{male})\psi_j^{female}\right] + (1-\omega_j^{single})\psi_j^{male}\psi_j^{female},$$

where ψ_j^{male} and ψ_j^{female} are survival probabilities for age j males and females respectively, and $\omega_j^{single}\omega_j^{male}$ is the fraction of single-adult households with an age-j male head of household. We obtain mortality rates by age and sex from Bell and Miller (2002) and derive $\{\psi_j\}_{j=1}^{J-1}$ by applying a quartic polynomial in age to the raw series.

Preferences: We compute the adult equivalent scale at each age (of the head of household) to convert households of varying sizes into a standardized measure,

$$\tilde{n}_j \equiv \left[\omega_j^{single}\cdot 1\right] + \left[(1-\omega_j^{single})\cdot 1.5\right] + (1/3)n_j^c$$

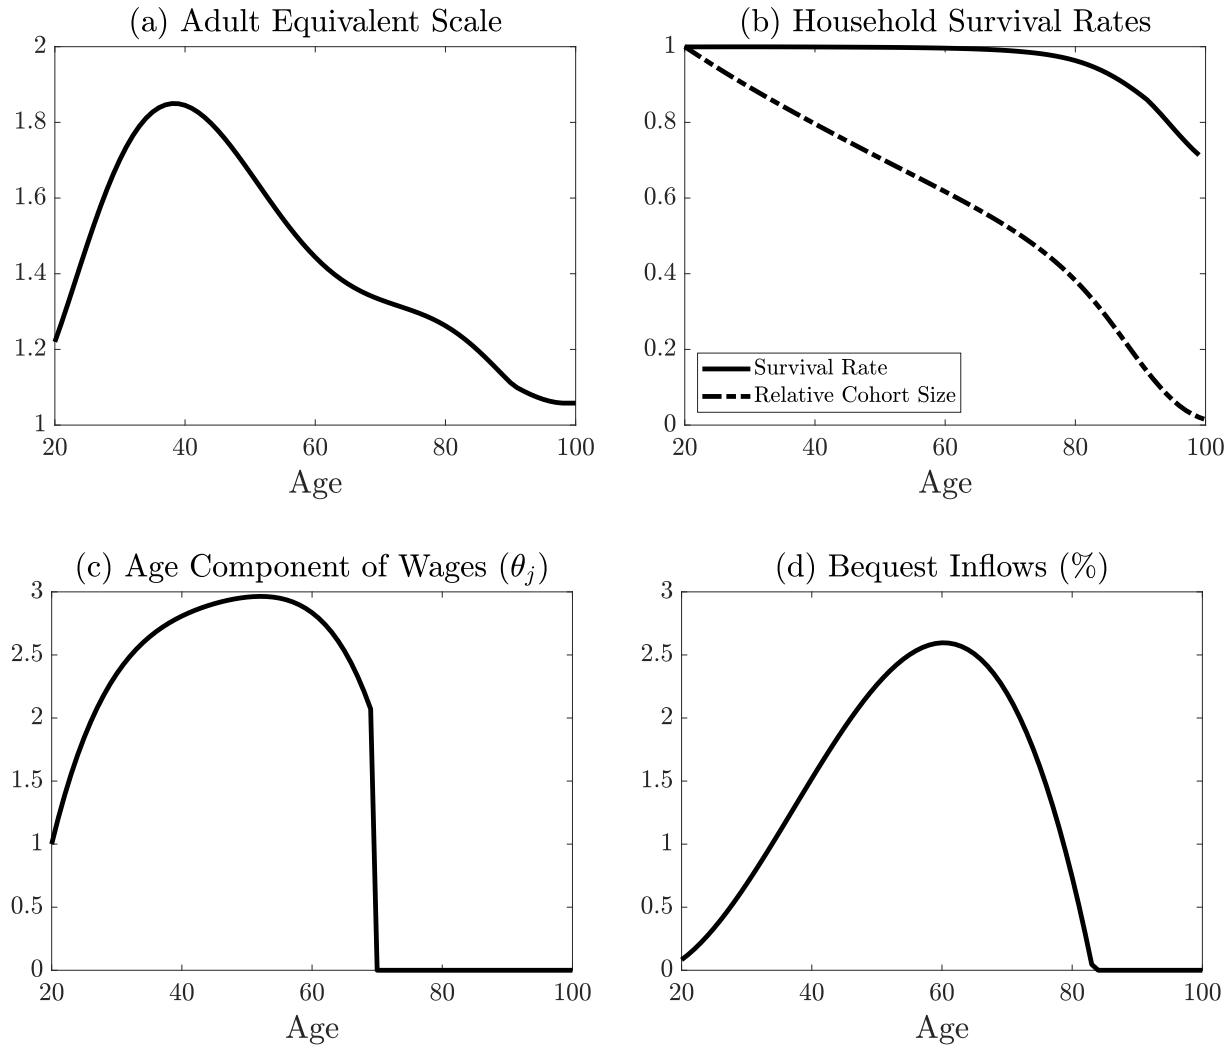
where ω_j^{single} is the fraction of single-adult households with an age-j head of household, and n_j^c is the average number of children in a household with an age-j head of household. For each household in the 2007 Survey of Consumer Finances, we observe a head of household, a spouse (if there is one), and children (if there are any). We compute the share of single-adult households with an age-j head of household, denoted ω_j^{single}, and the share of married households with age-j head of household, given by $1-\omega_j^{single}$. We derive $\{n_j\}_{j=1}^{J}$ by applying a quartic polynomial in age to the measured profile of $\{\tilde{n}_j\}_{j=1}^{J}$.

Bequests: When constructing the bequest distribution, we normalize the distribution by aggregate labor income. We use the HRS-AHEAD dataset, using estate inheritances and excluding intra-household bequests by dropping observations in which estates were transferred to a spouse. We use the CPI to convert to \$2002 and normalize by the Social Security's labor income adjustment (the 2002 Average Wage Index, https://www.ssa.gov/oact/cola/awidevelop.html).

The total level of bequests to type-κ households is determined by the total amount of wealth held by type-κ households upon death. For each type, these bequests are allocated to living households to match shares of bequests received by age in the Survey of Consumer Finances, according to the function $b_j(\kappa)$. In particular, we construct the function $b_j(\kappa)$ from the relationship

$$(1+g_n)b_j(\kappa) = \omega_j^b \cdot \sum_{j=1}^{J}\omega_j(1-\psi_j)\int a_j'(a,\varepsilon,m,d_{-1})\mathbf{d}\lambda_j(a,\varepsilon,m,d_{-1}|\kappa)$$

Figure A1: Household Size, Mortality, Wage, and Bequest Profiles



which must hold for each household-type κ and age-j, and where $\{\omega_j^b\}_{j=1}^J$ are the shares of bequests received from the Survey of Consumer Finances, scaled by the total bequests by type-κ households in the model. To compute $\{\omega_j^b\}_{j=1}^J$, we take total household-level bequests by age from Feiveson and Sabelhaus (2018, 2019) and apply a quartic polynomial in age and normalize to convert to lifetime shares by age such that $\sum_{j=1}^J \omega_j^b = 1$.

Medical Expenditures: We compute medical expenditures at each age as the weighted average for single and married households, controlling for the household's composition of men and women. To do so, denote average medical expenditures for age-j men and women by μ_j^m and μ_j^f, respectively, and recall that ω_j^{single} and ω_j^{male} denote the share of single households and share of single households with a male head, respec-

tively. Following DeNardi, French, and Jones (2010) and Kopecky and Koreshkova (2014), we compute out-of-pocket medical expenditures by sex and age (μ_j^m, μ_j^f) in the HRS-AHEAD dataset by regressing individual expenditures on a quartic polynomial in age and include individual-specific fixed effects in order to alleviate measurement error and survivorship bias. Accordingly, the average medical expenditures for a household with head of age-j are the weighted average of single male, single female and married household expenditures,

$$\tilde{\mu}_j = \omega_j^{single}\left[\omega_j^{male}\mu_j^m + (1-\omega_j^{male})\mu_j^f\right] + (1-\omega_j^{single})\left(\frac{1}{2}\mu_j^m + \frac{1}{2}\mu_j^f\right) .$$

We derive μ_j by applying a quartic polynomial in age to $\{\tilde{\mu}_j\}_{j=\bar{J}_{ret}}^{J}$.