Public Debt, Interest Rates and Wealth Inequality*

William B. Peterman
Federal Reserve Board

Erick Sager
Federal Reserve Board

May 30, 2023

Abstract

The U.S. federal debt-to-GDP ratio has almost doubled since the onset of the Great Recession, highlighting the importance of understanding the relationship between this debt and long term interest rates. Previous work finds empirically that a one percentage point increase in debt leads to a two-and-a-half basis point increase in interest rates. This paper revisits this relationship and finds the effect of public debt on the interest rate is twice as large, or five basis points, after addressing threats to identification from non-stationarity of public debt and endogeneity due to past debt obligations. While these estimates characterize the average effect from the historical changes in government debt, the relationship may vary with the type of policy that changes government debt and also may be affected by other structural changes in the economy. Thus, given that such estimates are often used for a specific policy evaluation, we examine if the estimated elasticity varies with respect to three dimensions (i) whether the change in debt is due to a legislative or macroeconomic shock, (ii) whether the change in debt is due to a change in discretionary outlays, mandatory outlays, or revenues, and (iii) how much wealth inequality exists in the economy. Overall, we find that when debt increases due to legislative changes in taxes the elasticity is twice as big, and that more wealth concentration leads to a smaller elasticity.

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. 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 hovering between thirty to fifty percent for three decades, federal government debt as a share of GDP has approximately doubled since the onset of the Great Recession. Given this recent large increase and the fact that under current policy the Congressional Budget Office projects debt to continue to rise by approximately 20 percentage points over the next ten years, it is important to understand the macroeconomic consequences of government debt. The literature largely examines these potential effects by focusing on the implications of government debt on long-term interest rate. For example, seminal work finds empirically that a one percent increase in the level of government debt leads to an increase in long term interest rates between two and five basis points, largely attributable to a crowding out of productive capital.[1] Previous empirical works assumes that there is one fixed relationship between government debt and interest rates. However, the relationship between public debt and interest rates could depend on both the source of debt change and also other fundamental factors within the economy.

In this paper we consider other channels through which public debt may impact interest rates and examine how the interest rate elasticity varies on three dimensions: (i) whether the change in debt is due to a legislative policies or to other macroeconomic shocks, (ii) whether the change in debt is due to a change in discretionary outlays, mandatory outlays or revenues, and (iii) how much wealth inequality exists in the economy. Examining the potential for the elasticity to vary along these dimension is of interest because the elasticity is often used to understand the macroeconomic implications of specific policy changes. Additionally, the aging population has already put – and is projected to continue to put – upward pressure on government deficits due to more spending on mandatory programs like social security. Finally, since the Great Recession, the concentration of wealth at the top of the distribution has increased dramatically and seems likely to persist.

We begin by revisiting the estimating strategy in these previous works but propose an improvement by focusing on the effect of changes in CBOs projections for government debt over the next five years as opposed to identifying the relationship from the levels. We show that prior estimates have suffered from bias due to non-stationarity in public debt-to-GDP ratios, which requires a reformulation of the estimation specification relative to prior work. We also show that prior estimates have not accounted for a source of endogeneity arising from existing debt obligations, namely that the

[1] See Gale and Orszag (2004), Engen and Hubbard (2004), Laubach (2009), as well as the more recent update in Gamber and Seliski (2019).

interest rates on existing debt are determined by shocks and variation in the data unrelated to fiscal policy. After addressing these threats to identification, we find that, on average, a one percentage point increase in government debt leads to about a 5 basis point increase in long term interest rates. The magnitude of this relationship is in the upward range of previous estimates. We find that our larger estimates are due in part to this new estimation strategy and also that we account for the potential for higher interest rates to increase government debt due to the effect on government borrowing costs.

Importantly, we find that both the source of the change in government debt along with the concentration of wealth have important implications for the relationship between government debt and interest rates. In particular, we find that as wealth inequality increases the effect of government debt on interest rates tends to be smaller, with a structural break around the time of the Great Recession. The smaller effect is of particular importance because it indicates that the large recent increase in overall government since 2008 may be putting less upward pressure on interest rates than previous estimates indicate. Moreover, depending on the size of the increase in inequality, it is possible that even with the increases in government debt since 2008 the overall effect of those particular increases on long term interest rates may be negative. In contrast, we find that changes in CBOs projection for government debt due to changes in revenues that are due to legislative action have a larger and more statistically significant affect on longer term interest rates, with estimates centered around 10bp. We posit that this is in part because markets view legislative changes in tax revenue as being particularly persistent as reductions are rarely unwound. These stronger empirical effects are important because as part of the Tax Cut and Job Act, legislative changes lowered tax revenue effective in 2018.

Finally, we use a quantitative life cycle model with incomplete markets that is calibrated to match the U.S. economy, following Peterman and Sager (2022). We explore how the model analogues of different types of fiscal policies can lead to differential effects on long term interest rates. Moreover, we examine the model mechanisms that lead a larger concentration of wealth inequality to put downward pressure on interest rates and lead to a lower elasticity.

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 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 previous 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.

The remainder of this paper is organized as follows. Section 2 sets up and estimates the relationship between interest rates and public debt. Section 3 details our model and characterizes the economic environment. Section 4 quantifies the model and discusses policy counterfactuals. Section 5 concludes.

2 Estimating the Interest Rate Elasticity of Debt

Strategy. We wish to estimate the response of the real interest rate to a change in public debt. Conceptually there are several threats to identification from naively regressing r_t on D_t, most importantly that at least in the short-term the interest rate varies with macroeconomic conditions and shocks that are unrelated to variation in public debt. 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. Given the plausibility of non-stationarity in the evolution of government debt (c.f. Barro (1979), Bohn (1998), and Campbell et al. (2023)) we construct forecast revisions, which are stationary and exhibit stable coefficient estimates across a host of empirical specifications in which forecasted levels exhibit instability.

Baseline Specification. The government's budget constraint is standard. 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 borrowing 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,

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

Our strategy for purging threats to estimation contained in short-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,

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

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,

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

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,

$$\Delta r_{t+j} = \alpha + \beta\Delta\eta_{t,t+j} + \gamma X_t + \varepsilon_t \quad (1)$$

where X_t is a set of time-varying controls for macroeconomic conditions that might confound the effect of fiscal shocks. We control for forecast revisions in total PCE inflation to effectively estimate the effect on the real interest, forecast revisions to GDP growth to adjust the interest rate for growth trends, the change in the dividend yield to control for the effect of equity returns on the interest rate, the changes of the stock of US Treasuries held by the Federal Reserve and, finally, the stock of US Treasuries held by foreign governments to control for composition of government debt.[2]

For the debt measure, we use the CBO's 5-year ahead debt forecast and take differences between successive CBO reports (comparing the same 5-year interval in each case) and express relative to 5-year ahead GDP at the time of later report. Finally, for the interest rate measure, we use updated yield curve estimates from Gürkaynak et al.

[2]This choice of control variables closely follows Laubach (2009) and Gamber and Seliski (2019), however relative to their methodologies we difference the control variables for use with our DiD estimator.

Table 1: Baseline Specification and Net Payments IV Estimates

	Δr_{t+J}					$\Delta \mathbb{E}_t \sum r_{t+j} D_{t-1+j}$
	(1)	(2)	(3)	(4)	(5)	(6)
Debt, $\Delta \eta_{t,t+j}$	0.053*** (0.017)			0.045** (0.018)		
Debt excluding Net Interest Payments		0.046** (0.020)				
Revenues (negative of)			0.041 (0.028)		0.057** (0.029)	
Mandatory Outlays			0.018 (0.070)		0.032 (0.075)	
Discretionary Outlays			-0.051 (0.043)		-0.027 (0.045)	
Net Interest Payments			0.291*** (0.077)		0.130 (0.870)	
Short-term Interest Rate						0.324*** (0.099)
Controls	✓	✓	✓	✓	✓	✓
Net Interest Instrument				✓	✓	
Observations	95	95	95	95	95	95
R^2	0.12	0.08	0.20	0.09	0.09	0.08

*Notes: The dependent variable in columns (1)-(5) is the change in 5-year ahead 10-year forward interest rate. The dependent variable in column (6) is the forecast revision to net interest payments. All sample periods are 1987m1-2020m3 and each observation corresponds to a CBO Budget Report. Standard errors in parentheses. * denotes significance at the 10% level, ** at the 5% level, and *** at the 1% level.*

(2007) to derive the 5-year ahead, 5-year interest rate.

Column (1) of table 1 shows that a 1 percent increase in public debt-to-GDP leads to an increase in the 5-year ahead, 5-year interest rate of 5.3 basis points. We take this as our baseline estimate, and notably it is at the upper end of prior estimates that specify the level of the interest and public debt-to-GDP ratio despite apparent non-stationarity in the data. While Gamber and Seliski (2019) find coefficient instability on samples estimated prior to and including the Great Recession, we find that our method is robust to such a choice of sample period.

Endogeneity of Existing Debt. The baseline specification provides a useful benchmark against existing results, especially given our strategy of studying forecast revisions and the changes in 5-year ahead interest rates. However, we uncover a form endogeneity in these existing estimates, which our baseline estimates inherit. We characterize the threat to identification and our solution now.

Although our baseline specification uses the debt-to-GDP ratio as a source of fiscal variation, debt can be decomposed into its constituent parts. We decompose debt by spending source. In particular, the coefficient β reflects the covariation between the change in the change in the interest rate at time $t + j$ and, concurrently, the forecast revisions to the $t + j$ flow of primary surplus relative to GDP, and the forecast revisions to net interest payments on existing debt (relative to GDP) at $t - 1 + j$. Using forecast decompositions from the CBO, we construct forecast revisions to government revenues (mainly inflows from taxation), mandatory government spending (such as entitlement programs), discretionary government spending (mainly due to changes in annual budgeting as well as emergency legislation), and net interest payments.

Notably, net interest payments includes the interest rate on previous debt obligations, which may reflect a host of other shocks. We deal with the potential endogeneity of net interest payments by measuring the effective interest rate on prior debt and measuring the effective interest rate in the CBO debt forecasts, and then orthogonalizing the net interest payment with respect to the effective interest rate at each period. This leaves us with a net interest payment measure that is correlated with the CBO's measures but orthogonal to the interest rate component that may embody the response to shocks not related to fiscal policy.

Column (2) of Table 1 estimates the relationship with the change in public debt forecasts when net interest payments are removed from the debt measure. Relative to the baseline, the interest rate effect is about a basis point smaller. Moreover, in column (3) we separately estimate the effect of each component of the public debt forecast revision and find that net interest payments are the only significant covariate. Thus, now we instrument for the net interest payment as described above – with first stage shown in column (6) – and find that when we regress the change in forward interest rates on a measure of debt in which the net interest payment component is now exogenous with respect to short term interest rates, then the interest rate effect is nearly the same as when net interest was excluded altogether. Finally in column (5), we separately estimate the effect of each component of the public debt forecast revision but now with exogenized net interest payments, and find that net interest payments are no longer significant but instead revenues are the only significant covariate. The interest rate effect on the overall tax revenue is nearly 6 basis points,

Table 2: Debt Decomposed by Type of Forecast Revision

	Δr_{t+J}			
	(1)	(2)	(3)	(4)
Debt, Legislative	0.047**		0.049*	
	(0.021)		(0.025)	
Debt, Economic	0.173***		0.165***	
	(0.048)		(0.061)	
Debt, Technical	-0.023		-0.014	
	(0.035)		(0.038)	
Revenues, Legislative		0.092**		0.117***
(negative of)		(0.044)		(0.041)
Revenues, Economic		0.102**		0.096
(negative of)		(0.058)		(0.066)
Revenues, Technical		0.016		0.048
(negative of)		(0.074)		(0.061)
Outlays, Legislative		0.007		0.013
		(0.037)		(0.036)
Outlays, Economic		0.231***		0.320***
		(0.065)		(0.111)
Outlays, Technical		0.024		0.051
		(0.068)		(0.067)
Controls	✓	✓	✓	✓
Instrumental Variables			✓	✓
Observations	95	95	95	95
R^2	0.25	0.19	0.20	0.25

*Notes: The dependent variable is the change in 5-year ahead 10-year forward interest rate. All sample periods are 1987m1-2020m3 and each observation corresponds to a CBO Budget Report. Standard errors in parentheses. * denotes significance at the 10% level, ** at the 5% level, and *** at the 1% level.*

slightly higher than the baseline effect in column (1).

Other Components of Debt. Next, we decompose debt by type of forecast revision. The CBO provides dollar values associated with debt forecasts due to legislative changes, economic changes and technical changes. Legislative changes refer to the

CBO's assessment of how newly passed legislation will affect debt, economic changes refer to how changes to the economy or the macroeconomic outlook affect the debt forecast, and technical changes mostly consist of internal and/or judgmental adjustments to the CBO forecast.

Column (1) of table 2 regresses the change in the 5-year ahead 5-year forward interest rate on public debt-to-GDP ratio forecast revisions that have been decomposed into the three factors described above. We find that legislative and economic justifications for the forecast revisions are significant, while technical adjustments are insignificant. In column(2) we decompose each of the justifications further into an outlay and a revenue concept, and we find that the significance of legislative changes to the debt forecast stems from revenue flows while the significance of economic changes stems from outlays.

Finally, the CBO does not consistently report how net interest payments factor into its forecast decomposition by legislative, economic and technical justification, and therefore cannot rule out that each of the variables may be endogenous. Accordingly, we repeat these two specifications from columns (1) and (2) after orthogonalizing each independent variable with respect to the short-term interest rate. In column (3) we find that the interest rate effects of the legislative, economic and technical forecast revisions to overall public debt are similar to those without instrumentation in column (1). However, the contributions of revenues and outlays by justification change more notably. In column (4) we find that the legislative component of revenue forecast revisions increases a few basis points while the economic component of outlays forecast revisions increases nearly ten basis points.

Heterogeneity by Wealth Inequality. The interest rate elasticity may vary with the wealth distribution. Such a relationship should inform models as to what underlying mechanisms are necessary to account for the evolution of interest rates over the past three decades. We study heterogeneity in the effect of fiscal shocks on interest rate changes by interacting the fiscal shock with a wealth inequality.

In our first specification, we measure wealth inequality as the ratio of the value of wealth held by the top 50% of the wealth distribution to that of the bottom 50%.[3] For ease of interpretation we demean the wealth inequality measure. Moreover, we lag wealth inequality in order to mitigate the effect of contemporaneous macroeconomic shocks on its evolution.[4] Finally, we see that this measure increased at an exceptionally fast rate during the great recession, and such variation could be an outlier

[3]See appendix A for robustness exercises with respect to inequality measures.

[4]We use the predicted value of from a regression of contemporaneous wealth inequality on four of its lags, following Hamilton (2018).

Table 3: Heterogeneity by Wealth

	Δr_{t+j}			
	(1)	(2)	(3)	(4)
Revenues (negative)	0.067**	0.057**	0.063**	0.069***
	(0.026)	(0.026)	(0.025)	(0.026)
Revenue × Net Worth (inequality measure)	0.001**			
	(0.01)			
Revenue × Assets (inequality measure)		-0.031**		
		(0.019)		
Revenue × Liabilities (inequality measure)		-0.330***		
		(0.130)		
Revenue × Net Worth (bottom 50% of distribution)			-0.201***	
			(0.056)	
Revenue × Net Worth (top 50% of distribution)			-0.142**	
			(0.072)	
Revenue × Assets (bottom 50% of distribution)				-0.655
				(0.689)
Revenue × Assets (top 50% of distribution)				-0.043
				(0.310)
Revenue × Liabilities (bottom 50% of distribution)				1.273***
				(0.362)
Revenue × Liabilities (top 50% of distribution)				-0.957**
				(0.571)
N	85	85	85	85
R^2	0.14	0.19	0.21	0.23

*Notes: The dependent variable is the change in 5-year ahead 10-year forward interest rate. All sample periods are 1990m6-2020m3 and each observation corresponds to a CBO Budget Report. Standard errors in parentheses. * denotes significance at the 10% level, ** at the 5% level, and *** at the 1% level.*

in our sample period. Hence, we interact wealth inequality with a dummy variable that equals zero during the great recession and one otherwise. The wealth inequality measure is thus,

$$\tilde{\sigma}_t \equiv \frac{\sum_i 1[i > p_{50}]a_{it}}{\sum_i 1[i \leq p_{50}]a_{it}} \times 1[t \notin (2007q1, 2008q3)] \ ,$$

with $\sigma_t = \tilde{\sigma}_t - (\sum_\tau \tilde{\sigma}_\tau)/T$. The resulting specification is,

$$\Delta r_{t+j} = \alpha + \beta \Delta \eta_{t,t+j} + \delta \Delta \eta_{t,t+j} \times \sigma_{t-1} + \gamma X_t + \varepsilon_t$$

we further focus on the revenues component of public debt, as this was the main channel in the previous empirical results. In table 3, column (1) shows that the interest rate does not vary much with overall net worth inequality. However, separately interacting asset inequality and liability inequality, in column (2), shows negative interest rate effects through both sources.

In our second specification, we gauge whether the top or bottom of the wealth distribution drives the wealth inequality channel. To do so, we separately include the numerator and denominator of the wealth inequality measure (interacting both with the time dummy to exclude the effect of the great recession). Column (3) of table 3 shows that while the an increase in wealth from both the bottom 50% and top 50% of the wealth distribution is associated with a less elastic interest rate. In column (4), splitting net worth into contributions from assets and liabilities shows that liabilities drive the interest rate effect from both the top and bottom of the distribution. Notably, an increase in liabilities from the bottom 50% of the wealth distribution leads to a more elastic interest rate, while an increase in liabilities from the top 50% of the wealth distribution leads to a less elastic interest rate.

3 Economic Environment

We will explore the economic effects of a change in debt within the quantitative life cycle model with incomplete markets, developed in Peterman and Sager (2022). In particular, we study a 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 examine why the interest rate elasticity varies by type of change to government debt and how wealth inequality interacts with these changes.

3.1 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 $g_y > 0$.

3.2 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.[5] Each period a new cohort 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:

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

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.[6] 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

[5]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.

[6]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 4. 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 **??**).

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:

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

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 and, thus, the optimal public policy, 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.[7]

3.3 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 $b_{base}^{ss}(m_{J_{ret}})$, is computed as a piecewise-linear function over the household's average

[7] 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.

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:

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

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 ($\underline{c}$).[8]

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:

$$G + Tr + (1 + g_y)B' - B = rB + R \tag{2}$$

where R is aggregate revenues from income taxation and G is the level of government expenditures that are independent from aggregate medical transfers, Tr.[9] The model's Social Security system is self-financing and therefore does not appear in the governmental budget constraint.

3.4 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:

$$\begin{aligned} V_j(a, \varepsilon, m, 1) = \max_{c, a', h} \quad & \left[u_j(c) - v(h, 1)\right] + [\beta(1 + g_w)^{1-\sigma}]\psi_j \sum_{\varepsilon'} \pi_j(\varepsilon'|\varepsilon) V_{j+1}(a', \varepsilon', m', 1) \\ & + [\beta(1 + g_w)^{1-\sigma}](1 - \psi_j)\phi(a') \end{aligned} \tag{3}$$

subject to

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

[8]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.

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

where g_w is the steady state rate of wage growth and σ is the coefficient of relative risk aversion (see Section 4 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:

$$V_j(a, \varepsilon, m, 0) = \max_{c,a'} \quad u_j(c) + [\beta(1 + g_w)^{1-\sigma}]\psi_j V_{j+1}(a', \varepsilon, m, 0) \\ + [\beta(1 + g_w)^{1-\sigma}](1 - \psi_j)\phi(a') \tag{4}$$

subject to

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

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 **??**.

3.5 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 (3) and (4) 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 $(a, \varepsilon, m, d_{-1})$ 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 (2), where aggregate income tax revenue is given by:

$$R \equiv \sum_{j=1}^{J} \omega_j \int \Upsilon \left(y_j(h_j(a, \varepsilon, m, d_{-1}), a, \varepsilon, d_j(a, \varepsilon, m, d_{-1})) \right) \mathbf{d}\lambda_j(a, \varepsilon, m, d_{-1}), \tag{5}$$

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:

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

$$= \sum_{j=1}^{J} \omega_j \int (1 - d_j(a, \varepsilon, m, d_{-1})) b_{ss}(m) \mathbf{d}\lambda_j(a, \varepsilon, m, d_{-1}). \tag{6}$$

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

$$K - B = \sum_{j=1}^{J} \omega_j \int \left(a + b_j(\kappa)\right) \ \mathbf{d}\lambda_j(a, \varepsilon, m, d_{-1}) \tag{7}$$

$$L = \sum_{j=1}^{J} \omega_j \int d_j(a, \varepsilon, m, d_{-1}) e(\varepsilon) h_j(a, \varepsilon, m, d_{-1}) \ \mathbf{d}\lambda_j(a, \varepsilon, m, d_{-1}) \tag{8}$$

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

$$C + (1 + g_y)K' + G + Tr = ZF(K, L) + (1 - \delta)K \tag{9}$$

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:

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

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 $j \in \mathbf{J}$.

4 Calibration

4.1 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 4 summarizes the target, source and value for each parameter, and Table 5 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,

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

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

Table 4: 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

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

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.

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).[10]

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.[11] 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 redis-

[10] 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.

[11] 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.

tributed 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.[12]

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.[13] 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 (v_{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)).[14] The value of ν_{max} is 8.8, which implies that a superstar earns approximately 9 times more

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

[13] 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.

[14] 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.

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.[15]

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 $\tau_1 = 0.768$, 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).[16] 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

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

[16] 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 5: 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

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.

4.2 Baseline Economy

The model matches the set of targeted moments very closely, as shown in Table 5. 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 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 en-

Figure 1: Wealth and Bequest Distributions, Model and Data

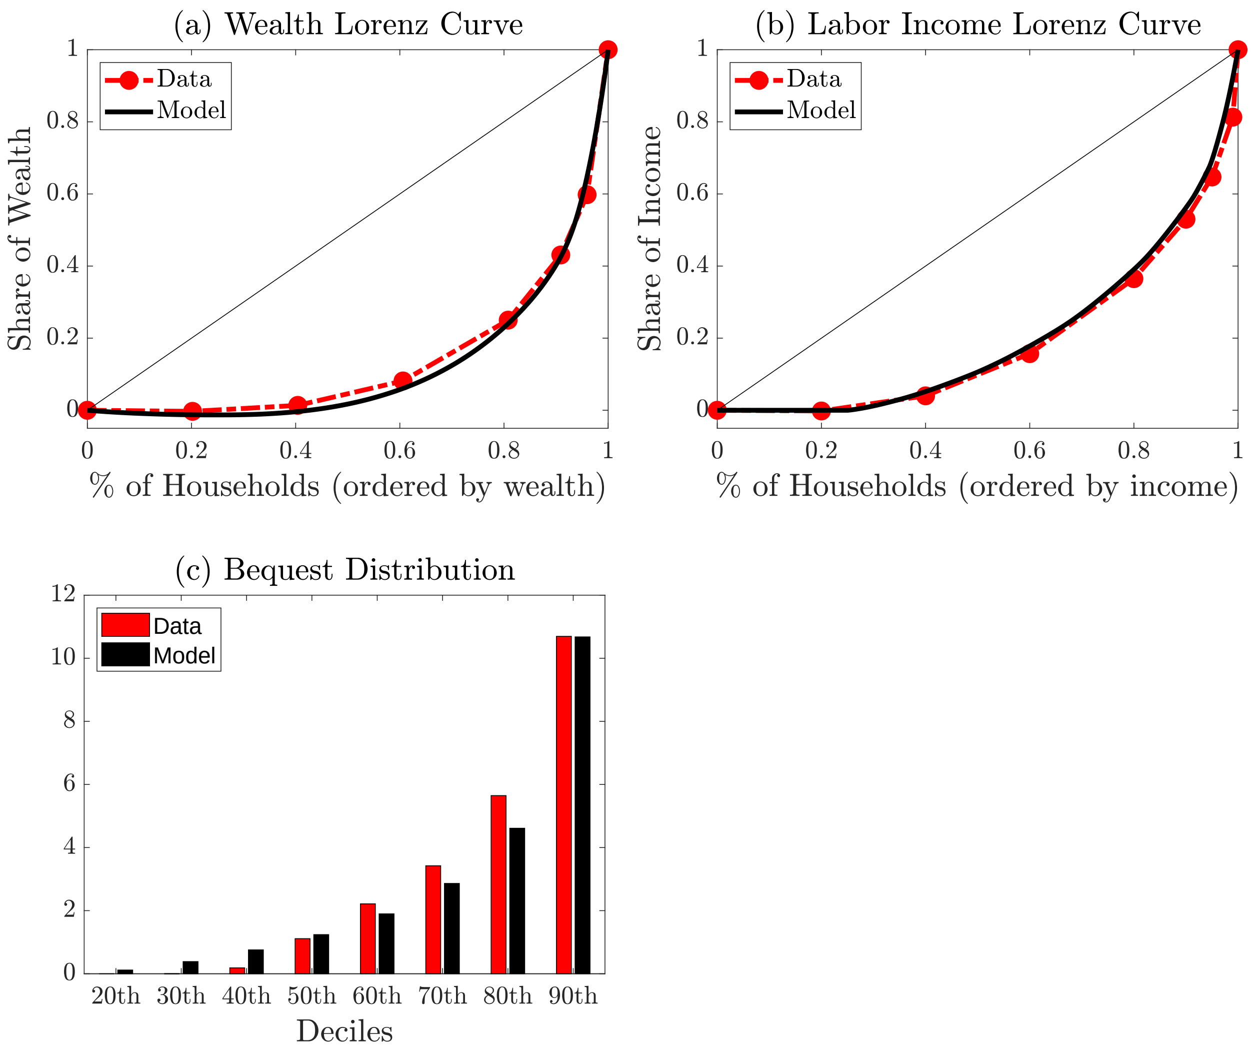

tire wealth distribution;[17] panel (b) shows that the model closely matches the Lorenz curve for labor income; and panel (c) shows that the model matches the distribution of bequests quite well.[18]

Next, because the evolution of households' allocations along the life cycle will be

[17] 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.

[18] 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 2: Life Cycle Profiles, Model and Data

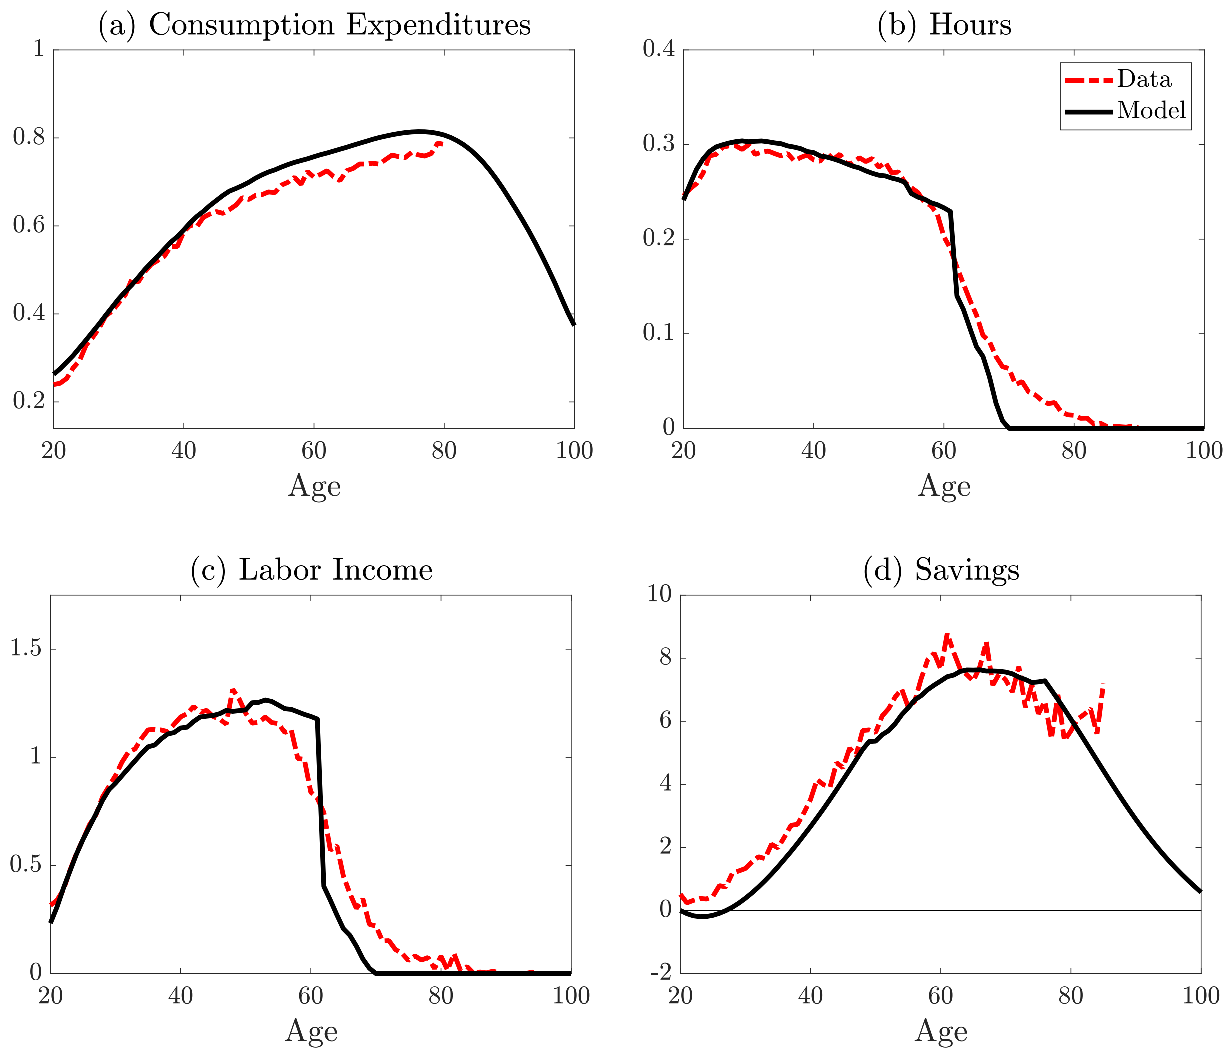

Notes: Solid lines depict the model's life cycle profile under the baseline public debt policy (67% of output). Dashed lines depict empirical life cycle profiles. Hours profiles are normalizd by total annual hours. All other model generated life cycle profiles are normalized by the model's aggregate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

a key determinant of optimal 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 over the available range of ages, both in terms of contour and level.[19] Likewise, from panels

[19]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

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

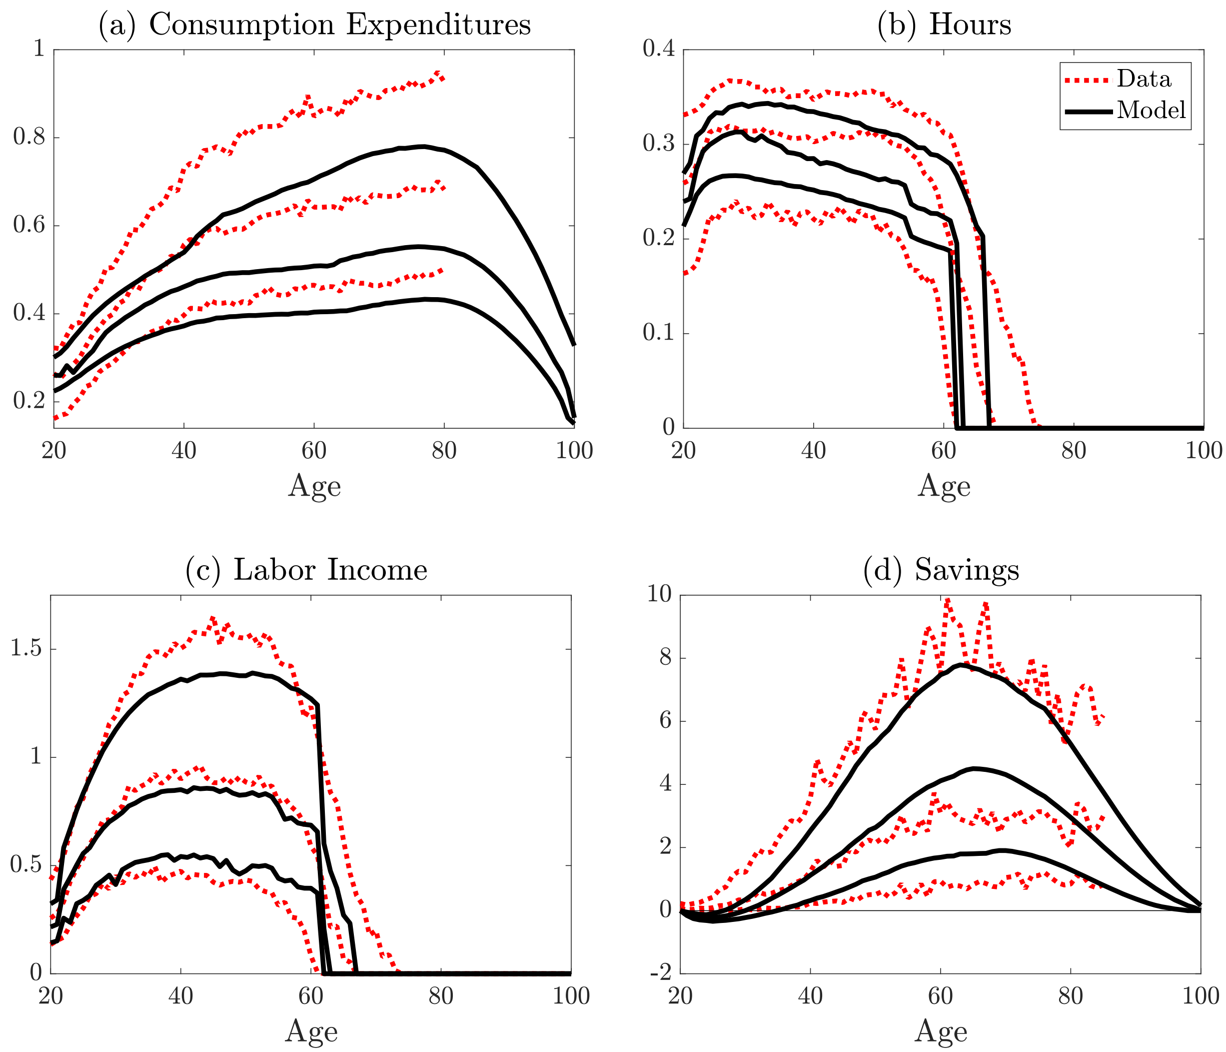

Notes: Red dashed lines depict data percentiles (25th, 50th, 75th) that have been normalized by PSID aggregate labor income. Black solid lines depict model percentiles (25th, 50th, 75th) that have been normalized by model generated aggregate labor income. The hours profile was normalized by total annual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

(b) and (c), average household hours and labor income in the model are very close to those in the data at each age.[20] Lastly, panel (d) demonstrates that the model

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.

[20] 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 hours (40 hours per week times 52 weeks per year) and the number of earners within the household

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 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 the life cycle.

4.3 Policy Counterfactuals

Given a model that is calibrated to recent U.S. data on households and macroeconomic aggregates, we can use the model to measure the impact of changes to fiscal policies that add to the public debt – such as, tax rates, transfers, government spending, and entitlement program transfer rates – on changes in the interest rate at varying maturities (horizons).

5 Conclusion

The U.S. federal debt-to-GDP ratio has almost doubled since the onset of the Great Recession, highlighting the importance of understanding the relationship between this debt and long term interest rates. Previous work finds that empirically a one percentage point increase in debt leads to a two and a half basis point increase in interest rates. This paper revisited this relationship and finds the effect of public debt on the interest rate is twice as large, or five basis points, after addressing threats to identification from non-stationarity of public debt and endogeneity due to past debt obligations. While these estimates characterize the average effect from the historical changes in government debt, the relationship may vary with the type of policy that is causing a change in government debt and also may be affected by other structural

(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.

changes in the economy. Thus, given that these estimate are often used for a specific policy evaluation, we examine if the estimated elasticity varies with respect to three dimensions (i) whether the change in debt is due to a legislative or macroeconomic shock, (ii) whether the change in debt is due to a change in discretionary outlays, mandatory outlays, or revenues, and (iii) how much wealth inequality exists in the economy. Overall, we find that when debt increases due to legislative changes in taxes the elasticity is twice as big and that more wealth concentration leads to a smaller elasticity.

References

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.

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.

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.

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.

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.

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.

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.

Hamilton, J. D. (2018): "Why You Should Never Use the Hodrick-Prescott Filter," *The Review of Economics and Statistics*, 100, 831–843.

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

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.

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

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.

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.