Fiscal Beliefs and the Effectiveness of Fiscal Consolidation

Nayoung Lee *

November 10, 2026

Abstract

This paper studies how the distribution of household fiscal beliefs shapes the effectiveness of tax-based fiscal consolidation. Motivated by empirical evidence of heterogeneous tax-hike expectations in the FRBNY Survey of Consumer Expectations, I incorporate fiscal beliefs into a tractable model. The model shows that, holding average beliefs fixed, the distribution of fiscal beliefs across households with different marginal propensities to consume (MPCs) and wealth substantially alters debt dynamics through aggregate demand and the tax base. I then quantify this mechanism in a two-group HANK model calibrated to cross-group differences in fiscal beliefs, MPCs, income, and wealth shares. Following a persistent one-percentage-point labor-income-tax increase, average fiscal beliefs raise the fiscal sacrifice ratio—cumulative output loss per unit of debt reduction—over the medium run relative to the full-information rational-expectations benchmark. However, belief heterogeneity offsets about one-third of this increase because beliefs closer to full-information rational expectations are held by households that account for a disproportionate share of aggregate demand.

JEL Classification: E62, D84, E21, H63.
Keywords: fiscal beliefs; heterogeneous expectations; fiscal policy; marginal propensities to consume; government debt.

1 Introduction

Fiscal consolidation is not a purely accounting exercise. A tax increase or spending reduction improves the primary balance mechanically, but its ultimate effect on government debt also depends on how households adjust consumption, how aggregate output responds, and how these changes feed back into the tax base. The strength of these endogenous fiscal feedbacks depends critically on how households understand the fiscal intervention.

Standard macroeconomic models typically assume that households observe fiscal shocks and correctly infer their implications for current disposable income, future taxes, and equilibrium activity. This full-information rational-expectations (FIRE) benchmark is difficult to reconcile with survey evidence documenting substantial heterogeneity in households' perceptions of public debt and future fiscal policy [Andrade et al., 2025, Bianchi et al., 2025]. Households differ not only in what they expect the government to do, but also in how strongly those expectations are reflected in their income and consumption plans. These differences raise a question that is

*Department of Economics, Indiana University. E-mail: nayolee@iu.edu

largely absent from standard analyses of debt sustainability: does it matter which households understand fiscal policy correctly?

This paper studies how the distribution of fiscal beliefs across households shapes the effectiveness of fiscal consolidation. The central argument is that average fiscal beliefs are not sufficient for evaluating fiscal policy. The same belief can have different aggregate consequences depending on the MPC, wealth, and income of the household that holds it. As a result, two economies with identical average fiscal beliefs can exhibit different paths for aggregate demand, tax revenue, and government debt when those beliefs are distributed differently across households.

I begin by presenting descriptive evidence from the Federal Reserve Bank of New York Survey of Consumer Expectations (SCE), combining the Public Policy Survey with variables from the core questionnaire. Tax-hike expectations vary systematically with respondents' perceived exposure to tax changes and with household income. Households that expect tax changes to affect them negatively assign higher probabilities to future income-tax increases, and low-income households report higher tax-hike expectations than middle- and high-income households across perceived-exposure groups. Furthermore, the SCE Household Survey shows that differences in income are associated with differences in MPCs.

These patterns highlight a key feature standard analyses of fiscal consolidation abstract from: fiscal beliefs vary systematically with household characteristics that govern fiscal transmission. To examine the theoretical mechanism, I incorporate this observed belief heterogeneity into the model. I then clarify how its interaction with MPCs and wealth shapes aggregate outcomes and determines the effectiveness of fiscal consolidation.

To this end, I first study the mechanism in an overlapping-generations model with government debt and a fiscal shock, building on Angeletos et al. [2024]. To depart from full-information rational expectations, I adopt the generalized non-FIRE framework of Angeletos et al. [2025]. Household beliefs are summarized by two parameters: λ^{PE} , which disciplines how accurately households perceive the fiscal shock, and λ^{GE}, which captures how households believe others perceive the shock, thereby governing general equilibrium effects. This flexible formulation nests a wide range of behavioral models, making it particularly suitable for my application. While my empirical findings document that belief heterogeneity exists, they do not uniquely identify the underlying behavioral mechanism. Therefore, I discipline beliefs using the SCE data without committing to a specific cognitive friction. Firms and the fiscal authority remain fully informed.

In this baseline, all households are ex ante identical: they share the same beliefs and the same marginal propensity to consume. I refer to it as the one-type model, a benchmark for the two-type model that incorporates the joint distribution of beliefs, MPCs, and wealth.

The one-type model isolates two channels through which household characteristics shape the effectiveness of fiscal consolidation following a tax shock. First, comparing a FIRE economy to a non-FIRE economy, debt stabilizes more effectively under non-FIRE beliefs. Intuitively, when households perceive the tax shock to be smaller than it actually is, aggregate demand adjusts less, mitigating tax base erosion and allowing debt to stabilize more rapidly. This constitutes the **belief channel**. Second, comparing a high-MPC economy to a low-MPC economy, a higher aggregate MPC yields higher remaining debt. Because high-MPC households respond more

strongly to tax shocks, the tax base erodes more severely, dampening debt reduction. This constitutes the **MPC channel**.

I then extend the framework to a two-type model which household differ in their fiscal beliefs, MPCs and government debt holdings. Key to this extension is that fiscal belief heterogeneity enters the model jointly with MPC and wealth heterogeneity. Holding the population average belief fixed, I compare a benchmark common-belief scenario against two heterogeneous-belief scenarios: one where low-MPC households hold beliefs closer to FIRE, and another where high-MPC households hold beliefs closer to FIRE. Relative to the common-belief baseline, belief heterogeneity weakens the effectiveness of fiscal consolidation in both cases.

However, the underlying mechanisms through which the distribution of fiscal beliefs impairs policy outcomes differ fundamentally. When low-MPC households hold beliefs closer to FIRE, weaker demand and higher debt arise through their larger wealth and population shares, which carry greater weight in the aggregate demand. In contrast, when high-MPC households hold beliefs closer to FIRE, higher debt is driven by their strong responsiveness to tax shocks (their high MPC), despite their small population and wealth shares. The two-type model thus demonstrates that average fiscal beliefs are insufficient for evaluating fiscal consolidation effectiveness. It is the distribution of fiscal beliefs that shapes the effectiveness of fiscal consolidation. Next, I quantify these mechanisms in a two-group HANK model with richer household heterogeneity.

The quantitative model is a two-group HANK model featuring heterogeneity in beliefs, MPCs, income, and wealth. I model fiscal consolidation as a one-percentage-point increase in the common labor income tax rate, which provides a tractable mapping to the SCE measure of expected tax rate changes. I discipline the first-order belief parameters (λ^{PE}) using the ratio of the subjective tax expectation response to the realized tax rate response following a high-frequency fiscal shock. For second-order beliefs (λ^{GE}) , I adopt a conservative approach by specifying them as a weighted average of first-order beliefs. I subsequently evaluate the robustness of this assumption using prior predictive analysis. The remaining parameters are calibrated to target the liquid wealth distribution in the Survey of Consumer Finances (SCF) and income shares from the Current Population Survey (CPS).

I compare three belief specifications: FIRE, a common-belief specification that preserves the population-average belief while removing cross-group dispersion, and calibrated heterogeneous beliefs. Comparing FIRE to the common-belief scenario isolates the **level effect** of fiscal beliefs, whereas comparing the common-belief scenario to the heterogeneous-belief scenario isolates the **dispersion effect** of fiscal beliefs.

To evaluate fiscal consolidation, I define the fiscal sacrifice ratio as the cumulative output loss per unit of debt reduction. This measure recognizes that a larger debt reduction does not necessarily represent a more effective consolidation if it is achieved through a disproportionately large and persistent contraction in economic activity. Following a persistent one-percentage-point labor income tax increase, average fiscal beliefs raise the sacrifice ratio from 1.55 under FIRE to 2.11 under common beliefs over a five-year horizon. Belief heterogeneity lowers it to 1.90, offsetting approximately one-third of this increase. This offset arises because households holding beliefs closer to full information account for a disproportionate share of aggregate demand, thereby pulling the economy toward FIRE. The sign and magnitude of the offset are robust to

an alternative timing of tax perception and to prior-predictive exercises that vary belief and structural parameters.

This paper contributes to three strands of literature. First, it contributes to the literature on fiscal consolidation. Alesina et al. [2015] use narratively identified multi-year fiscal plans to study the output effects of consolidation across OECD countries, while Erceg and Lindé [2013] compare tax- and spending-based consolidations in a currency union. I complement this literature by studying how household beliefs shape the output–debt trade-off of tax-based consolidation. The same statutory tax increase can generate different output and debt paths depending on how the fiscal consolidation and its future consequences are perceived across households.

Second, the paper contributes to the literature on fiscal policy under departures from full-information rational expectations. Existing work shows that cognitive discounting, sparse information, and higher-order belief distortions can affect fiscal multipliers and generate departures from Ricardian equivalence [Bianchi-Vimercati et al., 2024, Eichenbaum et al., 2026, Gabaix, 2020]. My contribution is threefold. First, I use the generalized non-FIRE framework to discipline beliefs with survey evidence without committing to a particular cognitive friction. Second, I shift the focus from average beliefs to their distribution across households. Third, I study government-debt dynamics and the output cost of debt reduction rather than fiscal multipliers alone.

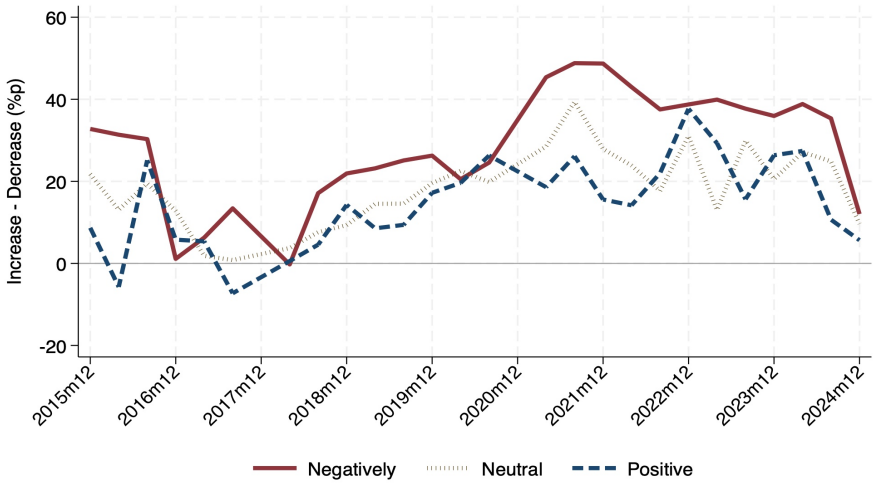
Third, the paper connects the HANK literature, which emphasizes the roles of heterogeneous MPCs, income exposure, and liquid wealth in aggregate-policy transmission [Auclert., 2019, Kaplan et al., 2018], with a growing literature that incorporates departures from full-information rational expectations into heterogeneous-agent economies [Angeletos et al., 2025, Farhi and Werning, 2019, Pfäuti and Seyrich, 2026]. The closest related paper is Guerreiro [2023], who endogenizes household attention as an optimal response to income cyclicality and shows that the resulting correlation between beliefs and exposure amplifies business-cycle fluctuations and makes government-spending multipliers depend on the composition of spending. I take a complementary approach along two dimensions. First, rather than deriving attention from an information-choice problem, I discipline perception of tax policy directly using SCE evidence. Second, rather than studying government-spending shocks, I study tax-based consolidation in an economy with government debt and a fiscal rule, where households' beliefs about future taxes shape macroeconomic consequences of fiscal consolidation. This shifts the focus from how much beliefs amplify fluctuations to how they change the output cost of reducing debt.

The remainder of the paper is organized as follows. Section 2 presents descriptive evidence on heterogeneity in household fiscal expectations and their association with household characteristics. Section 3 develops the generalized non-FIRE OLG framework, first with a representative household type and then with two types that differ in their MPCs, government debt holdings, and fiscal beliefs. Section 4 quantifies this mechanism using the HANK model. Section 5 concludes.

2 Empirical Evidence

This section documents heterogeneity in household fiscal beliefs using Federal Reserve Bank of New York Survey of Consumer Expectations Public Policy Survey. The survey elicits households' subjective expectations about future changes in government policies, including whether taxes or benefits are likely to increase, decrease, or remain unchanged. I focus on expectations about income tax rates, since income taxes constitute a central component of federal revenues. I measure net tax-increase expectations as the reported probability of an increase minus the reported probability of a decrease, expressed in percentage points. The survey also asks respondents whether they expect a given policy change to affect them personally in a positive, negative, or neutral way. I use this response to construct a measure of perceived exposure to fiscal policy. Further details on the questionnaire are provided in the Appendix A.

Figure 1. Tax-Hike Expectations by Perceived Exposure



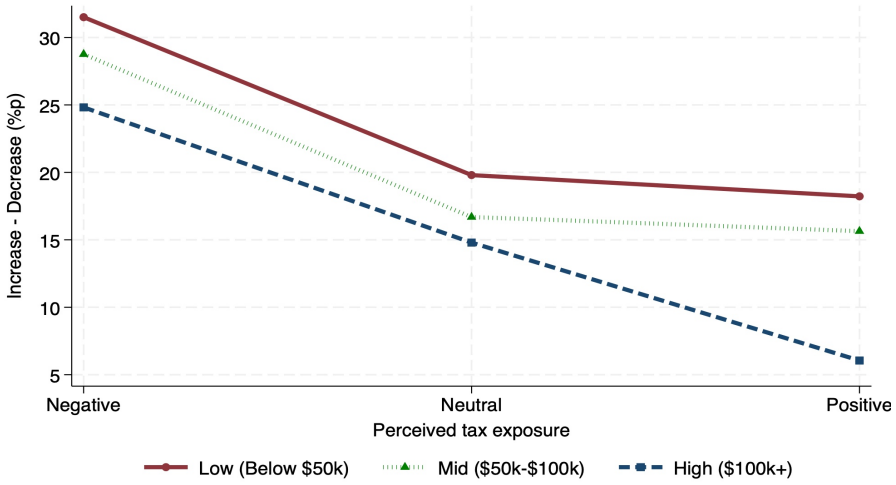
Source: FRBNY SCE Public Policy Survey. After winsorizing the data at the 1st and 99th percentiles, approximately 74%, 17%, and 9% of households in the full sample reported being negatively, neutrally, and positively affected, respectively.

Figure 1 plots average net tax-hike expectations by perceived exposure from 2015 to 2024. Households that expect tax changes to affect them negatively assign systematically higher probabilities to future income-tax increases than households expecting neutral or positive effects. This pattern indicates that fiscal expectations are heterogeneous across households and vary systematically with respondents' perceived exposure to fiscal policy.

Figure 2 further disaggregates net tax-hike expectations by income group. Two distinct patterns emerge from the data. First, within each income group, negative perceived exposure is associated with higher tax-hike expectations. In addition, low-income households report higher tax-hike expectations than middle- and high-income households across all perceived-exposure groups. These patterns suggest that fiscal beliefs vary systematically with household income.

Furthermore, these income differences are associated with differences in marginal propensities to consume. Table 1 presents the MPCs across the same income groups, estimated using the SCE Household Spending Survey following Kosar et al. [2025]. The survey asks households how they would allocate a hypothetical 10% rise or fall in income across spending, saving, and debt

Figure 2. Tax-Hike Expectations by Income Group and Perceived Exposure



adjustment, or a combination of the three, and to report the share assigned to each component. I then define the MPC as the share allocated to changes in spending.

Table 1. MPC Estimates by Income Group

Income Group	**Decrease in Income**	**Increase in Income**
Low	0.7617	0.1793
Mid	0.7365	0.1729
High	0.7118	0.1740
All	0.7400	0.1765

Source: FRBNY SCE Household Survey.

The estimates reveal heterogeneity in MPCs across income groups, particularly following a negative income shock. Specifically, the MPC for an income loss decreases monotonically with income, being highest for low-income households.This finding is central to this paper, as tax-based fiscal consolidation inherently entails a reduction in disposable income. Conversely, MPCs in response to a positive income shock are substantially smaller in magnitude, with minimal variation across groups.

Together, these findings demonstrate that tax-hike expectations vary systematically by household income and perceived policy exposure. Because lower-income households feature higher MPCs, the aggregate macroeconomic effects of fiscal expectations depend critically on the cross-sectional distribution of beliefs rather than just their mean level, with direct implications for aggregate demand and the tax base.

The empirical results therefore motivate a departure from the benchmark full-information rational expectations environment in which households share common fiscal beliefs. If households differ both in their fiscal beliefs and in the extent to which those beliefs are reflected in consumption, the effectiveness of fiscal consolidation may depend on the distribution of beliefs across households. I therefore use the survey evidence to motivate two qualitative features of the model: fiscal beliefs vary across households, and their association with spending differs across household

groups. The model developed below therefore introduces non-FIRE distortions in fiscal beliefs and examines how their interaction with MPC and wealth heterogeneity shapes debt dynamics and the effectiveness of fiscal consolidation.

3 Simple model

I incorporate fiscal beliefs into the tractable OLG framework of Angeletos, Lian, and Wolf (2024) to study how expectations drive aggregate demand, the endogenous tax base, and government debt. To model expectations, I adopt the generalized non-FIRE approach of Angeletos, Guerreiro, and Zhang (2025), which parameterizes households' first- and higher-order beliefs without committing to a particular cognitive friction. Section 3.1 first studies a one-type OLG benchmark in which households are ex ante identical in their beliefs and MPCs. Section 3.2 then extends this framework to a two-type model featuring heterogeneity in MPCs, fiscal beliefs, and debt holdings. This extension incorporates the joint distribution of beliefs, MPCs, and wealth, allowing fiscal beliefs to shape policy outcomes through aggregate demand and the tax base.

3.1 One-Type OLG

I build on the OLG framework of Angeletos et al. [2024], which features fiscal innovations, by incorporating a non-FIRE belief structure. In this setup, we distinguish between the non-FIRE beliefs, denoted by $\bar{E}$, and the standard rational expectations under FIRE, denoted by $\mathbb{E}$.

3.1.1 Model

Current consumption is determined by the marginal propensity to consume out of total wealth, which comprises financial wealth (d_t) and human wealth (the expected value of $y_{t+k} - t_{t+k}$). The final term captures the intertemporal substitution effect of varying interest rates

$$c_t = (1-\beta\omega)\left(d_t + \bar{E}_t\left[\sum_{k=0}^{\infty}(\beta\omega)^k(y_{t+k}-t_{t+k})\right]\right) - \left(\sigma\beta\omega - (1-\beta\omega)\frac{D}{Y}\right)\bar{E}_t\left[\sum_{k=0}^{\infty}(\beta\omega)^k r_{t+k}\right]$$

This setup is built on the OLG framework, so the on-impact MPC out of wealth equals $1-\beta\omega$, where $\omega \in (0,1)$ is the survival probability.

Inflation follows a standard New Keynesian Phillips Curve:

$$\pi_t = \beta\mathbb{E}_t[\pi_{t+1}] + \kappa y_t,$$

where y_t is the output gap, and $\kappa > 0$ measures the slope of the Phillips curve.

The government issues one-period real debt d_t, which evolves according to

$$d_t = (1+r_{t-1})d_{t-1} + g_t - T_t$$

Taxes follow the rule

$$T_t = \tau_y y_t + \tau_d(d_t + \varepsilon_t) - \varepsilon_t$$

where τ_y captures the tax-base channel through which output affects tax revenues, τ_d governs the debt feedback rule, and ε_t represents an exogenous fiscal shock. While I use lump-sum taxes in the simple model to build intuition, I incorporate proportional taxation in the quantitative model.

The central bank sets the interest rate according to a simple feedback rule responding to output.

$$r_t = \phi y_t$$

We abstract from active monetary policy ($\phi = 0$)[1], and goods market clearing condition $y_t = c_t$,

$$y_t = (1 - \beta\omega) \left(d_t + \bar{E}_t \left[\sum_{k=0}^{\infty} (\beta\omega)^k (y_{t+k} - T_{t+k}) \right] \right)$$

Following Angeletos et al. [2025], I make the following assumption about households' first- and second-order beliefs about a fiscal shock.

Assumption 1. Following a fiscal shock, a fraction $\lambda^{\text{PE}} \in [0, 1]$ of the population learns it immediately and perfectly, while the rest never learn it. Furthermore, everyone believes that the fraction of agents that learn the innovation is $\lambda^{\text{GE}} \in (0, 1]$, where λ^{GE} may or may not coincide with λ^{PE}. Then

$$\begin{aligned}
\bar{E}_t \varepsilon_{t+k} &= \lambda^{\text{PE}} \varepsilon_{t+k} \\
\bar{E}_t \bar{E}_t \varepsilon_t &= \lambda^{\text{GE}} \varepsilon_t \\
\bar{E}_t \bar{E}_{t+k} \varepsilon_{t+k} &= \lambda^{\text{GE}} \lambda^{\text{PE}} \varepsilon_{t+k}
\end{aligned}$$

Intuitively, λ^{PE} captures how accurately households perceive the shock, while λ^{GE} captures how households think others perceive the shock, thereby governing general equilibrium effects. The final condition implies that the period-t expectation of the average period-$t + k$ forecast can be also interpreted as the fraction of informed agents at t (λ^{PE}) times what these agents believe that the fraction of informed individuals will be at $t+h$ (λ^{GE}). The non-FIRE departure applies only to households. Firms and policy authorities retain full information.

Assumption 2. The fiscal shock ε_t follows an AR(1) process.

$$\varepsilon_{t+1} = \rho_\varepsilon \varepsilon_t + u_{t+1}$$

Assume a one-time fiscal innovation u_0, then

$$\bar{E}_t \bar{E}_{t+k} \varepsilon_{t+k} = \lambda^{\text{GE}} \lambda^{\text{PE}} \rho_\varepsilon^k \varepsilon_t$$

Throughout the quantitative exercises, I consider a one-time negative innovation, $u_0 < 0$. Because ε_0 enters the tax rule with a minus sign, this innovation raises current taxes and therefore represents a tax-based fiscal consolidation. Impulse responses are reported for a tax increase equal to 1 percent of GDP in absolute value.

[1] Following the original calibration in Angeletos et al. [2024], I set this parameter to exactly zero. This assumption can easily be relaxed.

With assumption 1 and assumption 2, the solution to this model becomes simple. The model can be solved using a standard guess-and-verify method. I guess that output and debt follow law-of-motion equations of the form:

$$y_t = \psi_d d_t + \psi_\varepsilon \underbrace{\bar{E}_t \varepsilon_t}_{\hat{\varepsilon}_t} \tag{1}$$

$$d_{t+1} = \phi_d d_t + \phi_\varepsilon \underbrace{\bar{E}_t \varepsilon_t}_{\hat{\varepsilon}_t} \tag{2}$$

where $\hat{\varepsilon}_t = \bar{E}_t \varepsilon_t$ denotes the average belief of a shock. Substituting the tax rule into household consumption and the government budget constraint yields the closed-form coefficients:

$$\psi_d = \frac{(1-\beta\omega)(1-\tau_d A_d)}{1-(1-\beta\omega)(1-\tau_y)A_d}, \qquad \psi_\varepsilon = \frac{(1-\beta\omega)\Big(\big[(1-\tau_y)\psi_d - \tau_d\big]A_\varepsilon - (\tau_d - 1)H\Big)}{1-(1-\beta\omega)(1-\tau_y)\lambda^{\mathrm{GE}}H},$$

$$\phi_d = \frac{1}{\beta}(1-\tau_d-\tau_y\psi_d), \qquad \phi_\varepsilon = \frac{1}{\beta}\left(-\tau_y\psi_\varepsilon - \frac{\tau_d - 1}{\lambda^{\mathrm{PE}}}\right).$$

where

$$H = \frac{\lambda^{\mathrm{PE}}}{1-\beta\omega\rho_\varepsilon}, \quad A_d = \frac{1}{1-\beta\omega\phi_d}, \quad A_\varepsilon = \frac{\beta\omega\phi_\varepsilon\lambda^{\mathrm{GE}}}{1-\beta\omega\phi_d}.$$

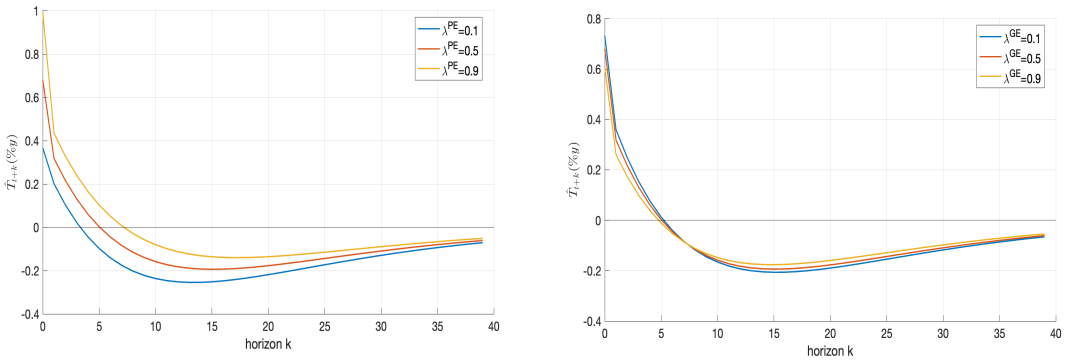
capture the aggregate demand and belief dynamics. Therefore, $\psi_d, \psi_\varepsilon, \phi_d, \phi_\varepsilon$ are solved as a fixed-point to this problem. Plug these into equations 1 and 2, and I can plot the impulse response functions. See Appendix B.1 for detailed derivations.

3.1.2 Results

Figure 3 plots the perceived tax path under alternative λ values in response to a one-time, unexpected tax increase. Rather than decaying according to a simple AR(1), the perceived path is non-monotonic. It combines the perceived shock itself, which decays at rate ρ_ε, with the endogenous response of future taxes to debt: as the tax increase reduces public debt, the debt-feedback rule (τ_d) translates lower debt into expected tax cuts.

Figure 3. Perceived Tax Path Under Alternative λ Values

(a) Alternative λ^{PE} values (b) Alternative λ^{GE} values



Notes: Impulse responses to a one-time tax increase of 1% of GDP. T is reported in percent of GDP, so a value of 1 corresponds to 1% of GDP. When varying λ^{PE}, λ^{GE} is fixed at 0.5, and vice versa.

A higher λ^{PE} raises the perceived tax path at all horizons (panel 3a). It scales up the perceived tax shock on impact and, by implying a larger contraction and a slower perceived decline in

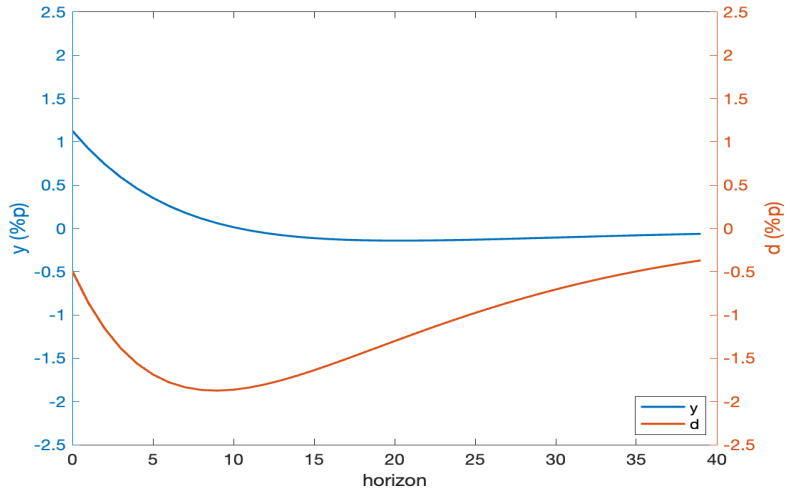
debt, dampens expected future tax cuts in the medium-run through the debt feedback rule. In contrast, λ^{GE} has only a small effect, because it leaves the impact response unchanged and operates solely through expectations of future general equilibrium endogenous outcomes, which are discounted at rate $\beta\omega$ (panel 3b).

Having established how belief parameters shape tax expectations, I conduct a counterfactual analysis to isolate how specific household characteristics alter the transmission of a fiscal shock. I examine two distinct specifications:

1. **Belief Channel**: Comparing a non-FIRE economy to a FIRE economy by varying λ values while holding ω constant.
2. **MPC Channel**: Comparing a high-MPC economy to a low-MPC economy by varying ω values while holding λ constant.

Figure 4 plots the differences in output and debt impulse responses under alternative values of λ, comparing non-FIRE fiscal beliefs against the FIRE benchmark. Non-FIRE beliefs lead to higher output (a milder contraction) and lower debt relative to the FIRE economy. Intuitively, when households perceive the tax shock as smaller than it actually is, aggregate demand adjusts less. Consequently, the tax base erodes less severely, facilitating more effective debt stabilization.

Figure 4. Differences in IRFs of Output and Debt: Non-FIRE vs. FIRE Economies

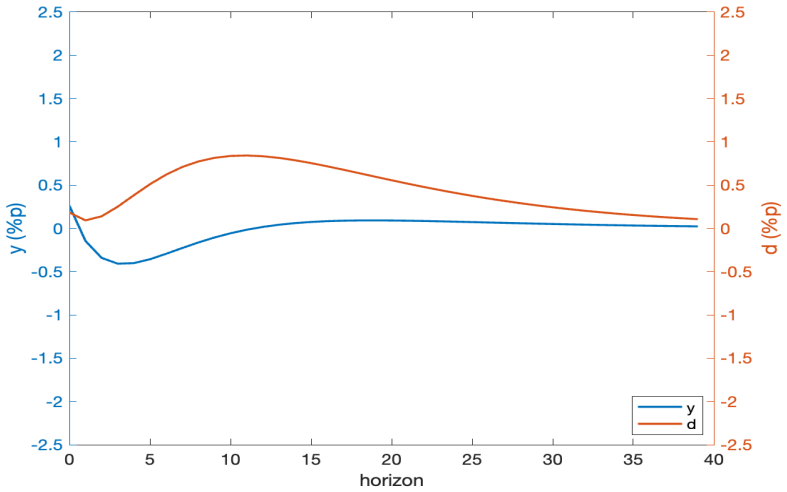


Notes: Differences in impulse responses relative to the FIRE benchmark following a one-time tax increase of 1% of GDP. Output (y) is reported in percentage points (%), and debt (d) is reported in percent of GDP. The cases with $\lambda^{\text{PE}} = \lambda^{\text{GE}} = 0.1$ and 1.0 represent non-FIRE and FIRE economy, respectively.

Figure 5 plots the differences in impulse responses under alternative values of ω, effectively comparing a high-MPC economy to a low-MPC economy. A higher-MPC economy produces a deeper demand contraction and a higher debt path relative to a lower-MPC economy. Because households respond more strongly to tax shocks when MPCs are high, aggregate demand and the endogenous tax base drop significantly, causing government debt to rise relative to the baseline.

This mechanism is corroborated by Table 2, which presents the model solution coefficients from the output and debt equations under alternative λ and ω values. Increasing λ leads to a larger contemporaneous output response (higher $\lambda\psi_\varepsilon$) and a weaker debt-reduction response (lower $\lambda\phi_\varepsilon$) to an actual fiscal shock (ε_t). Intuitively, when agents perceive a larger tax increase on

Figure 5. Differences in IRFs of Output and Debt: High-MPC vs. Low-MPC Economies



Notes: Differences in impulse responses following a one-time tax increase of 1% of GDP. Output (y) is reported in percentage points (%), and debt (d) is reported in percent of GDP. The survival probability is set to $\omega = 0.55$ and 0.865 for comparison.

impact (higher λ^{PE} for $\varepsilon < 0$), contemporaneous demand contracts more sharply, eroding the tax base ($\tau_y y$) and resulting in a smaller reduction in government debt.

Conversely, increasing ω (which lowers the MPC) yields a smaller output response (lower $\lambda^{\text{PE}}\psi_\varepsilon$) and a stronger debt-reduction response (higher $\lambda^{\text{PE}}\phi_\varepsilon$). When agents are less responsive to a tax shock (higher ω), aggregate demand contracts less, preserving the tax base and allowing government debt to decline more effectively.

Table 2. Coefficient Solutions Under Alternative λs and *omega*s

	Output Response to Shock (ψ)			**Debt Response to Shock (ϕ)**		
Param.	ψ_d	ψ_ϵ	$\lambda^{\text{PE}}\psi_\epsilon$	ϕ_d	ϕ_ϵ	$\lambda^{\text{PE}}\phi_\epsilon$
Panel A: Varying λ values						
λ value						
0.10	0.22	0.94	**0.09**	0.91	9.52	**0.95**
1.00	0.22	1.33	**0.13**	0.91	0.54	**0.05**
Panel B: Varying ω values						
ω value						
0.55	0.86	1.98	**0.19**	0.70	0.32	**0.03**
0.865	0.22	0.86	**0.08**	0.91	0.70	**0.07**

Note: ψ_ε and ϕ_ε are the coefficients on the perceived fiscal innovation, $\hat{\varepsilon}_t$, whereas $\psi_\varepsilon\lambda^{\text{PE}}$ and $\phi_\varepsilon\lambda^{\text{PE}}$ are the coefficients on the actual fiscal innovation, ε_t.

The counterfactual analysis establishes that beliefs and MPCs can materially alter the transmission of a fiscal shock. The one-type model, however, cannot address whether aggregate outcomes

depend on the joint distribution of fiscal beliefs, MPCs, and wealth. To study how the distribution of fiscal beliefs across households affects aggregate outcomes, I extend the model to two household types in the following section.

3.2 Two-Type OLG: The Allocation of Fiscal Beliefs

I extend the model to two household types that differ in their survival probabilities, ω_i, and hence in their MPCs and government debt holdings. I also allow the belief parameters, $(\lambda_i^{\text{PE}}, \lambda_i^{\text{GE}})$, to vary across types. This setup provides a framework for analyzing how belief heterogeneity, through its joint distribution with MPCs and wealth, shapes output and debt dynamics.

3.2.1 Model

Assume that each agent type $i \in \{H, L\}$ owns a fraction s_i of government debt:

$$d_t^i = s_i d_t, \qquad \text{with } s_H = s \quad \text{and} \quad s_L = 1 - s.$$

Note that the constants H, A_d, A_ε from Section 3.1 now carry type-specific superscripts $i \in \{H, L\}$:

$$H_i = \frac{\lambda_i^{\text{PE}}}{1 - \beta\omega_i\rho_\varepsilon}, \quad A_d^i = \frac{1}{1 - \beta\omega_i\phi_d}, \quad A_\varepsilon^i = \frac{\beta\omega_i\phi_\varepsilon\lambda_i^{\text{GE}}}{1 - \beta\omega_i\phi_d}, \quad B_d^i = \psi_d A_d^i.$$

Aggregate output is then obtained by summing consumption across household types:

$$\begin{aligned} y_t = \mu_H c_t^H + \mu_L c_t^L &= \sum_{i\in\{H,L\}} \mu_i(1-\beta\omega_i)\Big(s_i d_t + (1-\tau_y)S_y^i - \tau_d S_d^i - (\tau_d - 1)S_\varepsilon^i\Big) \\ &= \psi_d d_t + \psi_\varepsilon \hat{\varepsilon}_t, \end{aligned} \tag{3}$$

where $\hat{\varepsilon}_t \equiv \sum_i \mu_i \lambda_i^{\text{PE}} \varepsilon_t$ denotes the aggregate perceived fiscal shock.

Solving for the closed-form output coefficients yields

$$\psi_d = \frac{\Xi_s - \tau_d \Xi_d}{1 - (1-\tau_y)\Xi_d}, \qquad \psi_\varepsilon = \frac{\big((1-\tau_y)\psi_d - \tau_d\big)\Xi_A - (\tau_d - 1)\Xi_H}{\sum_i \mu_i \lambda_i^{\text{PE}} - (1-\tau_y)\Xi_G},$$

where the weighted aggregation operators are defined as

$$\Xi_s \equiv \sum_{i\in\{H,L\}} \mu_i(1-\beta\omega_i)s_i, \quad \Xi_d \equiv \sum_{i\in\{H,L\}} \mu_i(1-\beta\omega_i)A_d^i, \qquad \Xi_A \equiv \sum_{i\in\{H,L\}} \mu_i(1-\beta\omega_i)H_i A_\varepsilon^i,$$
$$\Xi_H \equiv \sum_{i\in\{H,L\}} \mu_i(1-\beta\omega_i)H_i, \quad \Xi_G \equiv \sum_{i\in\{H,L\}} \mu_i(1-\beta\omega_i)H_i\lambda_i^{\text{GE}}.$$

Finally, aggregating through the government budget constraint gives the debt law-of-motion coefficients:

$$\phi_d = \frac{1 - \tau_d - \tau_y\psi_d}{\beta}, \qquad \phi_\varepsilon = \frac{-(\tau_d - 1) - \tau_y\psi_\varepsilon \sum_i \mu_i \lambda_i^{\text{PE}}}{\beta \sum_i \mu_i \lambda_i^{\text{PE}}}.$$

As in the one-type benchmark, the vector of equilibrium coefficients $(\psi_d, \psi_\varepsilon, \phi_d, \phi_\varepsilon)$ is solved as a fixed point to this system. Substituting these back into equations (1) and (2) allows us to compute the impulse response functions. See Appendix B.2 for detailed derivations.

The calibration of the model parameters is summarized in Table 3. For the preference and policy parameters (β, τ_y, τ_d), I adopt the baseline values from Angeletos et al. [2024]. The persistence of the fiscal shock is set to a high value to clearly isolate the transmission mechanism of a one-time fiscal shock. The population shares of high- and low-MPC agents are calibrated based on summary statistics from the FRBNY SCE. In the survey, individuals with low income constitute approximately 35% of respondents. Accordingly, I set the population shares to $\mu_H = 1/3$ and $\mu_L = 2/3$. Regarding the survival probability parameters, I set ω_L following Angeletos et al. [2024], and calibrate ω_H (the high-MPC group) to match an aggregate average MPC of 0.25, following estimates from the tax-rebate literature (Parker et al. [2013]). Finally, to isolate the role of belief-MPC interaction, I abstract from heterogeneity in government debt holdings in the baseline by setting normalized per-capita holdings equal across the two groups ($s_1 = s_2 = 1$). The results are robust to allowing for unequal debt holdings.

Table 3. Calibration

Preferences & Policy			**Agents**		
Param.	Value	Description	Param.	Value	Description
β	0.99		(ω_H, ω_L)	$(0.65, 0.865)$	Mortality risks
τ_y	0.33	Tax rule	(μ_H, μ_L)	$(0.33, 0.67)$	Share of High/Low-MPC agents
τ_d	0.033	Tax rule	(s_H, s_L)	$(0.1, 0.9)$	Share of government debt holdings
ρ_ϵ	0.80	Deficit shock persistence			
$\frac{D}{Y}$	1.04	Steady-state Debt-to-GDP			

To study how the joint distribution of fiscal beliefs, MPCs, and wealth, shapes output and debt dynamics, I hold population-average beliefs fixed and redistribute beliefs across household types. In Scenario 1, high-MPC households hold beliefs closer to FIRE, whereas low-MPC households, hold more distorted beliefs. Scenario 2 reverses this allocation, assigning more distorted beliefs to the high-MPC type (Table 4). This counterfactual experiment evaluates how the allocation of belief heterogeneity across household types alters aggregate policy effectiveness and debt dynamics.

Table 4. Counterfactual analysis

Scenarios	$\lambda^{\mathbf{H}}$	$\lambda^{\mathbf{L}}$
Benchmark	0.5	0.5
Scenario 1: High-MPC strongly Non-FIRE	0.1	0.7
Scenario 2: High-MPC close to FIRE	0.9	0.3

The goal is to see how belief heterogeneity alters the aggregate outcomes.

3.2.2 Results

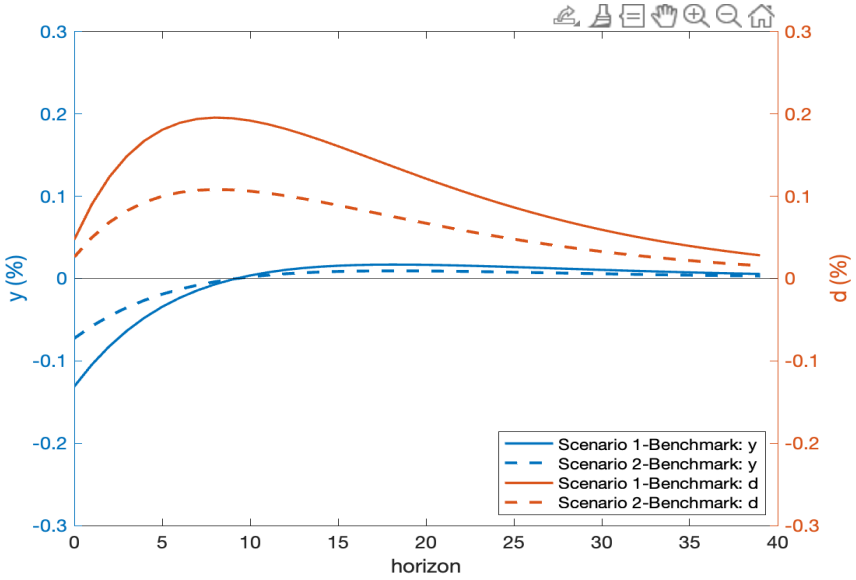
Figure 6 reports the differences in the impulse responses of output and government debt between each heterogeneous-belief scenario and the common-belief benchmark. Throughout this exercise, I impose $\lambda_i^{\text{PE}} = \lambda_i^{\text{GE}} = \lambda_i$ for $i \in \{H, L\}$ and hold the population-weighted average fiscal belief

fixed at

$$\bar{\lambda}^{\mathrm{PE}} = \sum_{i\in\{H,L\}} \mu_i \lambda_i^{\mathrm{PE}}, \quad \bar{\lambda}^{\mathrm{GE}} = \sum_{i\in\{H,L\}} \mu_i \lambda_i^{\mathrm{GE}}.$$

The counterfactuals therefore preserve average fiscal beliefs while changing their distribution across high- and low-MPC households.

Figure 6. Differences in Output and Debt Responses Relative to the Benchmark



Notes: Impulse responses to a one-time tax increase of 1% of GDP. y is reported in percentage points (%). d is reported in percent of GDP, so a value of 1 corresponds to 1% of GDP.

In both heterogeneous-belief scenarios, the debt response lies above that of the common-belief benchmark. Belief heterogeneity therefore weakens fiscal consolidation in both cases. Relative to the common-belief economy, aggregate demand contracts more sharply, the tax base erodes by more, and government debt declines by less.

While both scenarios weaken consolidation, they operate through distinct mechanisms. In Scenario 1, beliefs closer to FIRE are concentrated among low-MPC households. Although low-MPC households adjust consumption less strongly, their large population and wealth shares give them substantial weight in aggregate demand. Consequently, their more accurate perception of tax increase dampens consumption and the tax based, so the debt declines by less.

In Scenario 2, beliefs closer to FIRE are concentrated among high-MPC households. Despite their smaller population and government-debt shares, these households translate a larger fraction of perceived disposable-income changes into current consumption. Their more accurate perception of the tax increase therefore generates a sharper contraction in consumption and the tax base, weakening debt stabilization.

Table 5 confirms this ordering. The coefficients $\psi_\varepsilon \bar{\lambda}^{\mathrm{PE}}$ and $\phi_\varepsilon \bar{\lambda}^{\mathrm{PE}}$ in equations (1) and (2) measure the sensitivity of output and next-period debt to the fiscal shock. Since the tax increase is represented by a negative innovation, $(u_0 < 0)$, a larger value of $\psi_\varepsilon \bar{\lambda}^{\mathrm{PE}}$ implies a sharper output contraction, whereas a larger $\phi_\varepsilon \bar{\lambda}^{\mathrm{PE}}$ corresponds to a larger reduction in government debt.

Relative to the benchmark, the output coefficient rises from 0.590 to 0.661 in Scenario 1 and

to 0.733 in Scenario 2. The debt coefficient, in contrast, falls from 0.780 in the benchmark to 0.757 in Scenario 1 and to 0.733 in Scenario 2. Belief heterogeneity thus deepens the output contraction and weakens debt stabilization in both scenarios.

Table 5. Parameter Combinations and Equilibrium Coefficients ($\bar{\lambda}$ fixed)

	Parameters		Coefficients					
	λ^H	λ^L	ψ_d	ψ_ε	$\psi_\varepsilon \bar{\lambda}^{\text{PE}}$	ϕ_d	ϕ_ε	$\phi_\varepsilon \bar{\lambda}^{\text{PE}}$
Benchmark	0.50	0.50	0.2715	1.1790	**0.5895**	0.8863	1.5605	**0.7803**
Scenario 1	0.10	0.70	0.2715	1.3160	**0.6606**	0.8863	1.5071	**0.7566**
Scenario 2	0.90	0.30	0.2715	1.4718	**0.7329**	0.8863	1.4708	**0.7325**

Note: ψ_ε and ϕ_ε are the coefficients on the perceived fiscal innovation, $\hat{\varepsilon}_t$, whereas $\psi_\varepsilon \sum_i \mu_i \lambda_i^{\text{PE}}$ and $\phi_\varepsilon \sum_i \mu_i \lambda_i^{\text{PE}}$ are the coefficients on the actual fiscal innovation, ε_t.

The counterfactual analysis shows that population-average fiscal beliefs are not a sufficient statistic for fiscal policy outcomes. Holding average beliefs fixed, reallocating beliefs across household types changes aggregate consumption, the tax base, and government debt because the types differ in their MPCs, population shares, and government-debt holdings. The effectiveness of fiscal consolidation therefore depends on the joint distribution of these household characteristics and fiscal beliefs.

4 Quantitative model

While the simple model cleanly illustrates how the joint distribution of fiscal beliefs, MPCs, and government debt holdings shapes policy outcomes, it cannot quantify the importance of belief heterogeneity in a realistically calibrated economy. I next embed group-specific fiscal beliefs into a two-group HANK model. The calibration combines empirical evidence on fiscal beliefs from the Survey of Consumer Expectations with differences in liquid wealth from the Survey of Consumer Finances and labor-income shares from the Current Population Survey. Together, these ingredients allow the model to evaluate fiscal policy effectiveness when fiscal beliefs, MPCs, and household wealth and income are jointly distributed across household groups.

I model fiscal consolidation as a one-percentage point incrase in the common labor income tax rate, which provides a tractable mapping to the SCE measure of expected tax rate changes. I discipline the first-order belief parameters using the ratio of the subjective tax expectation response (from the SCE) to the realized tax rate response following a high-frequency fiscal shock. For second-order beliefs, I adopt a conservative approach by specifying them as a weighted average of first-order beliefs. I subsequently evaluate the robustness of this assumption using prior predictive analysis.

The quantitative exercise compares three belief specifications: full-information rational expectations (FIRE), common mean beliefs, and calibrated heterogeneous beliefs. This comparison separates the level effect and dispersion effect of fiscal beliefs. Moving from FIRE to common beliefs isolates the level effect: relative to FIRE, common beliefs raise the sacrifice ratio. Moving from common to heterogeneous beliefs then isolates the dispersion effect: heterogeneous beliefs

offset approximately one-third of the level effect associated with common beliefs. The sign and approximate magnitude of the offset are robust to an alternative timing of tax perception and to prior-predictive exercises that vary belief and model parameters.

The remainder of the section proceeds as follows. Section 4.1 sets out the environment and, in Section 4.1.5, states how the first- and second-order beliefs of Assumption 1 are embedded in a HANK model. Section 4.2 calibrates the model and belief parameters. Section 4.3 reports impulse responses to a persistent labor-income-tax consolidation under FIRE,under a common-belief counterfactual that holds average beliefs constant, and under calibrated heterogeneous beliefs. Section 4.4 evaluates the effectiveness of fiscal consolidation under different belief specifications using the sacrifice ratio. Section 4.5 assesses sensitivity to the tax-timing assumption and the model parameters.

4.1 Environment

Throughout this section, capital letters denote levels and lower-case letters denote deviations from the steady state (e.g., $y_t = \frac{Y_t - Y^{ss}}{Y^{ss}}$ and $c_t = \frac{C_t - C^{ss}}{C^{ss}}$). Fiscal variables, such as debt and taxes, are expressed as deviations relative to steady-state GDP: $b_t = \frac{B_t - B^{ss}}{Y^{ss}}$ and $\Delta T_t = \frac{T_t - T^{ss}}{Y^{ss}}$. Finally, interest rates and inflation rates are defined as absolute deviations from their steady-state values: $i_t = I_t - I^{ss}$, $\pi_t = \Pi_t - \Pi^{ss}$.

4.1.1 Households

There are two household groups, $g \in \mathcal{G} = \{H, L\}$. These correspond to the groups analyzed in the empirical section: Group H denotes high-MPC (low-income) households with annual income below \$50,000, and Group L collects the remainder. Within each group g, households face idiosyncratic labor productivity z and trade a single liquid asset, government debt. Preferences are given by

$$\overline{E}_{i,0} \sum_{t=0}^{\infty} \beta_g^t \left[u(C_{i,g,t}) - v(N_{i,g,t}) \right], \tag{4}$$

where $u(c) = \frac{c^{1-\sigma^{-1}}}{1-\sigma^{-1}}$ and $v(n) = \frac{\xi n^{1+\psi^{-1}}}{1+\psi^{-1}}$. The flow budget constraint for an optimizing household is

$$C_{i,g,t} + A_{i,g,t+1} = (1 - \tau_{g,t}) \varsigma_g W_t Z_{i,t} N_{i,g,t} + (1 + R_t) A_{i,g,t}, \qquad A_{i,g,t+1} \geq \underline{A}, \tag{5}$$

where $1 + R_t = \frac{(1+I_{t-1})P_{t-1}}{P_t}$ is the gross ex-post real return on government debt held from t to $t-1$, I_{t-1} is the net nominal interest rate set in period $t-1$, and $\Pi_t = P_t/P_{t-1} - 1$. To first order, $r_t = i_{t-1} - \pi_t$. We define group g's per-capita labor income share as $\varsigma_g = s_Y^g/\mu_g$, where $s_Y^g = Y^g/Y$ and μ_g denote group g's income share and population share, respectively. $Z_{i,t}$ follows a finite-state Markov process. The group-specific discount factors β_H and β_L govern asset accumulation of optimizing households. $\tau_{g,t}$ is the group-average labor-income tax rate defined in equation (9), and W_t is the real wage.

Each group also features a share π_g of structural hand-to-mouth (HtM) households, which contributes to cross-group MPC differences. These households hold no assets and consume their current after-tax labor income each period:

$$C_{i,g,t}^{\text{HtM}} = (1 - \tau_{g,t}) \varsigma_g W_t Z_{i,g,t} N_{i,g,t}, \qquad A_{i,g,t}^{\text{HtM}} = A_{i,g,t+1}^{\text{HtM}} = 0.$$

Group-level aggregate choices combine optimizer (O) and hand-to-mouth (HtM) households:

$$X_{g,t} = (1 - \pi_g) X_{g,t}^{\mathrm{O}} + \pi_g X_{g,t}^{\mathrm{HtM}}, \qquad X \in \{C, A\}. \tag{6}$$

4.1.2 Firms and labor unions

Output is linear in effective labor, $Y_t = \Theta_t N_t$. Following Auclert et al. [2024] which is based on Erceg et al. [2000], differentiated labor unions set wages subject to quadratic adjustment costs and employment is equally rationed across households. This ensures that a perceived tax change affects output solely through consumption, the channel through which the belief-wealth and belief–MPC mechanisms operates. Thus, households differ in assets, idiosyncratic productivity, wealth, income, and beliefs, but do not separately bargain over nominal wages.

With flexible goods prices and constant productivity, price inflation equals wage inflation. The linearized wage Phillips curve[2] can be written as

$$\pi_t = \overline{\beta} \pi_{t+1} + \kappa_w \left[\frac{1}{\sigma} \frac{C_t - C^{ss}}{C^{ss}} + \frac{1}{\psi} \frac{Y_t - Y^{ss}}{Y^{ss}} + \widetilde{\tau}_t^{\mathrm{eff}} \right] \tag{7}$$

where $\widetilde{\tau}_t^{\mathrm{eff}} = \sum_g s_g^Y d\tau_{g,t} / (1 - \bar{\tau}_g)$ is the labor-tax wedge.

Firms and unions have full information.

4.1.3 Monetary policy

Policy authorities have full information. Only household beliefs depart from FIRE. The monetary authority follows a Taylor rule that responds to both inflation and the output gap,

$$i_t = \rho_i i_{t-1} + (1 - \rho_i) (\phi_\pi \pi_t + \phi_y y_t) + u_t^i, \qquad \phi_\pi > 1, \quad \phi_y \geq 0. \tag{8}$$

where $y_t \equiv \log(Y_t / Y^{ss})$ is the output gap. I use the steady state Y^{ss} instead of the flexible-wage (natural) output level, Y_t^n. The ex ante real interest rate is $r_t = i_t - \mathbb{E}_t \pi_{t+1}$

4.1.4 Fiscal policy

I model fiscal consolidation as a uniform innovation to each group's labor-income tax rate. While the SCE measures expectations of overall effective tax burdens, labor earnings represent the primary income source and tax base for most households, rendering this specification a tractable and empirically grounded proxy. In this baseline setup, equal percentage-point changes in statutory rates generate differential tax incidence across groups through heterogeneity in effective tax rates, labor-income shares, and labor productivities.

The group-specific effective labor-income tax rate evolves according to

$$\tau_{g,t} = \bar{\tau}_g + x_t + \phi_b \frac{B_{t-1} - B^{ss}}{Y^{ss}}, \qquad g \in \mathcal{G}, \tag{9}$$

where x_t is a common percentage-point innovation to each group's labor-income tax rate. At impact, an innovation of $x_0 = 0.01$ generates direct revenue equal to one percent of quarterly steady-state output, with group g contributing $Y_g^{ss} x_0$ to this direct revenue effect.

[2]See Guerreiro (2023) for derivations. This is based on the optimization problem of labor union.

The fiscal innovation follows

$$x_t = \rho_x x_{t-1} + \varepsilon_t^{\mathcal{T}}, \qquad 0 \leq \rho_x < 1. \tag{10}$$

A positive innovation $\varepsilon_t^{\mathcal{T}} > 0$ is a tax-based fiscal consolidation. As in Section3, it generates a persistent increase in the tax burden.

Aggregate tax revenue is

$$T_t = \sum_{g \in \{L,H\}} \tau_{g,t} Y_{g,t} = \sum_{g \in \{L,H\}} \left(\bar{\tau}_g + x_t + \phi_b \frac{B_{t-1} - B^{ss}}{Y^{ss}} \right) Y_{g,t}$$

Let $b_t \equiv \frac{B_t - B^{\text{ss}}}{Y^{\text{ss}}}$, $\widehat{r}_t \equiv r_t - r^{ss}$, and $\Delta T_t \equiv \frac{T_t - T^{ss}}{Y^{ss}}$. A first-order approximation around the steady state yields

$$\Delta T_t = \sum_{g \in \mathcal{G}} \left[\underbrace{Y_g^{ss} x_t}_{\text{direct revenue}} + \underbrace{\bar{\tau}_g y_{g,t}}_{\text{tax-base feedback}} + \underbrace{\phi_b Y_g^{ss} b_{t-1}}_{\text{debt feedback}} \right]. \tag{11}$$

The government budget constraint is

$$B_t = (1 + R_t) B_{t-1} + G_t - T_t. \tag{12}$$

With government purchases fixed at their steady-state level, its first-order approximation is

$$b_t = (1 + R^{ss}) b_{t-1} + \frac{B^{ss}}{Y^{ss}} r_t - \Delta T_t. \tag{13}$$

Therefore, beliefs affect debt through consumption and output, the endogenous tax base, inflation and the real rate, and the induced debt-feedback taxes.

4.1.5 Beliefs

Beliefs are parameterized as in Assumption 1, with one modification. I allow $(\lambda_g^{\text{PE}}, \lambda_g^{\text{GE}})$ to differ across H and L. This generates belief–wealth and belief–MPC correlations studied in Section 3.2. In addition, the baseline timing is important. All households observe the current tax rate, labor income, and asset return before choosing consumption. Limited perception therefore applies to the continuation of the tax innovation and to its future aggregate consequences. This timing assumption follows the convention in the HANK formulations of Angeletos et al. [2025]. It also guarantees that the household budget and borrowing constraint are satisfied without an additional after-choice tax-settlement mechanism. I relax this assumption and examine a case in which households do not perceive the contemporaneous shock in Section 4.5.

In the OLG model of Section 3, aggregate demand is a closed-form function, allowing λ_g^{GE} to be substituted directly into the analytical solution.That shortcut is unavailable here, because household behavior is summarized by Jacobians of a heterogeneous-agent problem rather than by a single Euler equation [Auclert et al., 2021]. Second-order beliefs instead enter through a *subjective equilibrium*: the economy that an informed household believes it inhabits, and against which it forms the forecasts that drive its saving and consumption plan.

Let

$$\widehat{z}_g = \left(\widehat{y}_g, \widehat{\pi}_g, \widehat{i}_g, \widehat{r}_g, \widehat{b}_g \right)^\top$$

denote the *subjective equilibrium* perceived by an informed household in group g: the aggregate path that this household believes the economy will follow. The subjective equilibrium differs from the actual equilibrium whenever households' beliefs diverge from the economy's actual law of motion.

An informed household in group g believes that a fraction λ_g^{GE} of households in each group incorporates the fiscal shock into its forward-looking plan, while the remaining households respond only to realized current conditions. The parameter λ_g^{GE} therefore governs the extent to which the household takes endogenous general-equilibrium adjustment into account. Because an informed household in group g must forecast the behavior of both groups, the subjective equilibrium paths are determined jointly:

$$\widehat{z}_H = \Phi_H\left(\widehat{z}_H, \widehat{z}_L; \lambda_H^{\text{GE}}\right), \qquad \widehat{z}_L = \Phi_L\left(\widehat{z}_H, \widehat{z}_L; \lambda_L^{\text{GE}}\right). \tag{14}$$

where Φ_g collects the subjective household-response equations, market clearing, the wage Phillips curve, monetary policy, and the government budget constraint. Appendix E gives the full matrix representation and solution of this joint fixed-point problem.

Then the group-g's response aggregates informed and uninformed household's behavior using each group's *first-order* informed share λ_g^{PE}. Responses to realized inputs enter fully, whereas responses to a group's forward-looking plan under its perceived path $\hat{z}_g$ are scaled by its perception parameter λ_g^{PE}. To make this precise, write each household FIRE Jacobian as $J_g^X = \mathcal{E}_g^X + K_g^X$, where $\mathcal{E}_g^X$ captures responses to realized inputs up to the current period, and K_g^X captures responses to expected future paths. Specifically,

$$(\mathcal{E}_g^X)_{ts} = \begin{cases} (J_g^X)_{t-s,0} & \text{for } s \le t, \\ 0 & \text{for } s > t. \end{cases}$$

Informed households respond to both realized and expected path, while uninformed households respond only to realized inputs. Taking the λ_g^{PE}-weighted average of the two, group behavior in sequence-space form is

$$X_g = \mathcal{E}_g^X z + \lambda_g^{\text{PE}} K_g^X \hat{z}_g + \left(\mathcal{E}_g^{X\varepsilon} + \lambda_g^{\text{PE}} K_g^{X\varepsilon}\right)\varepsilon, \qquad X \in \{C, A\} \tag{15}$$

where $z = (y, \pi, i, r, b)^\top$ and aggregates are $X = \sum_g \mu_g X_g$. Thus λ_g^{PE} governs how much a group acts on its forecast, and λ_g^{GE} governs the equilibrium path it forecasts. Setting both weights to one gives FIRE. Appendix E reports the matrix implementation and equilibrium conditions.

4.1.6 General equilibrium

The equilibrium conditions in log-linearized form are:

$$y_t = \sum_g \mu_g c_{g,t} \qquad \text{goods market} \tag{16}$$

$$a_t = \sum_g \mu_g a_{g,t} \qquad \text{asset market} \tag{17}$$

$$\pi_t = \overline{\beta}\pi_{t+1} + \kappa_w \left[\frac{1}{\sigma} C_t + \frac{1}{\psi} Y_t + \widetilde{\tau}_t^{\text{eff}}\right] \qquad \text{wage Phillips curve} \tag{18}$$

$$i_t = \rho_i i_{t-1} + (1-\rho_i)\left(\phi_\pi \pi_t + \phi_y Y_t\right) \qquad \text{Taylor rule} \tag{19}$$

$$r_t = i_{t-1} - \pi_t \qquad \text{Fisher equation} \tag{20}$$

$$b_t = (1+R^{ss}) b_{t-1} + \frac{B^{ss}}{Y^{ss}} r_t - \Delta T_t. \qquad \text{government budget.} \tag{21}$$

Here r_t denotes the realized ex-post real return on nominal government debt carried from period $t-1$ to period t. Accordingly, $r_t = i_{t-1} - \pi_t$ is an accounting relation and does not contain an expectation operator. Under perfect foresight, this realized return coincides ex post with the real return anticipated in period $t-1$.

Household responses take the sequence-space form

$$c_g = \mathcal{E}_g^c \widehat{z} + \lambda_g^{\text{PE}} K_g^c \widehat{z}_g + \left[\mathcal{E}_g^{c\varepsilon} + \lambda_g^{\text{PE}} K_g^{c\varepsilon}\right]\varepsilon \tag{22}$$

$$a_g = \mathcal{E}_g^a \widehat{z} + \lambda_g^{\text{PE}} K_g^a \widehat{z}_g + \left[\mathcal{E}_g^{a\varepsilon} + \lambda_g^{\text{PE}} K_g^{a\varepsilon}\right]\varepsilon \tag{23}$$

where $z = (y, \pi, i, r, b)^\top$, J_g^X is the FIRE Jacobian, $\mathcal{E}_g^X$ captures responses to realized inputs up to the current period, and $K_g^X = J_g^X - \mathcal{E}_g^X$ captures responses to expected future paths for $X \in \{c, a\}$.

4.2 Calibration

4.2.1 Belief parameters

First-order beliefs. I measure the first-order belief parameter using the relative response of subjective and objective tax changes to the same fiscal-news shock. For income group g, define

$$\lambda_g^{\text{PE}} = \frac{\beta_g^{\text{survey}}}{\beta_g^{\text{objective}}} \quad \text{for each } g \in \{H, L\} \tag{24}$$

where β_g^{survey} denotes the response of expected household taxes for group g, and $\beta_g^{\text{objective}}$ denotes the response of the corresponding realized tax measure to the same fiscal shock. If households held FIRE beliefs, their expectations would move one-for-one with objective taxes, yielding $\lambda_g^{\text{PE}} = 1$. A value of λ_g^{PE} below 1 indicates that households only partially incorporate fiscal news into their expectations.

The two responses are estimated separately using the following local projections:

$$\text{objective}_{t+12|t,g} = \alpha_g^{\text{objective}} + \beta_g^{\text{objective}} \varepsilon_t + \Gamma_g^{\text{objective}\prime} W_{t-1} + u_{g,t}^{\text{objective}}, \tag{25}$$

$$\text{survey}_{t+12|t,g} = \alpha_g^{\text{survey}} + \beta_g^{\text{survey}} \varepsilon_t + \Gamma_g^{\text{survey}\prime} W_{t-1} + u_{g,t}^{\text{survey}}, \tag{26}$$

where ε_t is the fiscal shock and W_{t-1} contains predetermined controls. The estimated belief parameter is therefore

$$\widehat{\lambda}_g^{\mathrm{PE}} = \frac{\widehat{\beta}_g^{\mathrm{survey}}}{\widehat{\beta}_g^{\mathrm{objective}}} \quad \text{for each } g \in \{H, L\} \tag{27}$$

Measurement. The subjective tax-change variable, $\mathrm{survey}_{t+12|t,g}$, is constructed from Question 27 (Q27) of the monthly SCE Core Survey, which asks respondents about expected changes in total household tax payments over the next 12 months, holding income fixed. Because income is held fixed, the response is naturally interpreted as an expected change in the household's effective tax rate, rather than an expectation of statutory or marginal tax rates.

The question covers federal, state, and local income taxes, alongside property and sales taxes. I therefore construct an objective measure that aligns as closely as possible with this broad tax concept. Because no single data source tracks quarterly tax-liability changes across income groups, I combine aggregate and micro-level data. Specifically, the National Income and Product Accounts (NIPA) provide the quarterly aggregate effective tax rate, while the Consumer Expenditure Survey (CEX) provides the annual cross-sectional distribution of taxable income and tax liabilities. Combining these sources yields the quarterly group-specific objective tax rate

$$\tau_{g,t} = \kappa_g^{\mathrm{CEX}} \tau_t^{\mathrm{NIPA}}.$$

The objective tax-change variable, $\Delta\tau_{t+12|t,g}^{\mathrm{obj}}$, is then constructed by scaling the realized twelve-month change in the aggregate effective personal-tax rate over the corresponding twelve-month horizon by the group-specific tax-incidence loading, κ_g^{CEX}. Appendix C provides more details about estimating κ_g^{CEX} from CEX. Table 6 summarizes the data sources, variable definitions, and tax coverage used in this construction.

Table 6. Variables and Data Sources

Object	**Data and measure**	**Coverage**
Subjective tax expectation	SCE, Q27	Federal, state, and local income, property, and sales taxes
Objective agg. tax rate	NIPA, `W055RC1Q027SBEA/PI`	Federal, state, and local personal current taxes relative to personal income
Group tax incidence	CEX, IRS	Total income share relative to broad-base tax share that includes federal and state income tax, property tax, estimated sales tax share.

Notes: Federal and state incidence components are combined using their shares in each group's federal-plus-state income-tax liability. AGI-based proxy for state income-tax exposure relative to state gross income-tax liability

The fiscal-news shock, ε_t, is constructed from the daily high-frequency fiscal-news series of Fritsch and Mazlish. [2026]. Daily innovations are summed within each calendar quarter to match the

quarterly aggregation of the SCE responses. The resulting quarterly shock is normalized by its standard deviation over the common estimation sample. One limitation of this shock is that this is a fiscal-deficit news shock rather than a tax-policy shock. It can therefore alter households' expectations of future tax burdens without generating a comparably large change in realized taxes over the same twelve-month horizon. Consistent with this distinction, the estimated objective-tax responses are less precise than the survey responses, as shown in Table 7.

Estimates. Table 7 presents the estimated responses and the implied belief parameters. The point estimates lie below one in all specifications, ranging from 0.115 to 0.685. In the benchmark specification—excluding the Covid period—the estimated belief parameters are 0.310 for the lower-income group and 0.685 for the higher-income group. Both groups underperceive the shock: the high-MPC (lower-income) group incorporates around one-third of the fiscal news, whereas the low-MPC (higher-income) group incorporates around two-thirds. I use these ratios as the benchmark calibration values for the quantitative model.

Table 7. Household Fiscal Beliefs to Tax Changes by Group

Sample	Group	$\widehat{\lambda}_g^{\mathrm{PE}}$ (s.e.)	$\widehat{\beta}_{\mathrm{survey}}$ (s.e.)	$\widehat{\beta}_{\mathrm{objective}}$ (s.e.)	N
Baseline	H ($\leq$ \$50k)	0.115	0.188	1.629	46
		(0.086)	(0.099)	(1.168)	
	L ($>$ \$50k)	0.188	0.158	0.835	46
		(0.211)	(0.104)	(0.623)	
Exclude Covid	H ($\leq$ \$50k)	0.310	0.351	1.131	37
		(0.253)	(0.175)	(1.319)	
	L ($>$ \$50k)	0.685	0.380	0.555	37
		(0.789)	(0.058)	(0.686)	

The precision of $\widehat{\lambda}_g^{\mathrm{PE}}$ is limited by the objective-tax response in its denominator. Although the estimated objective-tax responses $\hat{\beta}_{\mathrm{objective}}$ are positive, none is statistically distinguishable from zero. The weak objective response is consistent with fiscal news affecting expectations before announced policies are implemented and reflected in realized effective tax rates (fiscal foresight, Leeper et al. [2013]). I therefore use the estimated ratios as benchmark calibration values and report the model's implications over alternative values of λ_g^{PE} in a neighborhood of these benchmarks (Table 14).

Second-order beliefs. Second-order beliefs capture how households expect aggregate economy-wide variables to respond to fiscal policy. In principle, they could be disciplined by comparing group-specific revisions in expected inflation or output following fiscal news with the corresponding realized responses. In practice, however, realized macroeconomic outcomes are driven by a range of non-fiscal shocks, so empirical estimates may not provide a clean measure of households' perceived general-equilibrium response to fiscal policy. I therefore do not estimate λ_g^{GE}. Instead, I impose a common second-order belief, assuming that informed households are *correct* about the population-average informed share. Specifically, I set their perceived share equal to

the estimated population average,

$$\lambda_H^{\mathrm{GE}} = \lambda_L^{\mathrm{GE}} = \mu_H \lambda_H^{\mathrm{PE}} + \mu_L \lambda_L^{\mathrm{PE}} = 0.5498. \tag{28}$$

This serves as the natural, agnostic benchmark. Section 4.5 reports the sacrifice-ratio decomposition over the admissible range $\lambda_g^{\mathrm{GE}} \in [0, 1]$. It shows that the quantitative findings are driven by the ordering of the first-order shares $\lambda_H^{\mathrm{PE}} < \lambda_L^{\mathrm{PE}}$, not by the level of λ^{GE}.

4.2.2 Model parameters

Table 8 reports the main quarterly parameters and targets. I keep the group classification from the motivating empirical evidence: low income group and mid/high income group. The population and labor-income shares come from the CPS. Government debt is the only traded asset, with Bss/Yss = 1.04 from Angeletos et al. [2024]. I set the debt-feedback parameter in the tax rule to $\phi_b = 0.12$, the minimum value that ensures equilibrium determinacy at the calibrated steady-state debt level, $B_{ss} = 1.04$. Appendix D provides the corresponding determinacy analysis.

Table 8. Quarterly calibration

Parameter or target	Value	Interpretation
A. Steady state and group composition		
Y^{ss}, B^{ss}/Y^{ss}	1, 1.04	Output normalization; liquid assets
μ_H, s_Y^H, s_H^{liquid}	0.361, 0.0617, 0.0645	H population, income, and asset shares
Aggregate one-quarter MPC	0.25	Targeted rebate MPC
B. Preferences and idiosyncratic risk		
EIS, Frisch	0.5, 0.5	Preference elasticities
ρ_z, σ_z, $\underline{a}$	0.966, 0.50, 0	Earnings process; no borrowing
C. Fiscal block		
r^{ss}, $\bar{\tau}_H$, $\bar{\tau}_L$, G^{ss}/Y^{ss}	0.005, 0.05, 0.18, 0.16	Steady fiscal block
ρ_τ, ϕ_b, ε_0^τ	0.80, 0.12, 0.01	Tax persistence, debt feedback, shock
D. Monetary policy and nominal rigidities		
ρ_i, ϕ_π, ϕ_y	0.70, 1.50, 0.125	Taylor rule
κ_w	0.0062	Wage Phillips curve

Using the 2022 SCF, I compute ℓ_g^{SCF}, the share of households in group g with non-positive liquid assets, as an empirical counterpart to the hand-to-mouth share π_g. The SCF measure of liquid assets includes transaction accounts, certificates of deposit, directly held mutual funds, stocks, and bonds, while excluding retirement accounts, housing, business equity, and other illiquid wealth. It is therefore close to the broad liquid-asset definition in Kaplan et al. [2014].

This liquid-poor share cannot be mapped directly into π_g, however, because some unconstrained households in the model optimally choose to hold zero or near-zero wealth. Setting $\pi_g = \ell_g^{\mathrm{SCF}}$ would therefore treat these optimizers as hand-to-mouth and overstate π_g. Instead, I use the SCF only to discipline the *relative* hand-to-mouth share across groups, and let the aggregate

MPC target pin down its overall level:

$$\pi_g = \varphi \, \ell_g^{\text{SCF}}, \qquad g \in \{H, L\},$$

where the common scaling factor φ preserves the cross-group ratio of liquid-poor shares from SCF. The scaling factor φ and the discount factors β_H and β_L are chosen jointly to match three moments: (i) a population-weighted one-quarter MPC out of transfers of 0.25, following estimates from the tax-rebate literature [Parker et al., 2013], (ii) asset-market clearing at B^{ss}, and (iii) group H's share of aggregate liquid assets in the SCF. The total mass at the borrowing limit, including both hand-to-mouth households and constrained optimizers, is untargeted and is reported below.

Table 9. Stationary household moments

Moment	H	L	Aggregate
Population share	0.3606	0.6394	1.0000
Optimizer discount factor	0.9605	0.9622	
Annualized discount factor, β_g^4	0.8512	0.8573	
Mean liquid assets per capita	0.1861	1.5217	1.0400
Steady consumption per capita	0.1730	1.2324	0.8503
One-quarter transfer MPC	0.2654	0.2413	0.2500
Four-quarter cumulative transfer MPC	0.4408	0.4256	0.4311
Structural HtM share within group	0.0741	0.0441	0.0549
Total zero-asset share within group	0.4645	0.3059	0.3631

Notes: Total zero-asset shares include structural HtM households and optimizing households at the borrowing limit.

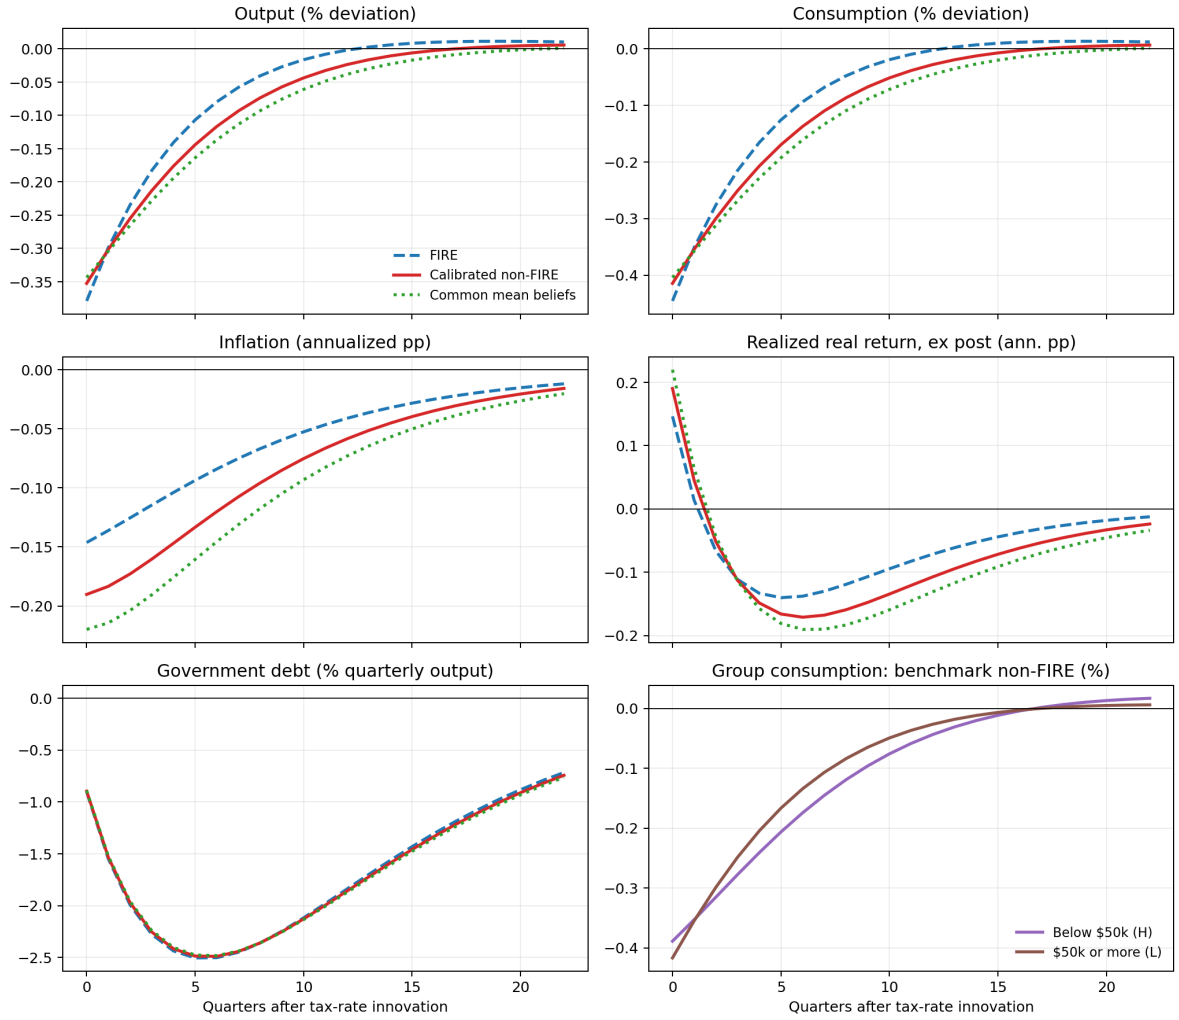
Table 9 reports the resulting stationary household moments. Because group-specific MPCs are not separately targeted, the implied one-quarter MPC difference is modest: the MPC is 0.2654 for group H and 0.2413 for group L. The ranking persists over a four-quarter horizon, with cumulative MPCs of 0.4408 and 0.4256, respectively. This muted MPC heterogeneity reflects the construction of the groups. To align the model with the SCE, households are sorted by income rather than directly by their MPCs. Because income is an imperfect proxy for liquidity and consumption responsiveness, income-defined groups need not exhibit sharply separated MPCs at every horizon. Nevertheless, the below-$50,000 group has the higher MPC at both horizons. The aggregate zero-asset share is 0.3631. This untargeted moment combines structural hand-to-mouth households with optimizing households at the borrowing limit. Its relatively high level partly reflects a familiar feature of one-asset HANK models: when calibrated to match both a high average MPC and a highly concentrated distribution of liquid wealth, these models tend to place a sizable fraction of optimizing households at or near the borrowing constraint [Kaplan and Violante , 2014].

4.3 Quantitative results

Figure 7 compares the responses to a persistent one-percentage-point increase in the labor-income tax rate under three belief specifications: full-information rational expectations (FIRE),

common mean beliefs, and calibrated heterogeneous beliefs. Under the common-belief specification, both groups hold first-order beliefs equal to the population average, $\lambda_H^{\text{PE}} = \lambda_L^{\text{PE}} = 0.525$, while higher-order beliefs remain uniform across groups as in the baseline. The common-belief specification therefore preserves the average beliefs while removing their cross-group dispersion. I organize the discussion around two steps: moving from FIRE to common beliefs isolates the *level* effect of fiscal beliefs. Moving from common to heterogeneous beliefs isolates its *dispersion* effect.

Figure 7. Impulse Responses Under Labor-income-tax Consolidation



Level effect of fiscal beliefs: from FIRE to common beliefs. The level effect dampens the initial output and consumption contraction. On impact, the contraction is smaller under the common-beliefs specification relative to the FIRE benchmark: consumption falls by 0.446% under FIRE but by only 0.404% under common beliefs, while output falls by 0.379% and 0.344%, respectively. This occurs because households underweight the continuation of the future tax increase, leading to a smaller immediate reduction in spending.

Beyond impact, however, the contraction becomes more pronounced under common beliefs than under FIRE. At subsequent horizons, consumption and output under common beliefs fall below their FIRE counterparts. This persistent contraction stems from the same underweighting mech-

anism: households fail to fully anticipate the future relief that follows. Although debt reduction lowers future taxes through the fiscal rule and real interest rates fall, non-FIRE households underweight this future relief, causing their spending to recover more slowly. Consequently, output returns to steady state in quarter 13 under FIRE, compared to quarter 22 under common beliefs.

Because the wage Phillips curve is forward looking and wage setters have full information, the more persistent contraction in demand under common-belief specification translates into a deeper decline in inflation. Impact inflation is -0.146 annualized percentage point under FIRE and -0.220 under common beliefs. Because the policy rule is inertial, the nominal rate initially falls by less than inflation, and the impact real rate rises by more under common beliefs (0.146 versus 0.220). As the policy rule subsequently lowers the nominal rate in response to weaker activity and lower inflation, the real rate turns negative from the second quarter onward. By $h = 11$, it falls to -0.145 under common beliefs, compared with -0.083 under FIRE.

Dispersion effect of fiscal beliefs: from common to heterogeneous beliefs. The heterogeneous-belief economy lies between FIRE and common beliefs for every variable and at every horizon. Belief heterogeneity therefore partially offsets the level effect of fiscal beliefs and moves the economy back toward FIRE.

The group-level responses clarify this result. Group L, the low-MPC group, holds beliefs closer to FIRE, with $\lambda_L^{\text{PE}} = 0.685$, and its consumption path consequently lies closer to the FIRE path. Group H, by contrast, holds more distorted beliefs, with $\lambda_H^{\text{PE}} = 0.303$, and deviates further from FIRE. Because group L accounts for a larger share of aggregate consumption through its higher income and wealth, its near-FIRE response carries greater weight in aggregate-demand and tax-base dynamics and pulls the economy back toward FIRE.

The dispersion effect demonstrates that *who* holds the beliefs matters. It underscores that evaluating the macroeconomic implications of fiscal policy requires accounting for the distribution of fiscal beliefs across households, rather than merely their average. In our quantitative setup, households with near-FIRE beliefs account for a larger share of aggregate demand and the tax base. Consequently, belief heterogeneity mitigates the level effect of fiscal beliefs and pulls aggregate dynamics back toward the FIRE benchmark.

To evaluate the effectiveness of fiscal consolidation, I now turn to debt dynamics.

Debt dynamics. Surprisingly, debt dynamics remain very similar across the three specifications, even though the paths of output and real rates differ markedly. At the end of year 5, debt is only 0.032 percent of quarterly output lower under heterogeneous beliefs and 0.049 percent lower under common beliefs than under FIRE, respectively. (With the inclusion of debt maturity, however, debt dynamics differ across specifications. See Appendix F for details.)

Beliefs affect government debt through three channels: (i) the real rate, which changes interest payments on outstanding debt, (ii) output, which changes the labor-income tax base, and (iii) lagged debt, which combines interest accrued and debt feedback tax rule. Let Δ denote the

deviation from the FIRE path. Differencing the linearized government budget constraint yields

$$\Delta b_t = \underbrace{B^{ss}\Delta r_t}_{\text{interest}} \underbrace{-\bar{\tau}\,\Delta Y_t}_{\text{tax base}} + \underbrace{(1 + r^{ss} - \phi_b)\,\Delta b_{t-1}}_{\text{lagged debt}}, \tag{29}$$

where $1 + r^{ss} - \phi_b = 0.885$. Table 10 reports this decomposition.

Table 10. Debt Dynamics Decomposition

Horizon	Interest $B^{ss}\Delta r_t$	Tax base $-\bar{\tau}\,\Delta Y_t$	Lagged debt $(1 + r^{ss} - \phi_b)\Delta b_{t-1}$	Total Δb_t
Panel A: Heterogeneous beliefs — FIRE				
Impact ($h = 0$)	+0.009	−0.005	0.000	+0.004
End of year 1 ($h = 3$)	−0.004	+0.005	+0.010	+0.011
End of year 3 ($h = 11$)	−0.011	+0.005	−0.017	−0.023
End of year 5 ($h = 19$)	−0.005	+0.002	−0.029	−0.032
Panel B: Common beliefs — FIRE				
Impact ($h = 0$)	+0.014	−0.006	0.000	+0.008
End of year 1 ($h = 3$)	−0.007	+0.008	+0.015	+0.016
End of year 3 ($h = 11$)	−0.016	+0.007	−0.026	−0.035
End of year 5 ($h = 19$)	−0.008	+0.003	−0.044	−0.049

Table 10 decomposes the debt gap into the three components of equation (29). The two flow channels—interest costs and the tax base—move in opposite directions and partially offset each other. On impact, higher real rates raise debt service, while the milder output contraction supports tax revenues. After the first year, these effects reverse: lower real rates reduce debt service, whereas the slower recovery erodes the tax base. The lagged-debt term accounts for most of the debt gap beyond the first year, because it accumulates past flow differences. However, this accumulation is limited. The debt-feedback coefficient of the fiscal rule shrinks any inherited gap by $\phi_b - r^{ss} = 11.5$ percent each quarter, preventing past flow differences from compounding.

Overall, because both flow channels are small and partly offset each other, and because the fiscal rule prevents the residual from compounding, beliefs that substantially alter aggregate activity leave the path of public debt nearly unchanged. Specifically, the debt gap remains below 0.05 percent of quarterly steady-state output through year 5 under both belief specifications. However, in a realistic environment incorporating a government debt maturity structure, debt dynamics differ across specifications (Appendix F).

4.4 Fiscal consolidation effectiveness and output costs

A larger decline in debt need not be a more desirable outcome if it is accompanied by a large and persistent contraction in activity. I therefore evaluate the consolidation along two dimensions: its *effectiveness* in reducing government debt and the output cost at which that reduction is achieved. Let $s \in \{F, C, H\}$ index the FIRE, common-belief, and heterogeneous-belief scenarios. I normalize the size of the initial fiscal innovation as

$$\bar{\varepsilon} \equiv \frac{100\varepsilon_0}{Y^{ss}} = 1 \tag{30}$$

so the innovation equals one percent of quarterly steady-state output. I define debt reduction, cumulative output loss, the fiscal sacrifice ratio, and fiscal effectiveness as

$$\mathcal{D}_s(h) = -\frac{100\, b_h^s / Y^{ss}}{100\, d\mathcal{T}_0^{direct} / Y^{ss}}, \tag{31}$$

$$\mathcal{L}_s(h) = -\frac{\sum_{t=0}^{h} 100\, y_t^s / Y^{ss}}{100\, d\mathcal{T}_0^{direct} / Y^{ss}}, \tag{32}$$

$$\mathcal{S}_s(h) = \frac{\mathcal{L}_s(h)}{\mathcal{D}_s(h)}, \tag{33}$$

$$\mathcal{E}_s(h) = \frac{1}{\mathcal{S}_s(h)} = \frac{\mathcal{D}_s(h)}{\mathcal{L}_s(h)}. \tag{34}$$

The sacrifice ratio $\mathcal{S}_s(h)$ measures the cumulative output loss per unit of debt reduction, capturing the output cost incurred to lower the debt stock. A lower $\mathcal{S}_s(h)$ therefore indicates that a given reduction in the debt stock is obtained at a smaller output cost. Fiscal effectiveness $\mathcal{E}_s(h)$ is its inverse and measures the debt reduction obtained per unit of cumulative output loss. I decompose the difference between the heterogeneous-belief economy and FIRE into a level component, $X_C - X_F$, and a dispersion component, $X_H - X_C$, for $X \in \{\mathcal{S}, \mathcal{E}\}$. Table 11 presents the main quantitative result.

Table 11. Beliefs and the Measures of Fiscal Consolidation

	Scenario			Decomposition of Δ		
Horizon	FIRE	Common	Hetero.	Level (F→C)	Dispersion (C→H)	Disp./ Level
Panel A: Sacrifice ratio $\mathcal{S}_s(h)$ (cumulative output loss per unit debt reduction)						
Impact ($h = 0$)	0.420	0.387	0.394	−0.033	0.007	−0.21
End of year 1 ($h = 3$)	0.481	0.510	0.498	0.028	−0.012	−0.41
End of year 3 ($h = 11$)	0.794	1.011	0.932	0.217	−0.079	−0.37
End of year 5 ($h = 19$)	1.552	2.111	1.903	0.559	−0.208	−0.37
Panel B: Fiscal effectiveness $\mathcal{E}_s(h)$ (debt reduction per unit of cumulative output loss)						
Impact ($h = 0$)	2.381	2.584	2.538	0.203	−0.046	−0.22
End of year 1 ($h = 3$)	2.079	1.961	2.008	−0.118	0.047	−0.40
End of year 3 ($h = 11$)	1.259	0.989	1.073	−0.270	0.084	−0.31
End of year 5 ($h = 19$)	0.644	0.474	0.525	−0.171	0.052	−0.30
Panel C: Cumulative output loss $\mathcal{L}_s(h)$ (pp-quarters)						
Impact ($h = 0$)	0.379	0.344	0.353			
End of year 1 ($h = 3$)	1.095	1.141	1.122			
End of year 3 ($h = 11$)	1.572	2.028	1.859			
End of year 5 ($h = 19$)	1.514	2.165	1.913			
Panel D: Debt reduction $\mathcal{D}_s(h)$						
Impact ($h = 0$)	0.903	0.890	0.896			
End of year 1 ($h = 3$)	2.276	2.240	2.254			
End of year 3 ($h = 11$)	1.980	2.005	1.995			
End of year 5 ($h = 19$)	0.975	1.026	1.005			

Level effect. On impact, the common-belief distortion lowers the sacrifice ratio from 0.420 to 0.387 and raises fiscal effectiveness from 2.281 to 2.584, because non-FIRE fiscal beliefs attenuate the immediate contraction in spending. The ranking reverses within the first year, as the contraction becomes more persistent under common beliefs. By the end of year 5, the sacrifice ratio increases from 1.552 to 2.111, by 0.559. In terms of fiscal effectiveness, the average fiscal beliefs lower the effectiveness measure from 0.644 to 0.474, a decline of about 27 percent. The level effect of fiscal beliefs thus weakens the effectiveness of fiscal consolidation.

Dispersion effect. Belief dispersion partially offsets the level effect for both measures across all horizons. On impact, belief heterogeneity slightly raises the sacrifice ratio (+0.007) and lowers effectiveness (−0.046). From the first year onward, as the output recovery under dispersion accelerates relative to common beliefs (Figure 7), belief dispersion lowers the sacrifice ratio and raises effectiveness, offsetting about one-third of the level effect on both the sacrifice ratio ($0.208/0.559 \simeq 0.37$) and effectiveness ($0.052/0.171 \simeq 0.30$) at $h = 19$. This offset reflects the mechanism in Section 4.3: under heterogeneous beliefs, the households that account for most aggregate spending hold beliefs close to FIRE, which brings the economy closer to FIRE.

Therefore, the effectiveness of fiscal consolidation depends on how fiscal beliefs are distributed across households, not merely on their average. Average beliefs weaken fiscal consolidation effectiveness, but ignoring fiscal belief heterogeneity overstates this weakening. This arises because beliefs closer to full-information rational-expectation are concentrated among the households that account for most aggregate demand. By offsetting roughly one-third of the level effect, belief dispersion implies a lower output cost of consolidation than common-belief models predict.

4.5 Robustness

I examine two key dimensions: the timing of tax perception and parameter values. Both exercises confirm the main findings: the average level of fiscal beliefs weakens fiscal consolidation effectiveness, and belief heterogeneity offsets roughly one-third of this weakening.

4.5.1 An alternative tax timing

In the baseline specification, households observe their current taxes, and misperception is limited to the persistence of the tax shock, following Angeletos et al. [2025]. Under the alternative timing assumption, however, uninformed households fail to correctly perceive the current tax shock. This unperceived tax liability is settled after choices are made: unconstrained optimizers absorb the surprise by reducing end-of-period assets, while constrained optimizers and structural hand-to-mouth households adjust current consumption one-for-one.

Table 12 reports the results. At impact, the settlement timing nearly eliminates differences in sacrifice ratios across belief specifications. The sacrifice ratio is 0.420 under FIRE, 0.421 under common beliefs, and 0.419 under heterogeneous beliefs. Consequently, both the level and dispersion components are economically negligible, rendering their ratio uninformative. Under the settlement timing, households that fail to perceive the current tax liability must nevertheless settle it after making their choices. Those at the borrowing constraint absorb the resulting forecast error through a one-for-one reduction in current consumption. In the calibrated

Table 12. Alternative Tax Timing

Panel A: Sacrifice ratio $\mathcal{S}_s(h)$

		Baseline timing		Settlement timing	
Horizon	FIRE	Common	Hetero.	Common	Hetero.
Impact ($h = 0$)	0.420	0.387	0.394	0.421	0.419
End of year 1 ($h = 3$)	0.481	0.510	0.498	0.568	0.539
End of year 3 ($h = 11$)	0.794	1.011	0.932	1.145	1.021
End of year 6 ($h = 23$)	2.287	3.143	2.826	3.576	3.119

Panel B: Decomposition into level (F→C) and dispersion (C→H)

	Baseline timing			Settlement timing		
Horizon	Level	Dispersion	Disp./Level	Level	Dispersion	Disp./Level
Impact ($h = 0$)	−0.033	+0.007	−0.21	+0.002	−0.002	−1.04
End of year 1 ($h = 3$)	+0.028	−0.012	−0.41	+0.086	−0.029	−0.34
End of year 3 ($h = 11$)	+0.217	−0.079	−0.37	+0.351	−0.124	−0.35
End of year 6 ($h = 23$)	+0.856	−0.317	−0.37	+1.290	−0.457	−0.35

Notes: $\mathcal{S}_s(h) = \mathcal{L}_s(h)/\mathcal{D}_s(h)$, normalized by the impact direct-tax revenue. Level measures $= \mathcal{S}_C - \mathcal{S}_F$ and Dispersion measures $= \mathcal{S}_H - \mathcal{S}_C$. A negative component lowers the sacrifice ratio.

model, this forced adjustment offsets most of the weaker impact consumption response that underperception of the future fiscal path generates under the baseline timing.

At longer horizons, settlement increases both the level component and the absolute magnitude of the dispersion component. At $h = 23$, the level component increases from 0.856 to 1.290, while the dispersion component becomes more negative, moving from -0.317 to -0.457. Settlement raises cumulative output losses because households at the borrowing constraint absorb unperceived liabilities through lower current consumption, whereas unconstrained households reduce assets, which depresses subsequent consumption. Nevertheless, dispersion offsets approximately 35 percent of the level component under settlement, close to the 37 percent offset under the baseline timing. The timing assumption therefore affects the magnitude of the belief effect, but has little effect on the relative importance of its dispersion component.

4.5.2 Prior predictive analysis.

The second exercise addresses the concern that the results reflect weakly identified or conventionally chosen parameters. I place priors over the belief parameters and the aggregate block, solve the model for each draw, and report the induced distribution of the decomposition.

Table 13 lists the priors. The first-order belief parameters are drawn from Beta priors centered on the survey estimates, with standard deviations of about 0.13; λ^{GE} is drawn uniformly on $[0.05, 1]$, essentially its whole admissible range, because it is not separately identified. The debt-feedback coefficient, tax-shock persistence, wage Phillips-curve slope, and Taylor-rule coefficients are drawn over conventional ranges. The prior does not impose $\lambda_H^{\mathrm{PE}} < \lambda_L^{\mathrm{PE}}$, so draws that reverse the estimated ordering of beliefs remain in the sample. This allows us to verify that the qualitative result stems from the empirical estimates rather than being an artifact of the

Table 13. Priors for the Prior Predictive Analysis

Parameter	Description	Prior	Baseline
Belief parameters			
λ_H^{PE}	Informed share, group H	Beta(mean 0.310, ν=12)	0.310
λ_L^{PE}	Informed share, group L	Beta(mean 0.685, ν=12)	0.685
λ^{GE}	Second-order (believed) share	U(0.05, 1.00)	0.525
Fiscal and nominal block			
ϕ_b	Debt feedback in the tax rule	U(0.10, 0.25)	0.12
ρ_τ	Tax-shock persistence	Beta(mean 0.80, ν=20)	0.80
κ_w	Wage Phillips-curve slope	LogNormal(med 0.0062, $\sigma_{\log}$=0.5)	0.0062
ϕ_π	Taylor rule, inflation	U(1.25, 2.50)	1.50
ϕ_y	Taylor rule, output	U(0.00, 0.25)	0.125
ρ_i	Taylor rule, inertia	U(0.50, 0.85)	0.70
Household block (drawn only in the extended exercise)			
$\overline{\mathrm{MPC}}$	One-quarter aggregate rebate MPC	U(0.15, 0.30)	0.25
B^{ss}	Aggregate liquid-asset supply	U(1.04, 2.60)	1.04

Notes: ν is the concentration of the Beta prior. The implied standard deviations are about 0.13 for λ_H^{PE} and λ_L^{PE} and about 0.09 for ρ_τ. With $r^{ss} = 0.005$, the ϕ_b range implies a half-life of the debt gap between about seven quarters at the lower end and two and a half quarters at the upper end, and contains the baseline value of 0.12.

prior specifications. Of 600 draws, all solve and 516 yield a determinate equilibrium. I use the determinate draws throughout.

The baseline-prior column of Table 14 reports the results. Three findings stand out. First, both components of the decomposition have robust signs beyond impact. The level component raises the sacrifice ratio in 81 to 94 percent of draws at the end of years 1, 3, and 6, while the dispersion component lowers it in 91 to 97 percent. The two components have opposite signs in 81 to 91 percent of draws, and the median ratio of dispersion to level is -0.28, -0.35, and -0.36 at these horizons, close to the offset of about 37 percent in the baseline calibration. On impact, the level component instead lowers the sacrifice ratio in most draws, consistent with the attenuation of the impact contraction in Section 4.3.

Second, the sign of the dispersion component is determined by the ordering of beliefs across groups. Among the 499 draws with $\lambda_L^{\mathrm{PE}} > \lambda_H^{\mathrm{PE}}$, the dispersion component is negative at $h = 23$ in every draw. Among the 16 draws that reverse this ordering, it is negative in none. The magnitude of the offset also increases with the belief gap (correlation -0.35). The mechanism is therefore governed by which group is better informed, as the analytical decomposition in Section 3.2 implies.

Third, the qualitative result is not sensitive to λ^{GE}. Its correlation with the dispersion component at $h = 23$ is modest ($+0.17$), and the dispersion component is negative in 96 to 98 percent of draws within each tercile of the λ^{GE} prior. This mitigates the concern that the restriction on λ^{GE} drives the sign of the dispersion effect.

Table 14. Prior Predictive Analysis: Baseline and Extended Priors

	Baseline priors		Extended priors	
Horizon	Level	Dispersion	Level	Dispersion
Share of draws in which the component lowers $\mathcal{S}$				
Impact ($h = 0$)	0.79	0.45	0.37	0.80
End of year 1 ($h = 3$)	0.19	0.91	0.10	0.96
End of year 3 ($h = 11$)	0.08	0.97	0.06	0.97
End of year 6 ($h = 23$)	0.06	0.97	0.09	0.97
Dispersion component at $h = 23$*, by sign of* $\lambda_L^p - \lambda_H^p$				
$\lambda_L^p > \lambda_H^p$	1.00		1.00	
$\lambda_L^p \leq \lambda_H^p$	0.00		0.00	
Correlation of the signed dispersion component at $h = 23$ *with*				
$\lambda_L^p - \lambda_H^p$	-0.35		-0.38	
λ^g	$+0.17$		$+0.15$	
MPC target	—		$+0.08$	
B^{ss}	—		$+0.26$	
Draws used / solved / requested	515/600/600		131/166/200	

Notes: Statistics rely on determinate draws with positive cumulative debt reduction across all three specifications, ensuring a well-defined sacrifice ratio. The second panel reports sample sizes of $n = 499$ and $n = 16$ (baseline) and $n = 127$ and $n = 4$ (extended, excluding 34 draws with unattainable MPC targets). Correlations reflect the relationship between drawn parameters and the dispersion component at $h = 23$. A negative correlation implies a larger offsetting effect. Dashes indicate parameters held fixed under that prior.

One caveat should be stated plainly. The analysis retains only draws that yield a determinate equilibrium, which excludes 14 percent of the baseline draws. In every retained draw, debt reduction is positive at every reported horizon, so the sacrifice ratio is well defined and no further selection is required. These results should therefore be interpreted as robustness patterns within this class of determinate, debt-reducing calibrations, rather than as a probability statement over all admissible economies.

Extending the priors to the household block. The baseline prior predictive exercise holds the aggregate MPC target and the liquid-asset supply fixed, because varying them requires re-solving the stationary household problem for each draw. As a more demanding check, I also draw these two parameters, the one-quarter aggregate rebate MPC over $[0.15, 0.30]$ and B^{ss} over $[1.04, 2.60]$, so that the degree of MPC heterogeneity and the level of household wealth are themselves randomized. Of 200 draws, 166 solve and 131 yield a determinate equilibrium. The 34 failures are parameter combinations for which no rule-of-thumb scale reproduces the drawn MPC target.

The extended-prior column of Table 14 shows that the qualitative conclusions are unchanged. Beyond impact, the level component raises the sacrifice ratio in 90 to 94 percent of draws, whereas the dispersion component lowers it in 96 to 97 percent. The median ratio of dispersion to level is -0.37, -0.37, and -0.35 at the end of years 1, 3, and 6, respectively.

Within the retained draws, the sign of the dispersion component again aligns perfectly with the ordering of beliefs: it is negative in all 127 draws with $\lambda_L^{\text{PE}} > \lambda_H^{\text{PE}}$ and in none of the four draws with the reverse ordering. The dispersion component is only weakly correlated with λ^{GE} (+0.15) and with the drawn MPC target (+0.08), while its correlation with B^{ss} is somewhat larger but still modest (+0.26). This difference appears to reflect the fact that varying the aggregate MPC target changes the MPCs of the two groups relatively similarly, whereas varying B^{ss} has a larger effect on their MPC gap. Thus, asset supply affects the magnitude of the dispersion component more than the aggregate MPC target does, but neither variation overturns its qualitative relationship with the ordering of attention.

Taken together, these exercises demonstrate that the model's central quantitative finding is robust across alternative tax timing assumptions and parameter values: average beliefs weaken fiscal consolidation effectiveness, but ignoring belief heterogeneity overstates this weakening by one-third.

5 Conclusion

This paper studies how the distribution of household fiscal beliefs shapes the effectiveness of tax-based fiscal consolidation. Using the FRBNY Survey of Consumer Expectations, I document that tax-hike expectations vary systematically across income and perceived-exposure groups. Furthermore, I find that differences in income are associated with differences in MPCs. These patterns motivate a framework in which fiscal beliefs are jointly distributed with MPCs, income, and wealth.

The generalized non-FIRE OLG model isolates the underlying mechanism. The one-type model shows that the effect of a tax increase on consumption, the endogenous tax base, and government debt depends on both how accurately households perceive the tax shock (the belief channel) and how strongly they respond to it (the MPC channel). The two-type extension shows that population-average beliefs are not sufficient to characterize these dynamics: holding average beliefs fixed, the joint distribution of fiscal beliefs, MPCs, and wealth shapes output and debt dynamics.

I quantify this mechanism in a two-group HANK model calibrated to cross-group differences in fiscal beliefs, MPCs, labor income, and liquid wealth, and consider a persistent one-percentage-point increase in the labor-income tax rate. Comparing FIRE with a common-belief specification that preserves the population-average belief measures the level effect of fiscal beliefs. Comparing the common-belief specification with calibrated heterogeneous beliefs measures their dispersion effect.

I find that the level effect raises the fiscal sacrifice ratio—cumulative output loss per unit of debt reduction—from 1.55 under FIRE to 2.11 by the end of year five. Belief heterogeneity lowers the ratio to 1.90, offsetting about one-third of this increase. The offset arises because beliefs closer to full information are held by the higher-income group, which accounts for most aggregate demand. The sign and approximate magnitude of the offset are robust to an alternative timing of tax perception, to a maturity structure for government debt, and to prior-predictive exercises

that vary belief, aggregate, and household parameters. In the prior-predictive exercises, the sign of the dispersion effect is determined without exception by which group holds beliefs closer to FIRE.

The paper therefore shows that average fiscal beliefs are not sufficient for evaluating fiscal consolidation; what also matters is who holds the beliefs closer to FIRE. Average belief distortions weaken the effectiveness of consolidation, but a model that ignores their distribution overstates this weakening.

Several limitations remain. First, the group-specific belief parameters are time-invariant. This restriction reflects the short and relatively infrequent survey sample: the SCE Public Policy Survey has been administered three times per year since November 2015. Identifying time variation in households' responsiveness to fiscal news would require a longer and more frequent series of tax expectations together with well-identified tax-specific fiscal shocks. The current data do not provide sufficient variation to distinguish changes in belief formation from changes in the fiscal environment. Second, income is an imperfect proxy for the MPC, and the calibrated MPC gap across income groups is modest. Third, the quantitative model features a single liquid asset and is solved to first order. Future research could combine direct measures of household MPCs with experimentally induced fiscal-information changes and examine the mechanism in a nonlinear two-asset HANK model. These extensions would permit a sharper separation of the response, labor-income, and balance-sheet components of fiscal transmission under heterogeneous beliefs.

References

Alesina, A., Favero, C., & Giavazzi, F. (2015). The output effect of fiscal consolidation plans. *Journal of International Economics*, 96(S1), S19–S42. DOI: 10.1016/j.jinteco.2014.11.003.

Andrade, P., Gautier, E., Mengus, E., & Moench, E. (2025). Household Beliefs about Fiscal Dominance. *Federal Reserve Bank of Boston Research Department Working Paper*, No. 25-2 (March). DOI: 10.29412/res.wp.2025.02.

Angeletos, G.-M., Guerreiro, J., & Zhang, D. R. (2025). From RANK to HANK, without FIRE. *NBER Working Paper*, No. 34596 (December). DOI: 10.3386/w34596.

Angeletos, G.-M., Lian, C., & Wolf, C. K. (2024). Can Deficits Finance Themselves? *Econometrica*, 92(5), 1351–1390.

Auclert, A. (2019). Monetary Policy and the Redistribution Channel. *American Economic Review*, 109(6), 2333–2367.

Auclert, A., Bardóczy, B., Rognlie, M., & Straub, L. (2021). Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models. *Econometrica*, 89(5), 2375–2408.

Auclert, A., Rognlie, M., & Straub, L. (2024). The Intertemporal Keynesian Cross. *Journal of Political Economy*, 132(12), 4068–4121. DOI: 10.1086/732531.

Bianchi-Vimercati, R., Eichenbaum, M. S., & Guerreiro, J. (2024). Fiscal Stimulus with Imperfect Expectations: Spending vs. Tax Policy. *Journal of Economic Theory*, 217, 105814. DOI: 10.1016/j.jet.2024.105814.

Bianchi, F., Dabla-Norris, E., & Khalid, S. (2025). Perceptions of Public Debt and Policy Expectations: Evidence from Cross Country Surveys. *IMF Working Papers*, 2025/197. DOI: 10.5089/9798229025751.001.

Eichenbaum, M. S., Guerreiro, J., & Obradovic, J. (2026). Ricardian Non-Equivalence. *NBER Working Paper*, No. 34691 (January). DOI: 10.3386/w34691.

Erceg, C. J., Henderson, D. W., & Levin, A. T. (2000). Optimal Monetary Policy with Staggered Wage and Price Contracts. *Journal of Monetary Economics*, 46(2), 281–313.

Erceg, C. J., & Lindé, J. (2013). Fiscal consolidation in a currency union: Spending cuts vs. tax hikes. *Journal of Economic Dynamics and Control*, 37(2), 422–445. DOI: 10.1016/j.jedc.2012.09.012.

Farhi, E., & Werning, I. (2019). Monetary Policy, Bounded Rationality, and Incomplete Markets. *American Economic Review*, 109(11), 3887–3928.

Fritsch, G., & Mazlish, J. (2026). High-frequency Fiscal Shocks. *Working Paper*.

Gabaix, X. (2020). A Behavioral New Keynesian Model. *American Economic Review*, 110(8), 2271–2327.

Guerreiro, J. (2023). Belief Disagreement and Business Cycles. *Working Paper*.

Kaplan, G., Violante, G. L., & Weidner, J. (2014). The Wealth Hand-to-Mouth. *NBER Working Paper*, No. 20073 (April). DOI: 10.3386/w20073.

Kaplan, G., & Violante, G. L. (2014). A Model of the Consumption Response to Fiscal Stimulus Payments. *Econometrica*, 82(4), 1199-1239. DOI: 10.3982/ECTA10528.

Kaplan, G., Moll, B., & Violante, G. L. (2018). Monetary Policy According to HANK. *American Economic Review*, 108(3), 697–743.

Koşar, G., & Melcangi, D. (2025). Subjective Uncertainty and the Marginal Propensity to Consume. *Federal Reserve Bank of New York Staff Reports, No. 1148, revised April 2026.*. DOI: 10.59576/sr.1148.

Leeper, E. M., Walker, T. B., and Yang, S.-C. S. (2013). Fiscal Foresight and Information Flows. *Econometrica*, 81(3), 1115–1145.

Parker, J. A., Souleles, N. S., Johnson, D. S., & McClelland, R. (2013). Consumer spending and the economic stimulus payments of 2008. *American Economic Review*, 103(6), 2530–2553. DOI: 10.1257/aer.103.6.2530.

Pfäuti, O., & Seyrich, F. (2026). A Behavioral Heterogeneous Agent New Keynesian Model. *EMPCT Working Paper Series*, No. 2026-01.

Appendix A FRBNY SCE Details

A.1 Survey Sample and Data Treatment

The empirical analysis and the calibration of belief parameters use the Federal Reserve Bank of New York Survey of Consumer Expectations (SCE). The SCE Core Survey is administered monthly and elicits consumers' expectations about a range of household-level and aggregate economic and financial conditions. The analysis combines the Core Survey over 2013–2024 with the Public Policy Survey over 2015–2024. The latter is a special module fielded three times per year and elicits subjective beliefs about prospective changes in government policies, including assistance programs, social insurance, labor-market policies, taxes, and government fees. The Public Policy Survey is administered to respondents drawn from the SCE panel and can therefore be linked to their Core Survey responses.

The empirical evidence uses two questions from the Public Policy Survey. QP1 measures beliefs about the direction of future changes in the average income tax rate, while QP2 measures respondents' perceived personal exposure to the policy change described in the survey.

QP1 asks respondents to assign probabilities to alternative changes in the average income tax rate over the next 12 months:

> *What do you think is the percent chance that over the next 12 months there will be an increase, decrease, or no change in the average income tax rate?*

Respondents report probabilities of an increase, a decrease, and no change, with an additional "no idea" option.

QP2 then asks about the perceived personal effect of the relevant tax change:

> *If the following policy changes were implemented in the next 12 months, how would you and your household be impacted?*

For the average income tax rate, the response categories are "very negatively," "somewhat negatively," "no impact," "somewhat positively," and "very positively." Respondents who assign a higher probability to a tax increase are asked about the personal impact of a tax increase; respondents who assign a higher probability to a tax decrease are asked about the personal impact of a tax decrease. When the two probabilities are equal, respondents are assigned approximately evenly across the two questions.

These two items provide the inputs for the measures of net tax-increase expectations and perceived tax exposure. Table 15 reports how the perceived-exposure sample is constructed. The main source of sample loss is that the QP2 tax-impact question is not available for every respondent-wave. Th number of missing response is 79 observations. The final perceived-exposure sample contains 22,919 respondent-wave observations, or 74.5 percent of the merged sample. Of these observations, 70.1 percent are asked about the effect of a tax increase and 29.9 percent are asked about the effect of a tax decrease. Equal-probability observations account for 25.4 percent of the final perceived-exposure sample. The final perceived-exposure sample in Table 15 is the sample used for the empirical evidence in Section 2.

Table 15. Construction of the Perceived-Exposure Sample

Observations	N	% of merged sample	% of final sample
Merged Core–Public Policy sample	30,769	100.00	
QP2 tax-impact question not recorded	7,771	25.26	
QP2 response missing	79	0.26	
Final perceived-exposure sample	**22,919**	**74.49**	**100.00**
Asked about the personal impact of a tax increase	16,075	52.24	70.14
Asked about the personal impact of a tax decrease	6,844	22.24	29.86
Equal increase and decrease probabilities	5,830	18.95	25.44

Notes: The final perceived-exposure sample includes respondents for whom the QP2 tax-impact question was available and the response was nonmissing. The equal-probability observations overlap the tax-increase and tax-decrease question groups and are therefore not an additional mutually exclusive category.

Winsorization For continuous point forecasts used elsewhere in the analysis, including expected income growth, expected consumption growth, inflation expectations, and expected debt growth, values below −100 percent and identified nonresponse or data-entry values (999, 1000, 1200, and 2025) are first set to missing. The remaining observations are winsorized separately within each survey wave at the survey-weighted first and ninety-ninth percentiles. This wave-specific procedure prevents a small number of typed extreme values from disproportionately affecting the moments of the continuous expectation variables.

No winsorization is applied to the tax probabilities, the net tax-increase expectation measure defined below, or the five-point perceived-exposure measure because these variables are bounded by construction.

A.2 Variable Construction

Net tax-increase expectations Respondents report subjective probabilities that the average income tax rate will increase, remain unchanged, or decrease. Let p_{it}^{inc} and p_{it}^{dec} denote respondent i's reported probabilities, in percentage points, of an increase and a decrease, respectively.

Using p_{it}^{inc} alone would not distinguish between respondents who assign the remaining probability primarily to no change and those who assign it primarily to a tax decrease. For example, two respondents may both assign a 40 percent probability to an increase while assigning very different probabilities to a decrease. The difference $p_{it}^{\text{inc}} - p_{it}^{\text{dec}}$ captures the directional tilt of beliefs toward an increase rather than a decrease. Consequently, the analysis defines this net difference as the benchmark proxy for household tax expectations.

$$B_{it} = p_{it}^{\text{inc}} - p_{it}^{\text{dec}}, \qquad B_{it} \in [-100, 100]. \tag{35}$$

A positive value indicates that the respondent assigns a higher probability to an income-tax-rate increase than to a decrease, whereas a negative value indicates that a decrease is viewed as more likely. A value of zero indicates that the respondent assigns equal probabilities to the two directions.

Table 16. Descriptive Statistics for Net Tax-Increase Expectations

		Observations	Share	Mean net expectation	Std.
Panel A: By perceived exposure					
	Negative	16,220	74.18	27.77	50.65
	Neutral	3,110	17.42	18.14	47.67
	Positive	1,598	8.40	13.86	51.66
Panel B: By income group and perceived exposure					
Low income	Negative	3,932	61.45	30.45	50.91
	Neutral	1,471	28.11	19.60	49.06
	Positive	556	10.44	16.84	52.56
Middle income	Negative	5,916	79.47	28.22	51.58
	Neutral	874	12.05	17.47	45.77
	Positive	580	8.48	15.50	51.52
High income	Negative	6,245	83.29	25.00	49.39
	Neutral	735	10.65	14.45	45.89
	Positive	451	6.06	5.12	48.89

Notes: Weighted shares in Panel A are calculated over the complete-case sample. Weighted shares in Panel B are calculated separately within each income group. Observations are unweighted counts, while means and standard deviations use SCE survey weights. Panel A uses 20,928 observations. Panel B uses 20,760 observations after dropping 168 observations with missing income-group information.

Perceived tax exposure The Public Policy Survey asks respondents how the anticipated change in the average income tax rate would affect them personally. The response is recorded on the following five-point scale:

Value	Response category
1	Very negatively
2	Somewhat negatively
3	No impact
4	Somewhat positively
5	Very positively

The direction referenced in the personal-impact question depends on the respondent's reported probabilities. Respondents who assign a higher probability to an increase than to a decrease are asked about the personal effect of a tax-rate increase. Those who assign a higher probability to a decrease are asked about the personal effect of a tax-rate decrease. When the two probabilities are equal, respondents are assigned approximately evenly between the two question wordings.

To place the responses on a common directional scale, I retain the original response under the tax-increase wording and reverse the response under the tax-decrease wording (**reverse coding**). If R_{it} denotes the original response, the aligned measure is

$$A_{it} = \begin{cases} R_{it}, & \text{under the tax-increase wording,} \\ 6 - R_{it}, & \text{under the tax-decrease wording.} \end{cases} \tag{36}$$

The transformation maps a positive perceived effect of a tax decrease into a negative perceived effect of the corresponding tax increase, and vice versa. Lower values of A_{it} therefore indicate that a tax-rate increase is perceived as personally harmful, whereas higher values indicate that it is perceived as personally beneficial.

For the descriptive figures, I aggregate the aligned measure into three categories:

$$\begin{aligned} \text{Negative exposure}: &\quad A_{it} \in \{1, 2\}, \\ \text{Neutral exposure}: &\quad A_{it} = 3, \\ \text{Positive exposure}: &\quad A_{it} \in \{4, 5\}. \end{aligned} \tag{37}$$

Reverse coding aligns the direction of responses across the tax-increase and tax-decrease questions, but does not imply that the two question versions generate identical response distributions. For this reason, I conduct reverse-coding diagnostics in the following part.

A.2.1 Reverse-Coding Diagnostics

The reverse coding in equation (36) aligns the directional interpretation of personal-impact responses across the two question formats. To evaluate whether this aligned measure provides a consistent basis for qualitative comparisons across income groups, I test for differential income effects between respondents asked about tax increases and those asked about tax decreases. Table 17 reports the results.

Table 17. Diagnostics for the Directionally Aligned Exposure Measure

	Separate samples		Pooled sample
	Increase question	Decrease question	All respondents
	(1)	(2)	(3)
Middle-income group	-0.300^{***}	-0.326^{***}	-0.335^{***}
	(0.026)	(0.036)	(0.036)
High-income group	-0.432^{***}	-0.438^{***}	-0.449^{***}
	(0.026)	(0.035)	(0.036)
Increase-question indicator			-0.344^{***}
			(0.031)
Middle-income group × increase question			0.038
			(0.043)
High-income group × increase question			0.022
			(0.042)
Observations	15,930	6,803	22,733
Joint p-value: income × question			0.664

Notes: All columns report ordered-probit estimates. Columns (1) and (2) use the tax-increase and tax-decrease question samples separately. Column (3) pools the two samples and includes an indicator for the increase-question wording and its interactions with income group. Standard errors clustered by respondent are reported in parentheses. $^{***}p < 0.001$.

Columns (1) and (2) show the same qualitative income ordering under both question wordings. Relative to low-income respondents, middle- and high-income respondents report lower values of the aligned exposure measure, indicating a more negative perceived effect of a tax increase. The estimated coefficients are also similar across the two samples. More formally, column (3) shows that the interactions between income group and the increase-question indicator are small and jointly insignificant ($p = 0.664$). Thus, there is no evidence that the relative income gradient in perceived exposure differs between the two question wordings.

These results support using the aligned measure for qualitative comparisons across income groups. They do not establish that the levels or distributions of responses are identical across the two wordings. Accordingly, I use the measure as a descriptive classification of perceived exposure, not as a framing-invariant measure or a source of causal identification.

Appendix B Equilibrium in OLG Models with Fiscal Beliefs

B.1 One-Type OLG

Output and Debt.

$$
\begin{aligned}
y_t &= (1-\beta\omega)\left(d_t + \bar{E}_t\left[\sum_{k=0}^{\infty}(\beta\omega)^k(y_{t+k}-t_{t+k})\right]\right)\\
&= (1-\beta\omega)\left(d_t + \bar{E}_t\left[\sum_{k=0}^{\infty}(\beta\omega)^k\left\{(1-\tau_y)y_{t+k}-\tau_d d_{t+k}-(\tau_d-1)\varepsilon_{t+k}\right\}\right]\right)\\
&= (1-\beta\omega)\left(d_t + \bar{E}_t\left[\sum_{k=0}^{\infty}(\beta\omega)^k\left\{(1-\tau_y)y_{t+k}-\tau_d d_{t+k}\right\}\right] - (\tau_d-1)\bar{E}_t\left[\sum_{k=0}^{\infty}(\beta\omega)^k\varepsilon_{t+k}\right]\right)
\end{aligned}
$$

- $S_\varepsilon = \sum_{k=0}^{\infty}(\beta\omega)^k\bar{E}_t\varepsilon_{t+k}$

$$
S_\varepsilon = \sum_{k=0}^{\infty}(\beta\omega)^k\bar{E}_t\varepsilon_{t+k} = \sum_{k=0}^{\infty}(\beta\omega)^k\lambda^p\varepsilon_{t+k} = \sum_{k=0}^{\infty}(\beta\omega\rho_\varepsilon)^k\lambda^p\varepsilon_t = \underbrace{\frac{\lambda^p}{1-\beta\omega\rho_\varepsilon}}_{H}\varepsilon_t
$$

- $S_d = \sum_{k=0}^{\infty}(\beta\omega)^k\bar{E}_t d_{t+k}$

$$
\begin{aligned}
\bar{E}d_{t+k+1} &= \phi_d\bar{E}_t d_{t+k} + \phi_\varepsilon\bar{E}_t\hat{\varepsilon}_{t+k}\\
S_d &= \sum_{k=0}^{\infty}(\beta\omega)^k\bar{E}_t d_{t+k} = d_t + \sum_{k=0}^{\infty}(\beta\omega)^{k+1}\bar{E}_t d_{t+k+1}\\
&= d_t + \sum_{k=0}^{\infty}(\beta\omega)^{k+1}(\phi_d\bar{E}_t d_{t+k} + \phi_\varepsilon\bar{E}_t\hat{\varepsilon}_{t+k}) = d_t + \beta\omega\phi_d S_d + \beta\omega\phi_\varepsilon\underbrace{\sum_{k=0}^{\infty}(\beta\omega)^k\bar{E}_t\hat{\varepsilon}_{t+k}}_{\lambda^g S_\varepsilon = \lambda^g H\varepsilon_t}
\end{aligned}
$$

$$
\therefore S_d = \underbrace{\frac{1}{1-\beta\omega\phi_d}}_{A_d}d_t + \underbrace{\frac{\beta\omega\phi_\varepsilon\lambda^g}{1-\beta\omega\phi_d}}_{A_\varepsilon}H\,\varepsilon_t
$$

- $S_y = \sum_{k=0}^{\infty}(\beta\omega)^k \bar{E} y_{t+k}$

$$
\begin{aligned}
y_t &= \psi_d d_t + \psi_\varepsilon \hat{\varepsilon}_t \\
S_y = \sum_{k=0}^{\infty}(\beta\omega)^k \bar{E}_t y_{t+k} &= \psi_d \sum_{k=0}^{\infty}(\beta\omega)^k \bar{E}_t d_{t+k} + \psi_\varepsilon \sum_{k=0}^{\infty}(\beta\omega)^k \bar{E}_t \hat{\varepsilon}_{t+k} \\
&= \psi_d S_d + \psi_\varepsilon \lambda^g S_\varepsilon \\
&= \psi_d (A_d d_t + A_\varepsilon \varepsilon_t) + \psi_\varepsilon \lambda^g H \varepsilon_t \\
&= \underbrace{\psi_d A_d}_{B_d} d_t + \underbrace{(\psi_d A_\varepsilon + \psi_\varepsilon \lambda^g H)}_{B_\varepsilon} \varepsilon_t
\end{aligned}
$$

Combined,

$$
\begin{aligned}
y_t &= (1-\beta\omega)\left(d_t + (1-\tau_y) S_y - \tau_d S_d - (\tau_d - 1) S_\varepsilon\right) \\
&= (1-\beta\omega)\left(d_t + (1-\tau_y)(B_d d_t + B_\varepsilon \varepsilon_t) - \tau_d (A_d d_t + A_\varepsilon \varepsilon_t) - (\tau_d - 1) H \varepsilon_t\right) \\
&= (1-\beta\omega)\left(\underbrace{\left[1 + (1-\tau_y) B_d - \tau_d A_d\right]}_{\text{Coefficient on } d_t} d_t + \underbrace{\left[(1-\tau_y) B_\varepsilon - \tau_d A_\varepsilon - (\tau_d - 1) H\right]}_{\text{Coefficient on} \varepsilon_t} \varepsilon_t\right) \\
&= \psi_d d_t + \psi_\varepsilon \hat{\varepsilon}_t \\
&= \psi_d d_t + \psi_\varepsilon \lambda^p \varepsilon_t
\end{aligned}
$$

Note

$$
\begin{aligned}
d_{t+1} &= \frac{1}{\beta}(d_t - t_t) + \frac{D}{Y}\underbrace{r_t}_{=0} - \frac{D}{Y}\underbrace{(\pi_{t+1} - \mathbb{E}_t[\pi_{t+1}]}_{=0} \\
&= \frac{1}{\beta}\left((1-\tau_d) d_t - \tau_y y_t - (\tau_d - 1)\varepsilon_t\right) \\
&= \frac{1}{\beta}\left((1-\tau_d) d_t - \tau_y (\psi_d d_t + \psi_\varepsilon \hat{\varepsilon}_t) - (\tau_d - 1)\frac{\hat{\varepsilon}_t}{\lambda^p}\right) \\
&= \phi_d d_t + \phi_\varepsilon \hat{\varepsilon}_t \\
&= \phi_d d_t + \phi_\varepsilon \lambda^p \varepsilon_t
\end{aligned}
$$

Therefore,

$$
\begin{aligned}
\phi_d &= \frac{1}{\beta}(1 - \tau_d - \tau_y \psi_d) \\
\phi_\varepsilon &= \frac{1}{\beta}\left(-\tau_y \psi_\varepsilon - \frac{\tau_d - 1}{\lambda^p}\right)
\end{aligned}
$$

Inflation

$$\pi_t = \beta \mathbb{E}_t \pi_{t+1} + \kappa y_t.$$

Following a one-time unanticipated shock, there is no future uncertainty. Thus, under rational expectations, $\mathbb{E}_t \pi_{t+1} = \pi_{t+1}$. Therefore,

$$\pi_0 = \kappa \sum_{t=0}^{\infty} \beta^t y_t = \kappa\left(\psi_d \cdot PV(d) + \psi_\varepsilon \cdot PV(\hat{\varepsilon})\right).$$

- $PV(\hat{\varepsilon})$.

  The perceived shock entering the realized paths satisfies $\hat{\varepsilon}_t = \lambda^p \varepsilon_t$ for all $t \geq 0$, with $\varepsilon_t = \rho^t \varepsilon_0$. Then

$$PV(\hat{\varepsilon}) \equiv \sum_{t=0}^{\infty} \beta^t \hat{\varepsilon}_t = \lambda^p \sum_{t=0}^{\infty} \beta^t \varepsilon_t = \lambda^p \frac{\varepsilon_0}{1-\beta\rho}.$$

- $PV(d)$.

$$PV(d) \equiv \sum_{t=0}^{\infty} \beta^t d_t.$$

  Given the debt law of motion $d_{t+1} = \phi_d d_t + \phi_\varepsilon \hat{\varepsilon}_t$, we have

$$d_{t+1} = \phi_d^{t+1} d_0 + \phi_\varepsilon \sum_{j=0}^{t} \phi_d^{t-j} \hat{\varepsilon}_j = \phi_d^{t+1} d_0 + \phi_\varepsilon \lambda^p \varepsilon_0 \sum_{j=0}^{t} \phi_d^{t-j} \rho^j.$$

  Equivalently, for $t \geq 1$,

$$d_t = \phi_d^t d_0 + \phi_\varepsilon \lambda^p \varepsilon_0 \sum_{j=0}^{t-1} \phi_d^{t-1-j} \rho^j.$$

  Therefore,

$$\therefore \; PV(d) = d_0 + \sum_{t=1}^{\infty} \beta^t d_t = \frac{1}{1-\beta\phi_d} \left[d_0 + \frac{\beta}{1-\beta\rho} \phi_\varepsilon \lambda^p \varepsilon_0 \right].$$

Substituting in,

$$\begin{aligned} \pi_0 &= \kappa \left(\psi_d \, PV(d) + \psi_\varepsilon \, PV(\hat{\varepsilon}) \right) \\ &= \kappa \left(\psi_d \frac{1}{1-\beta\phi_d} \left[d_0 + \frac{\beta}{1-\beta\rho} \phi_\varepsilon \lambda^p \varepsilon_0 \right] + \psi_\varepsilon \frac{\lambda^p \varepsilon_0}{1-\beta\rho} \right). \end{aligned}$$

Since $d_0 = -\frac{D}{Y}\pi_0$, substituting and solving for π_0 gives:

$$\therefore \; \pi_0 = \frac{\kappa \, \varepsilon_0 \, \dfrac{\lambda^p}{1-\beta\rho} \left[\psi_\varepsilon + \psi_d \dfrac{\beta\phi_\varepsilon}{1-\beta\phi_d} \right]}{1 + \kappa \, \psi_d \dfrac{D/Y}{1-\beta\phi_d}}.$$

Given π_0, we can determine the initial debt, $d_0 = -\frac{D}{Y}\pi_0$.

B.2 Two-Type OLG

Output and Debt.

Guess-and-verify:

$$\begin{aligned} y_t &= \psi_d d_t + \psi_\varepsilon \hat{\varepsilon}_t \\ d_{t+1} &= \phi_d d_t + \phi_\varepsilon \hat{\varepsilon}_t \end{aligned}$$

Note in two-type agent economy,

$$\hat{\varepsilon}_t = \bar{E}_t \varepsilon_t = \sum_{i=1}^{2} \mu_i \lambda_i^p \varepsilon_t$$

Then

$$y_t^i = (1-\beta\omega_i)\left(d_t^i + \bar{E}_t^i\left[\sum_{k=0}^{\infty}(\beta\omega_i)^k\left\{(1-\tau_y)y_{t+k} - \tau_d d_{t+k}\right\}\right] - (\tau_d - 1)\bar{E}_t^i\left[\sum_{k=0}^{\infty}(\beta\omega_i)^k \varepsilon_{t+k}\right]\right)$$

Note S_y, S_d, S_t in section 3.1 have i superscripts where

$$\begin{aligned}
S_\varepsilon^i &= \sum_{k=0}^{\infty}(\beta\omega_i)^k E_t\varepsilon_{t+k} = \frac{\lambda_i^p}{1-\beta\omega_i\rho_\varepsilon}\varepsilon_t = H_i\varepsilon_t \\
S_d^i &= \underbrace{\frac{1}{1-\beta\omega_i\phi_d}}_{A_d^i} d_t + \underbrace{\frac{\beta\omega_i\phi_\varepsilon\lambda_i^g}{1-\beta\omega_i\phi_d}}_{A_\varepsilon^i} H_i\varepsilon_t \\
S_y^i &= \psi_d S_d^i + \psi_\varepsilon\lambda_i^g S_\varepsilon^i \\
&= \psi_d(A_d^i d_t + A_\varepsilon^i H_i\varepsilon_t) + \psi_\varepsilon\lambda_i^g H_i\varepsilon_t \\
&= \psi_d A_d^i d_t + \left(\psi_d A_\varepsilon^i + \psi_\varepsilon\lambda_i^g\right) H_i\varepsilon_t
\end{aligned}$$

Substitute in $S_y^i, S_d^i, S_\varepsilon^i$:

$$\begin{aligned}
y_t &= \sum_{i=1}^{2}\mu_i(1-\beta\omega_i)\Big[s_i d_t + (1-\tau_y)\big(\psi_d A_d^i d_t + (\psi_d A_\varepsilon^i + \psi_\varepsilon\lambda_i^g)H_i\varepsilon_t\big) - \tau_d\big(A_d^i d_t + A_\varepsilon^i H_i\varepsilon_t\big) - (\tau_d-1)H_i\varepsilon_t\Big] \\
&= \left\{\sum_{i=1}^{2}\mu_i(1-\beta\omega_i)\Big[s_i + \big((1-\tau_y)\psi_d - \tau_d\big)A_d^i\Big]\right\} d_t \\
&+ \left\{\sum_{i=1}^{2}\mu_i(1-\beta\omega_i)H_i\Big[\big((1-\tau_y)\psi_d - \tau_d\big)A_\varepsilon^i + (1-\tau_y)\psi_\varepsilon\lambda_i^g - (\tau_d-1)\Big]\right\}\varepsilon_t \\
&= \psi_d d_t + \psi_\varepsilon\hat{\varepsilon}_t.
\end{aligned}$$

Coefficient on d_t:

$$\begin{aligned}
\psi_d &= \sum_{i=1}^{2}\mu_i(1-\beta\omega_i)\Big[s_i + (1-\tau_y)\psi_d A_d^i - \tau_d A_d^i\Big] \\
&= \underbrace{\sum_{i=1}^{2}\mu_i(1-\beta\omega_i)s_i}_{\equiv\Xi_s} + \big((1-\tau_y)\psi_d - \tau_d\big)\underbrace{\sum_{i=1}^{2}\mu_i(1-\beta\omega_i)A_d^i}_{\equiv\Xi_d}.
\end{aligned}$$

Equivalently,

$$\psi_d = \frac{\Xi_s - \tau_d\Xi_d}{1-(1-\tau_y)\Xi_d}.$$

Coefficient on ε_t:

$$\begin{aligned}
\psi_\varepsilon \sum_i \mu_i \lambda_i^p &= \sum_{i=1}^{2} \mu_i (1-\beta\omega_i) H_i \Big[\big((1-\tau_y)\psi_d - \tau_d\big) A_\varepsilon^i + (1-\tau_y)\psi_\varepsilon \lambda_i^g - (\tau_d - 1) \Big] \\
&= \big((1-\tau_y)\psi_d - \tau_d\big) \underbrace{\sum_{i=1}^{2} \mu_i (1-\beta\omega_i) H_i A_\varepsilon^i}_{\equiv \Xi_A} + (1-\tau_y)\psi_\varepsilon \underbrace{\sum_{i=1}^{2} \mu_i (1-\beta\omega_i) H_i \lambda_i^g}_{\equiv \Xi_G} \\
&\quad - (\tau_d - 1) \underbrace{\sum_{i=1}^{2} \mu_i (1-\beta\omega_i) H_i}_{\equiv \Xi_H} . \\
&= \big((1-\tau_y)\psi_d - \tau_d\big)\Xi_A + (1-\tau_y)\psi_\varepsilon \Xi_G - (\tau_d - 1)\Xi_H .
\end{aligned}$$

Therefore,

$$\psi_\varepsilon = \frac{\big((1-\tau_y)\psi_d - \tau_d\big)\Xi_A - (\tau_d - 1)\Xi_H}{\sum_i \mu_i \lambda_i^p - (1-\tau_y)\Xi_G}.$$

Finally, from the government budget constraint,

$$\begin{aligned}
d_{t+1} &= \frac{1}{\beta}\Big((1-\tau_d)d_t - \tau_y y_t - (\tau_d - 1)\varepsilon_t\Big) = \frac{1}{\beta}\Big((1-\tau_d)d_t - \tau_y(\psi_d d_t + \psi_\varepsilon \hat{\varepsilon}_t) - (\tau_d - 1)\varepsilon_t\Big) \\
&= \phi_d d_t + \phi_\varepsilon \hat{\varepsilon}_t, \\
\phi_d &= \frac{1 - \tau_d - \tau_y \psi_d}{\beta}, \qquad \phi_\varepsilon = \frac{-(\tau_d - 1) - \tau_y \psi_\varepsilon \sum_i \mu_i \lambda_i^p}{\beta \sum_i \mu_i \lambda_i^p}.
\end{aligned}$$

Note

$$\begin{aligned}
\Xi_s &\equiv \sum_{i=1}^{2} \mu_i (1-\beta\omega_i) s_i, \quad \Xi_d \equiv \sum_{i=1}^{2} \mu_i (1-\beta\omega_i) A_d^i, \quad \Xi_A \equiv \sum_{i=1}^{2} \mu_i (1-\beta\omega_i) H_i A_\varepsilon^i, \\
\Xi_H &\equiv \sum_{i=1}^{2} \mu_i (1-\beta\omega_i) H_i, \quad \Xi_G \equiv \sum_{i=1}^{2} \mu_i (1-\beta\omega_i) H_i \lambda_i^g.
\end{aligned}$$

Inflation

$$\pi_t = \beta \mathbb{E}_t \pi_{t+1} + \kappa y_t.$$

Following a one-time unanticipated shock, there is no future uncertainty. Thus, under rational expectations, $\mathbb{E}_t \pi_{t+1} = \pi_{t+1}$. Iterating forward and imposing the no-bubbles condition yields

$$\pi_0 = \kappa \sum_{t=0}^{\infty} \beta^t y_t = \kappa \left(\psi_d \cdot PV(d) + \psi_\varepsilon \cdot PV(\hat{\varepsilon})\right).$$

- PV($\hat{\varepsilon}$)

  The aggregate perceived shock entering the laws of motion is

$$\hat{\varepsilon}_t = P\, \varepsilon_t, \qquad P \equiv \sum_i \mu_i \lambda_i^p, \qquad \varepsilon_t = \rho^t \varepsilon_0.$$

Then

$$PV(\hat{\varepsilon}) \equiv \sum_{t=0}^{\infty} \beta^t \hat{\varepsilon}_t = P \sum_{t=0}^{\infty} \beta^t \varepsilon_t = P \frac{\varepsilon_0}{1-\beta\rho}.$$

- $PV(d)$

$$PV(d) \equiv \sum_{t=0}^{\infty} \beta^t d_t.$$

Given the debt law of motion

$$d_{t+1} = \phi_d d_t + \phi_\varepsilon \hat{\varepsilon}_t,$$

we have, for $t \geq 1$,

$$d_t = \phi_d^t d_0 + \phi_\varepsilon \sum_{j=0}^{t-1} \phi_d^{t-1-j} \hat{\varepsilon}_j = \phi_d^t d_0 + \phi_\varepsilon P \varepsilon_0 \sum_{j=0}^{t-1} \phi_d^{t-1-j} \rho^j \qquad (t \geq 1).$$

Therefore,

$$PV(d) = d_0 + \sum_{t=1}^{\infty} \beta^t d_t = \frac{1}{1-\beta\phi_d} \left[d_0 + \frac{\beta}{1-\beta\rho} \phi_\varepsilon P \varepsilon_0 \right].$$

Substituting into $\pi_0 = \kappa(\psi_d PV(d) + \psi_\varepsilon PV(\hat{\varepsilon}))$,

$$\pi_0 = \kappa \left(\psi_d \frac{1}{1-\beta\phi_d} \left[d_0 + \frac{\beta}{1-\beta\rho} \phi_\varepsilon P \varepsilon_0 \right] + \psi_\varepsilon \frac{P\varepsilon_0}{1-\beta\rho} \right).$$

Since $d_0 = -\frac{D}{Y}\pi_0$, substituting and solving for π_0 yields

$$\therefore \ \pi_0 = \frac{\kappa \, \varepsilon_0 \, \dfrac{P}{1-\beta\rho} \left[\psi_\varepsilon + \psi_d \dfrac{\beta\phi_\varepsilon}{1-\beta\phi_d} \right]}{1 + \kappa \, \psi_d \dfrac{D/Y}{1-\beta\phi_d}}.$$

Given π_0, we can determine the initial debt, $d_0 = -\frac{D}{Y}\pi_0$.

Appendix C Measuring Group-Specific Tax Changes Using the CEX

SCE asks households about the change in their taxes over the next twelve months while holding income fixed, essentially the effective tax rate, so the objective benchmark must measure a change in the effective tax rate at a monthly frequency and separately for each group. No existing data source provides both. Group-specific effective tax rates, $r_{g,t} = T_{g,t}/Y_{g,t}$, are available only at an annual frequency. The CEX is fielded quarterly, but its income and tax questions refer to the preceding twelve months, and the BLS imputes income-tax liabilities on an annual basis. The CBO series are also released with a lag of several years. The NIPA, in contrast, record tax payments quarterly, but only for the household sector as a whole.

I therefore proceed in two steps. First, I measure the quarterly realized change in the *aggregate* effective tax rate over the survey horizon using the NIPA. Second, I map this aggregate change into group-specific changes using a group-specific loading, κ_g, estimated from the annual CEX cross-section. The second step requires an assumption about how an aggregate change in the tax rate is distributed across groups, which I state and discuss below.

Step 1: Aggregate realized tax change. Let $r_t = \texttt{W055}_t/\texttt{PI}_t$ denote the aggregate effective personal tax rate in quarter t, where W055 is personal current taxes and PI is personal income (BEA, NIPA). For a survey response in quarter t, the realized change over the survey's twelve-month horizon is

$$z_t = 100 \times \left[\log\left(\frac{1}{4} \sum_{j=0}^{3} r_{t+j} \right) - \log r_{t-1} \right],$$

which compares the average effective rate over the four quarters beginning with the survey quarter to the rate in the last quarter before the survey. Because z_t is a log difference, it approximates the percent change in the effective rate.

Step 2: From aggregate to group-specific changes. The aggregate rate is an income-weighted average of group rates,

$$r_t = \sum_g s_{g,t}\, r_{g,t}, \qquad s_{g,t} = Y_{g,t}/Y_t,$$

so a quarterly observation of Δr_t gives one restriction on the G unobserved group-specific changes $\{\Delta r_{g,t}\}_{g=1}^{G}$. The restriction alone cannot pin down how the aggregate change is divided across groups. Identifying the group-specific changes therefore requires additional structure, and I impose the following two assumptions.

(A1) *Equal percentage-point changes:* all groups experience the same change in their effective tax rate, $\Delta r_{g,t} = \Delta r_t$ for all g.

(A2) *Stable composition:* income shares s_g are constant over the sample period, and κ_g is evaluated at sample-average effective tax rates, $\kappa_g = \bar{r}/\bar{r}_g$.

Assumption (A1) is stated in percentage points, whereas z_t measures a proportional (log) change. To express (A1) in the same units as z_t, divide both sides of $\Delta r_{g,t} = \Delta r_t$ by r_g. To a first order,

$$d \log r_g = \frac{dr_g}{r_g} = \frac{dr}{r_g} = \frac{r}{r_g} \cdot \frac{dr}{r} = \kappa_g\, d \log r, \qquad \kappa_g \equiv \frac{r}{r_g}.$$

Applying this to the twelve-month change in Step 1 gives the group-specific realized tax change

$$z_{g,t} = \kappa_g\, z_t. \tag{38}$$

to a first-order. The loading κ_g can be written in terms of observable shares:

$$\kappa_g = \frac{r}{r_g} = \frac{T/Y}{T_g/Y_g} = \frac{Y_g/Y}{T_g/T} = \frac{\text{income share of group } g}{\text{tax share of group } g}. \tag{39}$$

By construction, $\sum_g (T_g/T)\, \kappa_g = 1$. Assumption (A1) implies that a group with a below-average effective tax rate ($\kappa_g > 1$) experiences a larger *proportional* change in its tax rate for a given aggregate change. The alternative assumption of equal proportional changes across groups corresponds to $\kappa_g = 1$ for all g.

Measurement. I estimate κ_g from the BLS Consumer Expenditure Survey (CEX) Interview Survey for 2015–2023, which overlaps with the SCE sample period. Groups are defined as in the main text: group H consists of low-income / high-MPC households whose income is below \$50,000 and group L consists of all remaining households. Tax liabilities include (i) federal and state income taxes, taken from the BLS tax-imputation files (`ntaxi`), which are computed with NBER TAXSIM, (ii) property taxes on owner-occupied dwellings, and (iii) general sales taxes, imputed with the IRS Optional State Sales Tax Tables using each household's state of residence, income, and family size. Income is CEX total income before taxes. All shares use CEX sample weights. I compute κ_g separately for each year and average across years. The `ntaxi` files are not yet available for 2024 and 2025, so I apply the 2015–2023 average to those years, consistent with (A2).

Estimates. Table 18 reports the estimates. Group H accounts for 13.8 percent of income but only 8.0 percent of broad-base tax liabilities, reflecting the progressivity of the tax system. Its loading therefore exceeds one ($\kappa_H = 1.736$), while $\kappa_L = 0.936$ is slightly below one. Under (A1), a given aggregate change in the effective tax rate thus translates into a proportional change for group H that is about 1.85 times as large as for group L. Noe that because κ_g is the average of annual ratios, it need not equal the ratio of the averaged shares shown in the table.

Table 18. Group tax loadings, CEX 2015–2023 average

	H	L
Income share, Y_g/Y	0.138	0.862
Broad-base tax share, T_g/T	0.080	0.920
Tax loading, κ_g	1.736	0.936

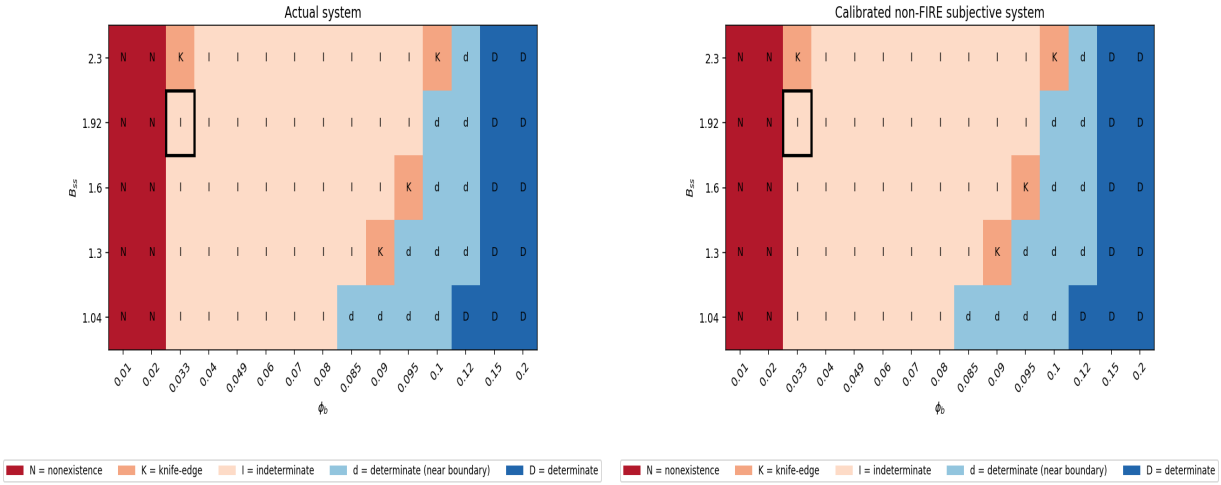
For group H, whose effective tax rate is below average, (A1) therefore implies a larger proportional response than (A1') equal proportional changes ($\kappa_g = 1$).

Caveats. The main caveat is that the two steps use different tax bases. NIPA personal current taxes, which define z_t, cover income and personal property taxes but classify general sales and real-estate taxes as taxes on production and imports, whereas κ_g is computed on the broad CEX base. Because sales taxes are regressive, the broad base raises the tax share of group H and lowers κ_H relative to an income-tax base. In addition, CEX income taxes are TAXSIM-imputed statutory liabilities rather than payments, and the CEX under-represents top incomes.

Appendix D Debt Feedback and Equilibrium Determinacy

This appendix characterizes the determinacy region over combinations of the steady-state debt level, B_{ss}, and ϕ_b, for a target MPC of 0.25 in both the actual and subjective economies. We determine equilibrium existence and determinacy using the winding-number method. In the figures, N denotes nonexistence of an equilibrium, K the boundary between nonexistence and indeterminacy, I indeterminacy, d determinacy with a root close to one, and D determinacy. In both economies, when $B_{ss} = 1.04$, equilibrium is determinate only if $\phi_b \geq 0.12$. We therefore set $\phi_b = 0.12$.

Figure 8. Determinacy Regions in the Actual and Subjective Economies



This calibration implies a half-life of government debt of approximately five to six quarters, which is faster than typical empirical estimates. The relatively strong fiscal feedback is required by a reinforcing debt–inflation mechanism in the model. When government debt rises, the tax rule raises taxes. Higher taxes increase the tax wedge entering the wage Phillips curve, putting upward pressure on inflation. The Taylor rule then raises the nominal interest rate, increasing debt-service costs and thereby further raising government debt. Without sufficiently strong feedback from debt to taxes, this feedback loop can cause debt to diverge. For the target MPC of 0.25 and $B_{ss} = 1.04$, a value of $\phi_b \geq 0.12$ is therefore necessary to ensure a determinate equilibrium.

Appendix E HANK with Belief–MPC Correlation

This appendix derives the subjective and actual equilibrium used in Section 4.1. The logic has two steps. Informed households must first forecast the aggregate path, which requires taking a stand on how many other households are informed. Solving that forecasting problem is the *subjective equilibrium*. Only then can the *actual* equilibrium be computed, by aggregating informed and uninformed behavior with the true informed shares.

Second-order beliefs. Collect the second-order beliefs in a matrix Q,

$$Q = \begin{bmatrix} q_{HH} & q_{HL} \\ q_{LH} & q_{LL} \end{bmatrix},$$

where q_{ij} is the share of group j that group i believes to be informed. Assumption 1 makes this belief a statement about a population-wide fraction rather than about a particular group, so $q_{ij} = \lambda_i^{\mathrm{GE}}$ for both j, and

$$Q = \begin{bmatrix} \lambda_H^{\mathrm{GE}} & \lambda_H^{\mathrm{GE}} \\ \lambda_L^{\mathrm{GE}} & \lambda_L^{\mathrm{GE}} \end{bmatrix}.$$

Household responses. Let $z = (y, r, b, \pi, i)^\top$, let J_h^X denote the FIRE Jacobian,

where $z = (y, \pi, i, r, b)^\top$, J_g^X is the FIRE Jacobian, $\mathcal{E}_h^X$ captures responses to realized inputs up to the current period, and $K_g^X = J_g^X - \mathcal{E}_g^X$ and $K_h^{X\varepsilon} = J_h^{X\varepsilon} - \mathcal{E}_h^{X\varepsilon}$ capture responses to expected future paths for $X \in \{c, a\}$ and ε, respectively. An informed household in group h acts on its forecast $\hat{z}_h$; an uninformed household responds only to realized inputs:

$$\begin{aligned} X_h^I &= \mathcal{E}_h^X \cdot z + K_h^X \cdot \hat{z}_h + J_h^{X\varepsilon} \cdot \varepsilon \\ X_h^U &= \mathcal{E}_h^X \cdot z + \mathcal{E}_h^{X\varepsilon} \cdot \varepsilon \end{aligned}$$

Averaging within the group with the true informed share λ_h^{PE},

$$\begin{aligned} X_h &= \lambda_h^{\mathrm{PE}} \cdot X_h^I + (1 - \lambda_h^{\mathrm{PE}}) \cdot X_h^U \\ &= \mathcal{E}_h^X z + \lambda_h^{\mathrm{PE}} K_h^X \hat{z}_h + \left[\mathcal{E}_h^{X\varepsilon} + \lambda_h^{\mathrm{PE}} K_h^{X\varepsilon}\right]\varepsilon. \end{aligned}$$

Fixed-point expectations of the informed. An informed household in group h must forecast the behavior of *both* groups. It knows the information structure and the Jacobians, so it can compute what an informed household in group g would do. The only thing it does not know is *who* is informed. Formally this is the assumption $E_h^I[\hat{z}_g] = \hat{z}_g$. The belief structure is common knowledge, so group g's forecast is not itself uncertain from h's point of view. Under it,

$$\begin{aligned} E_h^I[X_g^U] &= \mathcal{E}_g^X \hat{z}_h + \mathcal{E}_g^{X\varepsilon}\varepsilon \\ E_h^I[X_g^I] &= \mathcal{E}_g^X \hat{z}_h + K_g^X \hat{z}_g + J_g^{X\varepsilon}\varepsilon \end{aligned}$$

Group h's informed households expect a fraction q_{hg} of group g to be informed, so

$$\begin{aligned} E_h^I[X_g] &= q_{hg} E_h^I[X_g^I] + (1 - q_{hg}) E_h^I[X_g^U] \\ &= \mathcal{E}_g^X \hat{z}_h + q_{hg} K_g^X \hat{z}_g + \left[\mathcal{E}_g^{X\varepsilon} + q_{hg} K_g^{X\varepsilon}\right]\varepsilon, \end{aligned}$$

and, aggregating across groups,

$$\widehat{X}_h = \sum_{g \in \{H,L\}} \mu_g \left\{ \mathcal{E}_g^X \hat{z}_h + q_{hg} K_g^X \hat{z}_g + \left[\mathcal{E}_g^{X\varepsilon} + q_{hg} K_g^{X\varepsilon}\right]\varepsilon \right\}.$$

Because $\widehat{X}_h$ depends on both $\hat{z}_h$ and $\hat{z}_g$, the two groups' forecasts must be solved jointly:

$$\widehat{X}_h = \widehat{X}_h(\hat{z}_H, \hat{z}_L), \qquad \widehat{X}_L = \widehat{X}_L(\hat{z}_H, \hat{z}_L).$$

Subjective general equilibrium. Observer h closes its perceived economy with the same equilibrium conditions as the true economy. Let L and F denote the lag and lead operators, let $\bar{\tau}$ be the steady-state aggregate tax rate, and let $\varpi \equiv \sum_g \mu_g \varsigma_g \omega_g$ be the direct-revenue

coefficient. Then $\hat{z}_h$ solves

$$\begin{aligned}
&\widehat{X}_h = \sum_{g\in\{H,L\}} \mu_g \left\{\mathcal{E}_g^X \hat{z}_h + q_{hg} K_g^X \hat{z}_g + \left[\mathcal{E}_g^{X\varepsilon} + q_{hg} K_g^{X\varepsilon}\right]\varepsilon\right\} && X \in \{c, a\}\\
&\widehat{a}_h = \widehat{b}_h && \text{asset market}\\
&\hat{\pi}_h = \overline{\beta} F\hat{\pi}_h + \kappa_w\left[\tfrac{1}{\text{EIS}}\widehat{c}_h + \tfrac{1}{\text{Frisch}}\widehat{y}_h + \widetilde{\tau}_h^{\text{eff}}\right] && \text{wage Phillips curve}\\
&\hat{i}_h = \rho_i L\hat{i}_h + (1-\rho_i)\left(\phi_\pi\hat{\pi}_h + \phi_y\widehat{y}_h\right) && \text{Taylor rule}\\
&\hat{r}_h = L\hat{i}_h - \hat{\pi}_h && \text{Fisher equation}\\
&\left[I - (1+R^{ss}-\phi_b)L\right]\hat{b}_h + \bar{\tau}\widehat{y}_h - \frac{B^{ss}}{Y^{ss}}\hat{r}_h = -\varpi\varepsilon && \text{government budget}\\
&x = (1-\rho_x L)^{-1}\varepsilon && \text{tax shock}
\end{aligned}$$

Goods-market clearing $\widehat{y}_h = \widehat{c}_h$ then holds by Walras's law and is not imposed separately. Stacking the two observers, the system is linear in $(\hat{z}_H, \hat{z}_L)$ and depends on the second-order beliefs only through Q:

$$A(Q)\begin{bmatrix}\hat{z}_H\\ \hat{z}_L\end{bmatrix} = R(Q)\varepsilon \qquad \Longrightarrow \qquad \begin{bmatrix}\hat{z}_H\\ \hat{z}_L\end{bmatrix} = A(Q)^{-1}R(Q)\varepsilon.$$

Setting Q to a matrix of ones recovers the FIRE path, which is the nesting check reported in the text.

Actual equilibrium. With the perceived paths in hand, the actual equilibrium aggregates informed and uninformed behavior using the *true* informed shares λ_g^p rather than the believed shares Q:

$$X = \sum_{g\in\{H,L\}} \mu_g \cdot X_g = \sum_g \mu_g \left\{\mathcal{E}_g^X z + \lambda_g^{\text{PE}} K_g^X \hat{z}_g + \left[\mathcal{E}_g^{X\varepsilon} + \lambda_g^{\text{PE}} K_g^{X\varepsilon}\right]\varepsilon\right\},$$

together with the same wage Phillips curve, asset-market clearing, Taylor rule, Fisher equation, and government budget constraint, now evaluated at the realized path z. Writing the resulting system as

$$Az = R\varepsilon + B_p\begin{bmatrix}\hat{z}_H\\ \hat{z}_L\end{bmatrix} \qquad \Longrightarrow \qquad z = A^{-1}\left\{R + B_p A(Q)^{-1}R(Q)\right\}\varepsilon,$$

which makes the two-stage structure explicit: $A(Q)^{-1}R(Q)$ is the forecasting stage governed by second-order beliefs, and A^{-1} is the equilibrium stage, with the true informed shares entering through B_p.

Appendix F Government Debt Maturity Structure

This section incorporates a maturity structure for government debt. The main findings remain unchanged: average fiscal beliefs increase the output cost of fiscal consolidation, whereas fiscal belief heterogeneity offsets this cost by one-third. These results are also robust to alternative tax timing specifications and prior predictive analysis. Section F.1 outlines the long-term-debt environment, Section F.2 presents the implied output and debt dynamics, and Section F.3 evaluates fiscal-consolidation effectiveness.

F.1 Environment: Fiscal Policy

To model the maturity structure of government debt, I assume that debt consists of nominal bonds whose coupons decay geometrically at rate ρ. Let B_t denote the number of bonds outstanding at the end of period t and Q_t their nominal price, so that the real market value of debt is $\mathcal{B}_t \equiv Q_t B_t / P_t$. Setting $\rho = 0$ recovers one-period debt.

Bond price. No arbitrage between the long-term bond and a one-period nominal bond with rate i_t requires

$$Q_t(1 + i_t) = 1 + \rho\, Q_{t+1}, \tag{40}$$

so that $Q_t = \sum_{k\geq 0} \rho^k \prod_{j=0}^{k} \frac{1}{1+i_{t+j}}$. In the steady state, $Q = 1/(1 + i^{ss} - \rho)$, or equivalently $(1 + i^{ss})Q = 1 + \rho Q$. Let q_t denote the log deviation of Q_t, let i_t and π_t denote deviations from the steady state, and define

$$\omega \equiv \frac{\rho Q}{1 + \rho Q} = \frac{\rho}{1 + i^{ss}}.$$

Log-linearizing (40) yields $q_t = -i_t + \omega\, q_{t+1}$, and solving forward,

$$q_t = -\sum_{k\geq 0} \omega^k\, i_{t+k}. \tag{41}$$

The log bond price moves inversely with the entire expected path of short-term rates, with weights that decline geometrically at rate ω.

Budget constraint. The government pays the coupon on its outstanding bonds and finances the deficit by issuing new bonds:

$$Q_t B_t = (1 + \rho Q_t) B_{t-1} + P_t(G - T_t).$$

Dividing by P_t and multiplying and dividing the first term by $Q_{t-1}P_{t-1}$ expresses the constraint in terms of the real market value of debt:

$$\mathcal{B}_t = (1 + R^b_t)\, \mathcal{B}_{t-1} + G - T_t, \qquad 1 + R^b_t \equiv \frac{1 + \rho Q_t}{Q_{t-1}(1 + \Pi_t)}, \tag{42}$$

where R^b_t is the ex-post real holding return on government debt between $t-1$ and t: a bond bought at Q_{t-1} pays a coupon of one and leaves ρ bonds worth ρQ_t. Let $b_t \equiv (\mathcal{B}_t - \mathcal{B}^{ss})/Y^{ss}$ and $\Delta T_t \equiv (T_t - T^{ss})/Y^{ss}$. To first order, (42) becomes

$$b_t = (1 + r^{ss})\, b_{t-1} + \frac{\mathcal{B}^{ss}}{Y^{ss}}\, r^b_t - \Delta T_t, \tag{43}$$

which has the same form as with one-period debt, except that the return on debt now reflects changes in the bond price:

$$r^b_t = \omega\, q_t - q_{t-1} - \pi_t. \tag{44}$$

Realized returns. After the shock in period 0, the path is deterministic. For $t \geq 1$, the price q_{t-1} already reflects the realized path, so (41) implies $q_{t-1} = -i_{t-1} + \omega q_t$. Substituting into (44),

$$r^b_t = i_{t-1} - \pi_t, \qquad t \geq 1,$$

which is the same return as on one-period nominal debt. In period 0, the pre-shock price is at its steady state, $q_{-1} = 0$, so

$$r_0^b = \omega\, q_0 - \pi_0 = -\pi_0 - \sum_{k\geq 0} \omega^{k+1}\, i_k.$$

Long-term debt thus introduces an additional revaluation channel on impact. In addition to the revaluation through inflation, $-\pi_0$, the bond price term ωq_0 revalues outstanding debt in response to the change in the entire expected path of nominal rates. When the shock lowers that path, the market value of outstanding debt rises, offsetting part of the debt reduction.

Calibration of the maturity structure. I calibrate ρ to match the empirical maturity profile of marketable U.S. Treasury debt. Using security-level data from the Monthly Statement of the Public Debt (MSPD) from January 2001 through March 2026, I calculate the monthly weighted average maturity (WAM) of marketable securities, weighting each issue by its outstanding face value. The sample-average WAM is 64.5 months (21.5 quarters). Under a geometric decay structure, the undiscounted average maturity of the payment stream equals $1/(1-\rho)$, yielding $\rho = 1 - 1/21.5 = 0.954$ at a quarterly frequency. Accounting for discounting yields the effective decay factor $\omega = \rho/(1+i^{ss}) \approx 0.949$. This corresponds to a Macaulay duration, which measures the present-value-weighted maturity of the debt stream, of $1/(1-\omega) \approx 19.5$ quarters. I use this Macaulay duration as the empirical target to calibrate ω, which directly governs the model's debt-revaluation dynamics.

F.2 Quantitative results

Figure 9 compares the responses to a persistent one-percentage-point increase in the labor-income tax rate under three belief specifications: full-information rational expectations (FIRE), common mean beliefs, and calibrated heterogeneous beliefs. Introducing long-term government debt leaves the outcome rankings across belief specifications unchanged. Relative to FIRE, common beliefs produce a smaller contraction on impact but a more persistent recovery, while heterogeneous beliefs move the economy back toward FIRE. The mechanism is therefore the same as in the baseline model: beliefs closer to FIRE are concentrated among the households that account for most aggregate spending.

Debt dynamics. Debt maturity matters primarily through the revaluation of outstanding debt on impact. The tax increase lowers the expected path of nominal interest rates, raising bond prices and offsetting part of the debt reduction generated by higher tax revenue. This revaluation is larger under non-FIRE beliefs because their more persistent disinflation induces a larger decline in the expected interest-rate path. Impact debt consequently falls by 0.412 percent of quarterly output under FIRE, compared with 0.252 under heterogeneous beliefs and 0.152 under common beliefs. The revaluation occurs only on impact, and the fiscal feedback rule subsequently closes the resulting debt gaps. By the end of year 3, debt reduction differs by less than 3 percent across the three specifications(Table 19).

Figure 9. Impulse Responses Under Labor-income-tax Consolidation

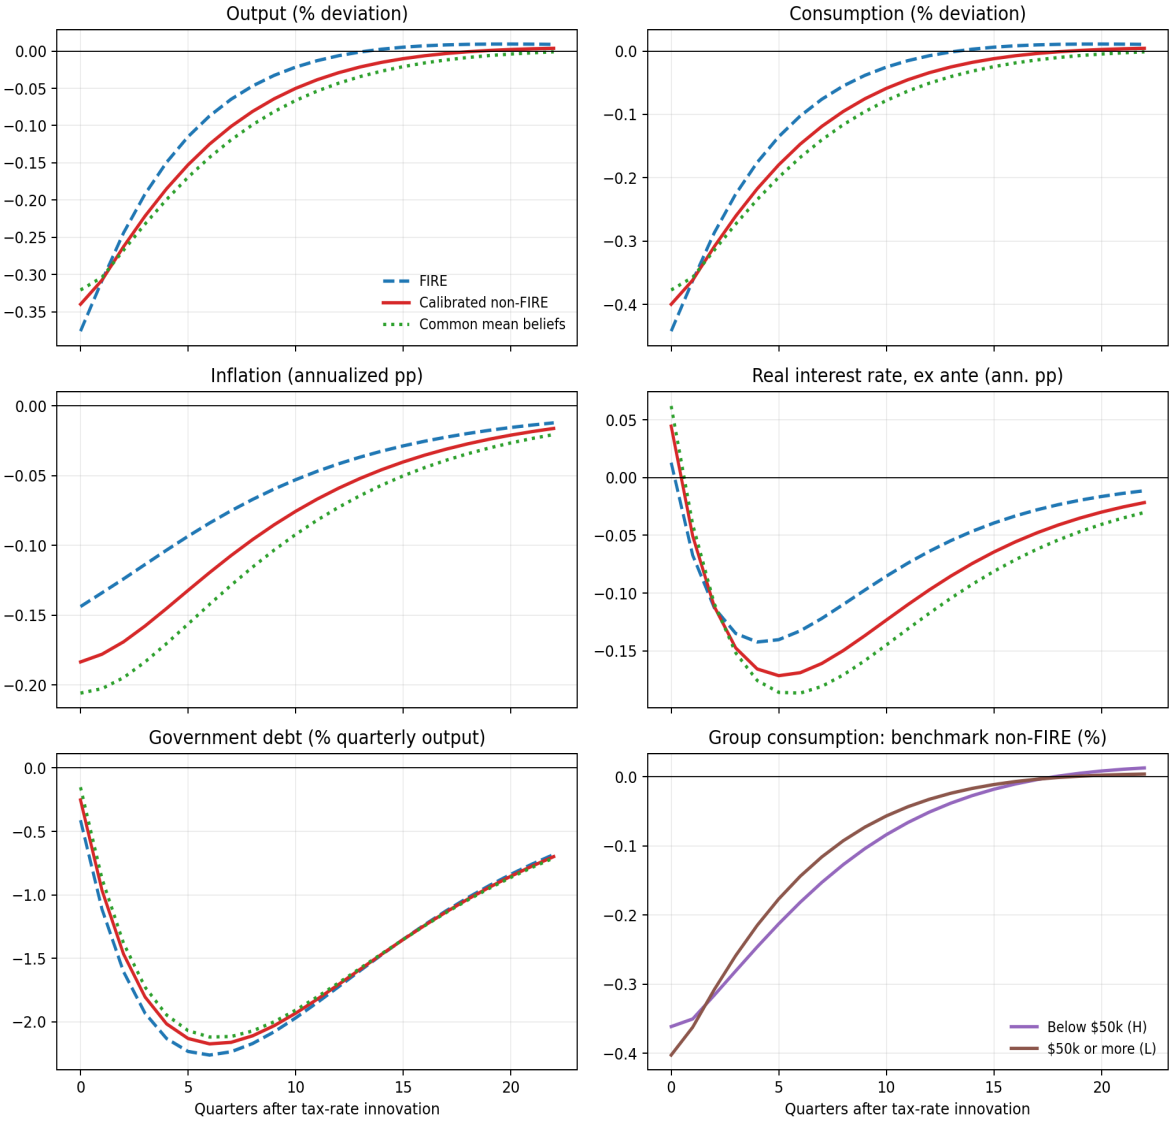


Table 19. Impact Debt Decomposition across Belief Specifications

$$b_0 = \underbrace{\frac{\mathcal{B}^{ss}}{Y^{ss}}\,\omega q_0}_{\text{Bond price}} - \underbrace{\frac{\mathcal{B}^{ss}}{Y^{ss}}\,\pi_0}_{\text{Inflation}} - \Big(\underbrace{\bar{\tau}\, y_0}_{\text{Tax base}} + \underbrace{x_0}_{\text{Direct}} \Big)$$

	Impact Debt	Revaluation Terms		Tax Revenue Terms	
Specification	b_0	Bond Price	Inflation	Tax Base	Direct
FIRE	−0.412	0.493	0.037	−0.058	1.000
Hetero.	−0.252	0.647	0.048	−0.053	1.000
Common	−0.152	0.745	0.054	−0.050	1.000

Notes: Entries are expressed in percent of quarterly steady-state output (Y^{ss}). The table presents the initial debt decomposition according to the equation above, where $b_0 = \text{Bond Price} - \text{Inflation} - (\text{Tax Base} + \text{Direct})$. The debt-feedback term is zero on impact ($t = 0$).

F.3 Fiscal consolidation effectiveness and output costs

Table 20 presents the main quantitative result. Beyond the first year, debt reduction differs by less than 3 percent across specifications (Panel D), so differences in the sacrifice ratio and in fiscal effectiveness are driven almost entirely by differences in cumulative output loss (Panel C). On impact and in the first year, by contrast, the revaluation of long-term debt makes debt reduction smaller under non-FIRE beliefs.

Table 20. Beliefs and the Measures of Fiscal Consolidation

	Scenario			Decomposition of Δ		
Horizon	FIRE	Common	Hetero.	Level (F→C)	Dispersion (C→H)	Disp./ Level
Panel A: Sacrifice ratio $\mathcal{S}_s(h)$ (cumulative output loss per unit debt reduction)						
Impact ($h = 0$)	0.913	2.111	1.347	1.197	−0.763	−0.64
End of year 1 ($h = 3$)	0.580	0.651	0.627	0.071	−0.024	−0.34
End of year 3 ($h = 11$)	0.892	1.139	1.058	0.247	−0.081	−0.33
End of year 5 ($h = 19$)	1.744	2.343	2.142	0.599	−0.201	−0.34
Panel B: Fiscal effectiveness $\mathcal{E}_s(h)$ (debt reduction per unit of cumulative output loss)						
Impact ($h = 0$)	1.095	0.474	0.742	−0.621	0.268	−0.43
End of year 1 ($h = 3$)	1.725	1.536	1.594	−0.189	0.059	−0.31
End of year 3 ($h = 11$)	1.120	0.878	0.945	−0.243	0.067	−0.28
End of year 5 ($h = 19$)	0.573	0.427	0.467	−0.147	0.040	−0.27
Panel C: Cumulative output loss $\mathcal{L}_s(h)$ (pp-quarters)						
Impact ($h = 0$)	0.376	0.321	0.340			
End of year 1 ($h = 3$)	1.120	1.124	1.131			
End of year 3 ($h = 11$)	1.650	2.055	1.928			
End of year 5 ($h = 19$)	1.616	2.222	2.014			
Panel D: Debt reduction $\mathcal{D}_s(h)$						
Impact ($h = 0$)	0.412	0.152	0.252			
End of year 1 ($h = 3$)	1.932	1.726	1.804			
End of year 3 ($h = 11$)	1.849	1.804	1.822			
End of year 5 ($h = 19$)	0.927	0.949	0.940			

On impact, long-term-debt revaluation magnifies the differences in sacrifice ratios across belief specifications: the ratio is 0.913 under FIRE, 2.111 under common beliefs, and 1.347 under heterogeneous beliefs. As discussed above, these differences primarily reflect the effect of bond-price revaluation on impact debt reduction. At longer horizons, the central quantitative result is essentially unchanged. By the end of year 5, the sacrifice ratio rises from 1.744 under FIRE to 2.343 under common beliefs, a level effect of 0.599. Introducing heterogeneous beliefs lowers the ratio to 2.142, a dispersion effect of −0.201, thereby offsetting 34 percent of the level effect. The corresponding fiscal-effectiveness measure delivers the same ranking. Thus, a realistic maturity structure changes the short-run debt response but not the paper's central conclusion: average belief distortions raise the output cost of fiscal consolidation, while their distribution across households offsets roughly one-third of this increase.