

Household Income Inequality and Optimal Trend Inflation ^{*}

Jong Soo Kim [†] Daeha Cho [‡] Kwang Hwan Kim [§]

January 30, 2026

Abstract

This paper studies the optimal inflation target in a tractable heterogeneous agent New Keynesian (THANK) model with an occasionally binding zero lower bound, incorporating income inequality, income risk, and countercyclical fiscal policy. We analytically characterize how cyclical and long-run inequality, income risk, and fiscal policy shape the welfare trade-off of trend inflation across households. Greater cyclical inequality and income risk increase the marginal benefits of inflation, especially for hand-to-mouth (HtM) households, while long-run inequality reduces marginal benefits. Quantitatively, income inequality and income risk modestly raise the utilitarian optimal inflation rate relative to the representative agent benchmark. We further show that countercyclical fiscal policies lower the marginal benefit of inflation for all households and reduce the optimal inflation target. Finally, adding wealth inequality in the form of private or government debt, in addition to income inequality, raises the optimal inflation target.

Key Words: Zero lower bound, Welfare, Optimal inflation target, Income inequality

JEL Classification: E31, E32, E52

^{*}We thank the editor Matteo Iacoviello and two anonymous referees for their insightful comments and suggestions. We also thank the participants at the Yonsei Macroeconomics Workshop, the Spring 2025 Midwest Macroeconomics Meeting, and ESWC 2025. This research was supported by the BK21 FOUR program funded by the Ministry of Education (MOE, Korea) and the National Research Foundation of Korea (NRF), as well as by the Hanyang University research fund (HY-202400000003791).

[†]School of Economics, Yonsei University. Email: jskim0915@yonsei.ac.kr

[‡]College of Economics and Finance, Hanyang University. Email: daehac@hanyang.ac.kr

[§]School of Economics, Yonsei University. Email: kimkh01@yonsei.ac.kr

1 Introduction

Many central banks in advanced countries set an inflation target, and reaching it in the long run is one of their most important goals. While policymakers and academics debate what the optimal inflation target should be, there is an understanding that a higher inflation target provides central banks more room to accommodate macroeconomic shocks by moving the range of policy rates further away from the zero lower bound (ZLB). Much of the debate relies on the trade-off between the marginal benefit of raising the inflation target and its marginal cost, often overlooking the role of household inequality. Figure 1 displays the evolution of labor income inequality in the U.S. Not only has income inequality exhibited a long-term upward trend, but it also moves countercyclically, raising the question of whether inequality should play a more central role in the conduct of monetary policy.¹ While a growing body of literature explores the implications of household inequality for optimal monetary policy, most of these studies focus on the optimal short-run policy rate, that is, the period-by-period interest rate that maximizes household welfare, rather than on the optimal long-run inflation rate.²

The goal of this paper is to answer how the inflation target should be set in an economy with an occasionally binding ZLB under a given Taylor rule, taking household income inequality into account.³ Unlike in a representative agent setting, the answer would depend on how the cost-benefit trade-off of trend inflation differs across income levels. The extent to which this cost-benefit trade-off differs across households depends on cyclical inequality, long-run (steady-state) inequality, income risk, and the degree of countercyclical fiscal policy, all of which have been regarded as important features in determining aggregate dynamics in heterogeneous agent models.⁴

We provide an analytical characterization of the relationship between these elements—cyclical income inequality, long-run income inequality, idiosyncratic income risk, and countercyclical fiscal policy—and the cost-benefit trade-off of trend inflation. In doing so, we use a tractable heterogeneous agent New Keynesian (THANK) model that builds on the model developed by Bilbiie (2024), which provides tractability in isolating the role of each element through simple parameterization. Cyclical inequality is modeled as the sensitivity of individual income to aggregate output, while long-run inequality is modeled by the steady-state income difference between household

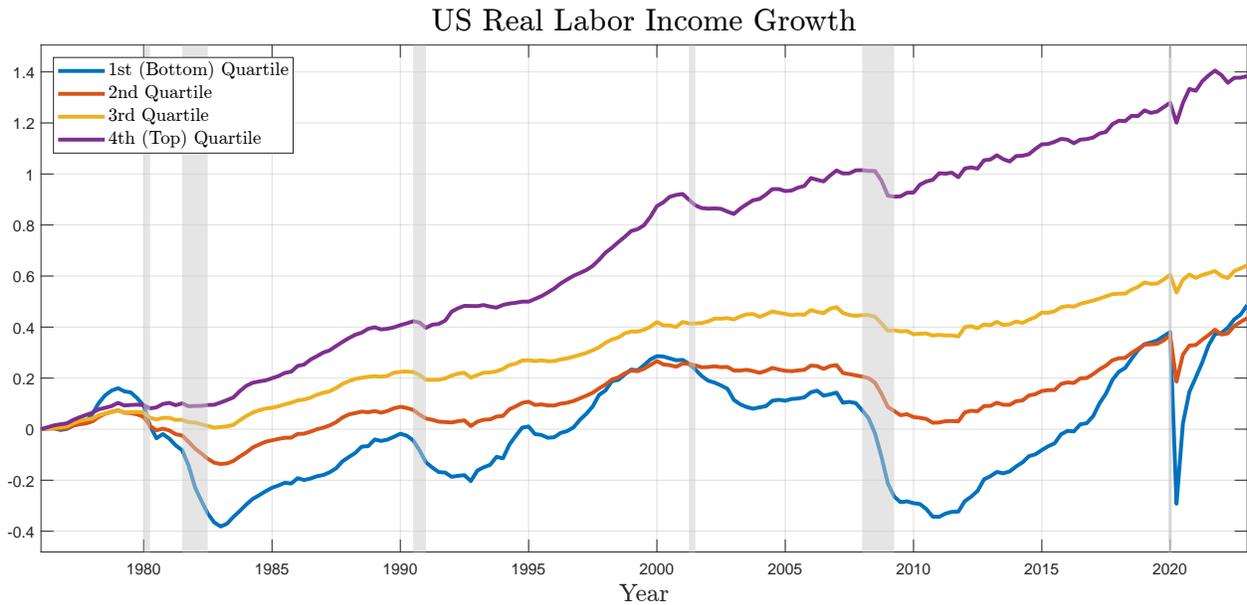
¹The dynamics of labor income growth across income percentiles are from Realtime Inequality, constructed by Thomas Blanchet, Emmanuel Saez, and Gabriel Zucman.

²See Bhandari et al. (2021), Davila and Schaab (2022), and McKay and Wolf (2023) for recent contributions on optimal short-run monetary policy.

³We keep the Taylor rule fixed and vary only the steady-state inflation rate. Hence, our results characterize the welfare-maximizing inflation rate under a given policy rule, rather than the Ramsey-optimal steady-state inflation rate.

⁴See Bilbiie (2020) and Auclert (2019) for the implication of cyclical inequality for the propagation of shocks, Fernández-Villaverde et al. (2025) for the interplay between long-run income inequality and macro volatility, McKay, Nakamura and Steinsson (2016) for the interplay between income risk and the efficacy of forward guidance, McKay and Reis (2016) and Hagedorn, Manovskii and Mitman (2019) for the stabilization effect of countercyclical fiscal policy, among many others.

Figure 1: Evolution of Income Inequality



Notes. The figure displays the dynamics of labor income growth across income percentiles (source: Realtime Inequality, constructed by Thomas Blanchet, Emmanuel Saez, and Gabriel Zucman).

types. Idiosyncratic income risk is modeled as the type-switching probability between low-income and high-income households, which are labeled as Hand-to-Mouth (HtM) households and Savers, respectively. Countercyclical fiscal policy is introduced as the sensitivity of government spending and targeted transfers with respect to aggregate output.

A virtue of our approach is that the welfare effect of changing the inflation target for each household type can be explicitly decomposed into its benefits and costs, providing a clear economic interpretation. This would be obscure in quantitative heterogeneous agent New Keynesian (HANK) models. To be more specific, we derive an individual-level quadratic welfare function and demonstrate that it consists of two components: a steady-state component, the change of which captures the marginal cost of increasing the inflation target, and a business cycle component, the change of which captures its marginal benefit. These two components are, in turn, functions of cyclical income inequality, long-run income inequality, idiosyncratic income risk, and the degree of countercyclical fiscal policy.

The steady-state component captures the welfare loss due to increased price dispersion in the steady state resulting from a positive inflation target. Raising the inflation target increases price dispersion, thereby reducing steady-state output. This reduction in steady-state output lowers steady-state utility unevenly across household types. The business cycle component—a linear combination of mean output losses, output variance, and inflation variance—captures the welfare

costs of business cycles at a given level of the inflation target. Raising the inflation target increases welfare in two ways. First, it reduces mean output losses and the volatility of output and inflation by lowering the frequency of hitting the ZLB. Second, these changes in aggregate moments have different welfare gains across households, as households have different weights attached to these moments.

The main findings are as follows. First, greater cyclical income inequality and higher income risk increase the marginal benefits of a higher inflation target for both HtMs and Savers. Under countercyclical income inequality, HtMs, who exhibit high marginal propensities to consume (MPCs) and face more volatile income streams, amplify aggregate demand fluctuations and increase the frequency of hitting the ZLB. This leads to more volatile inflation and output, as well as lower average output. Consequently, the welfare gains from raising the inflation target increase with the degree of cyclical inequality. Among the two groups, HtMs experience larger marginal benefits than Savers, since their consumption and income are more sensitive to aggregate output. The marginal benefits also increase with greater income risk for all households. Higher income risk strengthens the precautionary saving motive, which lowers the steady-state nominal interest rate. As a result, the frequency of hitting the ZLB rises, leading to more volatile output and inflation and larger mean output losses.

Second, greater long-run income inequality reduces the marginal benefits of trend inflation for both HtMs and Savers, while increasing the marginal costs for HtMs and decreasing them for Savers. Regarding marginal benefits, long-run income inequality dampens aggregate volatility by reducing the consumption share of HtMs, thereby weakening their contribution to aggregate demand fluctuations. This dampening effect partially offsets the amplification effect of cyclical inequality. Regarding marginal costs, under long-run inequality HtMs face a higher marginal utility of consumption because of their lower steady-state consumption. As a result, a given decline in steady-state income due to greater price dispersion generates a larger utility loss for HtMs. Conversely, because Savers have a lower marginal utility of consumption under greater long-run inequality, the same decline in steady-state income results in a smaller utility loss for them.

Third, given the calibrated levels of cyclical and long-run income inequality and income risk, the utilitarian optimal inflation rate is 1.95%, only 45 basis points higher than in the counterfactual representative agent New Keynesian (RANK) model. These quantitative results suggest that although the marginal cost of trend inflation is high, particularly for HtMs, the marginal benefit is sufficiently large to outweigh the cost. Introducing private debt raises the utilitarian optimal inflation rate to 2.40% for the following reason. When nominal policy rates are zero during ZLB episodes, deflation increases the real cost of debt through Fisherian debt deflation. This effect raises the burden on indebted HtMs, causing aggregate consumption and output to fall further relative to the case without wealth inequality.

Table 1: Directional Effects on the Marginal Benefits and Costs of Trend Inflation

Channel	Marginal benefits		Marginal costs	
	Savers	HtMs	Savers	HtMs
Cyclical income inequality	↑	↑	–	–
Long-run inequality	↓	↓	↓	↑
Idiosyncratic income risk	↑	↑	–	–
Countercyclical fiscal policy	↓	↓	–	–

Notes. ↑ denotes an increase, ↓ denotes a decrease, and – denotes no first-order effect.

Fourth, countercyclical government spending and targeted transfers toward HtMs reduce the marginal benefits of trend inflation for both HtMs and Savers, thereby lowering the optimal inflation rate at both the individual and utilitarian levels. This result holds even under progressive taxation, where Savers bear a larger share of the tax burden used to finance government spending or transfers. These countercyclical fiscal policies generate two opposing effects on the marginal benefit of trend inflation under progressive taxation. On the one hand, by imposing a larger tax burden on Savers during recessions, such policies make a given level of aggregate output volatility more costly for them than under less progressive taxation. This effect increases the marginal benefit of trend inflation for Savers, since less severe recessions associated with higher trend inflation reduce their tax burden. On the other hand, countercyclical fiscal interventions reduce the severity of ZLB episodes, thereby lowering output and inflation volatility and mitigating mean output losses. This reduces the marginal benefit of raising the inflation target for both Savers and HtMs. Under our calibration, the latter effect dominates the former.

Fifth, incorporating positive government debt raises the utilitarian optimal inflation rate for a given degree of government spending cyclicity, especially under low tax progressivity. When the ZLB binds, ex-post real rates rise because inflation falls. The higher real rate implies larger government debt servicing costs, which require higher taxes. While Savers can partially offset the decline in temporary disposable income, constrained HtMs experience a sharp fall in consumption. As a result, aggregate output contracts more severely in an economy with a higher level of government debt, increasing the marginal benefits of trend inflation.

Table 1 summarizes the directional effects of cyclical income inequality, long-run income inequality, idiosyncratic income risk, and countercyclical fiscal policy on the marginal benefits and costs of trend inflation for each household type.

Related Literature The present paper contributes to the growing HANK literature that explores the role of inequality in amplifying aggregate shocks. Existing quantitative HANK models typically feature a nondegenerate distribution of income and wealth, in which cyclical inequality,

long-run inequality, and income risk are intertwined.⁵ This interdependence makes it difficult to isolate the effects of each component, which we address in a tractable framework. [Bilbiie \(2008, 2024\)](#) and [Bilbiie, Känzig and Surico \(2022\)](#) examine the implications of income inequality for aggregate demand amplification in a model similar to ours, but they focus solely on cyclical inequality and abstract from the ZLB. In contrast, we separately analyze the effects of cyclical and long-run inequality and explore their normative implications.

Moreover, the present paper contributes to a growing body of research that incorporates the ZLB into HANK models. [Schaab \(2020\)](#) examines the propagation of macroeconomic uncertainty near the ZLB through its interaction with countercyclical unemployment risk. [Cho and Ma \(2024\)](#) study the role of inflation-indexed loan contracts in stabilizing business cycles. However, all of these studies assume constant inflation targets and therefore do not investigate the benefits and costs of raising the inflation target. [Fernández-Villaverde et al. \(2025\)](#) show that the benefits of raising the inflation target are greater in HANK than in RANK due to the more frequent occurrence of the ZLB. Nevertheless, they do not derive the optimal inflation target or decompose its costs and benefits across households.

Lastly, this paper contributes to the literature on optimal trend inflation that balances the benefits and costs of positive steady-state inflation. [Coibion, Gorodnichenko and Wieland \(2012\)](#), [Carreras et al. \(2016\)](#), [L’Huillier and Schoenle \(2023\)](#), and [Cho et al. \(2025\)](#) derive the optimal inflation target in ZLB models with Calvo pricing, while [Blanco \(2021\)](#) does so using a menu cost model. [Kim and Ruge-Murcia \(2009\)](#), [Fagan and Messina \(2009\)](#), [Abo-Zaid \(2013\)](#), [Benigno and Ricci \(2011\)](#), and [Mineyama \(2022\)](#) examine the role of downward nominal wage rigidity in rationalizing a positive inflation target. These studies, however, do not incorporate countercyclical income inequality or examine the heterogeneous effects of raising inflation targets. Some studies do investigate the disaggregate effects of higher inflation in the absence of the ZLB constraint. For example, [Doepke and Schneider \(2006\)](#) and [Adam and Zhu \(2016\)](#) study how unexpected inflation or deflation generates gains and losses through the revaluation of nominal assets and liabilities. In contrast, in our work the welfare effects of inflation operate through changes in business cycle moments and differential exposure to those moments across income groups.

The remainder of the paper is structured as follows. Section 2 presents the THANK model and its calibration. Section 3 derives individual welfare functions and explores the implications of cyclical and long-run inequality, income risk, and private debt for the optimal inflation targets. Section 4 examines the role of fiscal policy and government debt in shaping the optimal inflation targets. Section 5 concludes.

⁵See [McKay and Reis \(2016\)](#), [Auclert, Rognlie and Straub \(2020\)](#), [Cho \(2023\)](#), and [Bayer, Born and Luetticke \(2024\)](#) for quantitative contributions, among others.

2 THANK Model

We extend a tractable HANK model sketched in [Bilbiie, Känzig and Surico \(2022\)](#) and [Bilbiie \(2024\)](#) by allowing for positive trend inflation. Unlike them, we incorporate both countercyclical income inequality and steady-state income inequality to highlight the role of each factor in shaping the optimal trend inflation.

2.1 Model Description

2.1.1 Households

There exists a unit mass of households that derive utility from consumption C_t and disutility from labor N_t , with CRRA preferences represented by $\frac{C_t^{1-\sigma}-1}{1-\sigma} - \chi \frac{N_t^{1+\varphi}}{1+\varphi}$, where σ is relative risk aversion and φ is the inverse of the Frisch elasticity of labor supply. Households are divided into two types, Savers (S) and Hand-to-Mouth spenders (H), who differ in income volatility and the level of income: HtMs are exposed to a more volatile income stream and have a smaller steady-state income. We incorporate idiosyncratic income risk, which is the probability of households of one type switching to the other. The exogenous type-switching occurs according to a Markov chain, under which the probability of remaining a Saver is s , and that of remaining Hand-to-Mouth is h . Accordingly, the probability of transitioning from S to H is $1 - s$ and that from H to S is $1 - h$. In the stationary equilibrium, around which we approximate the model with respect to aggregate shocks, the stationary household distribution is as follows: $(1 - \lambda)$ for Savers and λ for Hand-to-Mouth, where $\lambda = \frac{(1-s)}{2-s-h}$.

For tractability, we assume households are members of a family whose head maximizes the family's welfare—the equally weighted utility of all households—while facing limited risk-sharing across household types. Households can be thought of as belonging to one of two states or "islands": all Savers are on island S , and all HtMs are on island H . At the beginning of the period, the family head pools resources within the island and distributes them equally to all households within the island. Once aggregate shocks are revealed, the family head determines the consumption and saving choices for each household on each island. Subsequently, households discover their type for the next period and transition to the corresponding island, carrying their risk-free bonds with them.

Flows of risk-free bonds across islands are as follows. The mass of HtMs becoming Savers each period is $(1 - h)\lambda$, while the remaining $h\lambda$ remain hand-to-mouth. Similarly, the mass of Savers that become HtMs is $(1 - s)(1 - \lambda)$. Hence, the beginning-of-period $t + 1$ per capita bond holdings, B_t^j , for household type $j \in \{S, H\}$ follow the relations:

$$(1 - \lambda)B_t^S = s(1 - \lambda)Z_t^S + (1 - h)\lambda Z_t^H \quad \text{and} \quad \lambda B_t^H = (1 - s)(1 - \lambda)Z_t^S + h\lambda Z_t^H,$$

where Z_t^j denotes the end-of-period t per capita quantity of bonds. Dividing each equation by the population measure yields:

$$\begin{aligned} B_t^S &= sZ_t^S + (1-s)Z_t^H \\ B_t^H &= (1-h)Z_t^S + hZ_t^H. \end{aligned} \quad (1)$$

The problem of the family head is as follows:

$$U(B_{t-1}^S, B_{t-1}^H) = \max_{C_t^S, C_t^H, Z_t^S, Z_t^H} (1-\lambda) \frac{(C_t^S)^{1-\sigma} - 1}{1-\sigma} + \lambda \frac{(C_t^H)^{1-\sigma} - 1}{1-\sigma} - \chi \frac{N_t^{1+\varphi}}{1+\varphi} + \beta \mathbf{E}_t [U(B_t^S, B_t^H)]$$

subject to the flows of bonds (1) and the budget constraints:

$$C_t^S + Z_t^S + TG_t^S = \theta_t^S ((1-\nu)W_t N_t + D_t) + \frac{R_{t-1}q_{t-1}}{\Pi_t} B_{t-1}^S + \mathcal{T}_t^S - TQ_t^S \quad (2)$$

$$C_t^H + Z_t^H + TG_t^H = \theta_t^H ((1-\nu)W_t N_t + D_t) + \frac{R_{t-1}q_{t-1}}{\Pi_t} B_{t-1}^H + \mathcal{T}_t^H - TQ_t^H \quad (3)$$

$$\theta_t^H = \theta_H \left(\frac{Y_t}{\bar{Y}} \right)^{\gamma_Y^H} \quad (4)$$

$$(1-\lambda)\theta_t^S + \lambda\theta_t^H = 1 \quad (5)$$

with borrowing constraints $Z_t^S, Z_t^H \geq 0$. W_t and D_t denote the real wage and dividends, respectively. Here, we assume that the utilitarian family head chooses equal labor inputs across household types, so that $N_t^S = N_t^H = N_t$. The household income side consists of three components: the sum of labor and dividend income, financial income from nominal bonds, and the transfer component. The main difference in the income process between the two household types arises from the first component, which depends on θ_t^j . First, γ_Y^H determines how much HtMs' income share covaries with aggregate output. More positive γ_Y^H exposes HtMs to greater income volatility by raising their income share during economic expansions and lowering it during recessions. Naturally, Savers' income share θ_t^S covaries negatively with aggregate output under positive γ_Y^H . Second, θ_H governs long-run income differential between Savers and HtMs. A lower θ_H means a smaller share of steady-state aggregate output goes to HtMs, resulting in a lower long-run income level for them compared to Savers. We assume that both households receive dividend income in proportion to their level of labor income share.⁶

TG_t^j and \mathcal{T}_t^j are taxes for government spending and transfer income, respectively. Depending

⁶In the previous version, we kept θ_H and γ_Y^H identical across household types and introduced unequal dividend distribution to generate countercyclical income inequality. We reached conclusions similar to those in the current setting: the utilitarian optimal inflation rates are higher in THANK than in RANK. However, because dividends are counterfactually countercyclical in sticky-price models, we considered modeling countercyclical inequality through unequal dividend distribution to be unrealistic. Therefore, as in McKay and Wolf (2023), we instead assume proportional dividends so that countercyclical income inequality is driven by differences in income shares across household types, where θ_t^j is the share of aggregate income $Y_t (\equiv (1-\nu)W_t N_t + D_t)$ received by type j .

on how fiscal policies respond to business cycles, TG_t^j and \mathcal{T}_t^j can deviate from their steady-state levels and adjust differently between the two household groups. Finally, q_t is the aggregate risk-premium shock that drives a wedge between the policy interest rate and the households' effective return on bond holding. A positive shock to q_t makes households save, generating a negative demand and possibly leading to ZLB episodes. In the case of positive bond holdings, a positive shock to q_t acts like a subsidy shock in the sense that it increases wealth by $\frac{R_{t-1}}{\Pi_t}(q_{t-1} - 1)B_{t-1}^j$. To neutralize this subsidy-like effect, we introduce lump-sum taxes $TQ_t^j = \frac{R_{t-1}}{\Pi_t}(q_{t-1} - 1)B_{t-1}^j$. In this way, the risk-premium shock purely represents the Euler equation wedge.

The optimality conditions for the above dynamic programming problem are

$$(C_t^S)^{-\sigma} = \beta \mathbf{E}_t \left[\frac{R_t q_t}{\Pi_{t+1}} (s(C_{t+1}^S)^{-\sigma} + (1-s)(C_{t+1}^H)^{-\sigma}) \right] + \frac{\Theta_t^S}{1-\lambda} \quad (6)$$

$$(C_t^H)^{-\sigma} = \beta \mathbf{E}_t \left[\frac{R_t q_t}{\Pi_{t+1}} ((1-h)(C_{t+1}^S)^{-\sigma} + h(C_{t+1}^H)^{-\sigma}) \right] + \frac{\Theta_t^H}{\lambda} \quad (7)$$

$$\Theta_t^S \geq 0, \quad Z_t^S \geq 0, \quad \Theta_t^S Z_t^S = 0 \quad (8)$$

$$\Theta_t^H \geq 0, \quad Z_t^H \geq 0, \quad \Theta_t^H Z_t^H = 0, \quad (9)$$

where Θ_t^S and Θ_t^H are the Lagrangian multipliers associated with budget constraints (2) and (3). Since HtMs are borrowing-constrained, $\Theta_t^H > 0$. We assume zero net supply of risk-free bonds, which implies $Z_t^S = 0$. Finally, the optimality condition with respect to labor is:

$$W_t C_t^{-\sigma} = \chi N_t^\varphi. \quad (10)$$

2.1.2 Firms

The firm side follows the structure of the canonical New Keynesian framework with Calvo price rigidity. Final goods producers use intermediate goods as inputs and produce final goods according to a CES technology:

$$Y_t = \left(\int Y_t(i)^{\frac{\varepsilon_p - 1}{\varepsilon_p}} di \right)^{\frac{\varepsilon_p}{\varepsilon_p - 1}}.$$

They sell final goods to households in a competitive market. Profit maximization by final goods producers yields the demand function for intermediate good i and the aggregate price index:

$$Y_t(i) = \left(\frac{P_t(i)}{P_t} \right)^{-\varepsilon_p} Y_t$$

$$P_t = \left(\int P_t(i)^{1-\varepsilon_p} di \right)^{\frac{1}{1-\varepsilon_p}}.$$

A continuum of intermediate goods firms produces differentiated goods in a monopolistically

competitive market using the linear technology $Y_t(i) = A_t N_t(i)$, where A_t is the total factor productivity subject to an exogenous shock. The profit maximization problem of a price-resetting firm is:

$$\max_{P_t(i)} E_t \left[\sum_{k=0}^{\infty} (\theta_p \beta)^k \left(\frac{C_{t+k}}{C_t} \right)^{-\sigma} \left(\frac{P_t(i)}{P_{t+k}} Y_{t+k}(i) - MC_{t+k} Y_{t+k}(i) \right) \right],$$

where the real marginal cost, MC_t , is $MC_t = (1 - \nu) \frac{W_t}{A_t}$, with ν being the subsidy rate that eliminates steady-state distortion. θ_p denotes the probability of being unable to adjust the price. The optimality conditions for the dynamic profit maximization problem are:

$$X_{1,t} = (C_t)^{-\sigma} MC_t Y_t + \theta_p \beta E_t [\Pi_{t+1}^{\varepsilon_p} X_{1,t+1}] \quad (11)$$

$$X_{2,t} = (C_t)^{-\sigma} Y_t + \theta_p \beta E_t [\Pi_{t+1}^{\varepsilon_p - 1} X_{2,t+1}] \quad (12)$$

$$\Pi_t^* = \frac{\varepsilon_p}{\varepsilon_p - 1} \frac{X_{1,t}}{X_{2,t}} u_t, \quad (13)$$

where u_t is the cost-push shock.

2.1.3 Government

The central bank sets the gross policy rate, R_t , following the Taylor rule subject to a ZLB constraint:

$$R_t^* = (R_{t-1}^*)^{\varphi_r} \left(\left(\frac{\Pi_t}{\bar{\Pi}} \right)^{\varphi_\pi} \left(\frac{Y_t}{\bar{Y}} \right)^{\varphi_y} \right)^{1 - \varphi_r}$$

$$R_t = \begin{cases} R_t^* & \text{if } R_t^* \geq 1 \\ 1 & \text{if } R_t^* < 1, \end{cases}$$

where R_t^* is the shadow rate. φ_r governs the interest rate smoothing, while φ_π and φ_y measure the sensitivity of policy rates to deviations of inflation and output from their respective steady states. The fiscal authority undertakes government spending and transfer policies following a feedback rule that responds to the deviation of output from its steady state:

$$G_t = G \left(\frac{Y_t}{\bar{Y}} \right)^{\gamma_Y^G}$$

$$\mathcal{T}_t^H = \mathcal{T}^H + \gamma_Y^T (Y_t - \bar{Y}).$$

Both γ_Y^G and γ_Y^T govern the countercyclicality of government spending and transfer policy, respectively. More negative values of γ_Y^G and γ_Y^T indicate that public spending and targeted transfers toward HtMs increase more aggressively in recessions. The way government spending and tar-

geted transfers toward HtMs are financed is determined by the following rules:

$$\begin{aligned} G_t &= (1 - \lambda)TG_t^S + \lambda TG_t^H \\ TG_t^S &= TG^S + \frac{1 - \alpha}{1 - \lambda} (G_t - G) \\ (1 - \lambda)\mathcal{T}_t^S + \lambda\mathcal{T}_t^H &= 0. \end{aligned}$$

Government spending is financed through lump-sum taxes levied on both Savers and HtMs. These taxes consist of two components: a constant steady-state component that funds steady-state government spending, and a cyclical component that adjusts with deviations of government spending from its steady-state level. The parameter α governs the allocation of the cyclical tax burden between the two household groups. When $\alpha = 0$, Savers bear the entire tax burden associated with cyclical government spending. In contrast, $\alpha = \lambda$ implies an equal tax burden across Savers and HtMs. Additionally, the targeted transfer policy requires Savers to finance the full amount of transfers to HtMs.

2.1.4 Market Clearing

The aggregate demand for goods equals the aggregate supply of goods:

$$(1 - \lambda)C_t^S + \lambda C_t^H + G_t = Y_t = (1 - \nu)W_t N_t + D_t.$$

Combining the bond market clearing condition with each household's budget constraint yields:

$$\begin{aligned} C_t^S + TG_t^S &= \theta_t^S Y_t + \mathcal{T}_t^S = Y_t^S \\ C_t^H + TG_t^H &= \theta_t^H Y_t + \mathcal{T}_t^H = Y_t^H, \end{aligned}$$

where Y_t^j denotes the total income of household type j . We characterize inequality as a triplet of three parameters, (γ_Y^H, μ_y, s) . First, γ_Y^H , which we refer to as the parameter governing cyclical income inequality, captures the extent to which HtMs' income is more volatile than that of Savers. Second, $\mu_y \equiv \frac{\theta^S Y + \mathcal{T}^S}{\theta^H Y + \mathcal{T}^H}$, which represents the steady-state income ratio between the two household types, captures the degree of long-run income inequality. Lastly, the switching probability s between Savers and HtMs reflects uninsurable individual income risk. When $\mu_y > 1$, Savers internalize the possibility of transitioning into a low-income state with higher income volatility, which strengthens the precautionary savings motive in the deterministic steady state, thereby lowering the real interest rate.⁷

The full set of equilibrium conditions is given in the Online Appendix. We solve the model us-

⁷In quantitative HANK models, the idiosyncratic productivity shock processes simultaneously affect cyclical fluctuations in income shares, long-run income inequality, and the equilibrium level of precautionary motives. In contrast, our framework isolates these elements using a parsimonious parameterization, enabling clear analysis of their distinct effects on household welfare.

Table 2: MODEL PARAMETERS

Parameter	Interpretation	Value
External calibration		
σ	Relative risk aversion	2
φ	Inverse of labor supply elasticity	1
χ	Labor disutility	1
θ_p	Calvo parameter	0.75
ε_p	Elast. of substitution b/w goods	7.67
λ	Measure of Hand-to-Mouth Households	0.3
s	Probability of staying at Saver	0.9998
φ_r	Interest rate smoothing	0.75
φ_π	Inflation responsiveness	2.5
φ_y	Output responsiveness	0.1
s_g	Government Spending to Output Ratio at SS	0.2
γ_y^G	Output Responsiveness of Government Spending	0
γ_y^T	Output Responsiveness of Transfer	0
ρ_q	AR(1) of risk premium shock	0.9
ρ_a	TFP Shock Persistence	0.96
ρ_u	Cost-Push Shock Persistence	0.89
σ_a	TFP Shock Volatility	0.0041
σ_u	Cost-Push Shock Volatility	0.0008
Internal calibration		
β	Discount factor	0.9958
γ_Y^H	Cyclical Income Inequality Parameter	1.5
θ^S	Long-Run Income Inequality Parameter (Saver)	1.1475
θ^H	Long-Run Income Inequality Parameter (HtM)	0.6557
σ_q	Std. of risk premium shock	0.00255

ing the perturbation method, approximating the equilibrium conditions around the deterministic steady state. To accommodate the occasionally binding ZLB constraint in our stochastic simulation, we use the Dynare OBC software developed by [Holden \(2016, 2023\)](#).

2.2 Calibration

We calibrate the model parameters both externally and internally. As for the externally calibrated parameters, we set their values to the standard ones commonly used in the literature. We set the relative risk aversion, σ , to 2, while the inverse of the labor supply elasticity, φ , and the disutility of labor, χ , are set to 1. The elasticity of substitution between a pair of intermediate goods, ε_p , is chosen to match the 15% markup. The HtM population measure, λ , is set to 30% to match the portion of hand-to-mouth households reported in [Kaplan, Violante and Weidner \(2014\)](#). The Calvo parameter, θ_p , matches the price duration of 4 quarters, in line with evidence provided by [Nakamura and Steinsson \(2008\)](#). The measure of income risk, s , is 0.9998, implying that Savers face a 0.02% probability of becoming hand-to-mouth, to match the steady-state 0.19% real interest rate gap between the THANK economy and the counterfactual RANK economy. This real rate gap is equal to the value obtained in the quantitative HANK model of [Fernández-Villaverde et al. \(2025\)](#). The steady-state annual inflation is 2%, consistent with the inflation target set in the U.S. The steady-state government spending to output ratio, s_g , is 20% to match the observed long-run

average of government spending to GDP ratio in the U.S. As for the monetary policy rule, we set the interest smoothing parameter, φ_r , and the responsiveness of the policy rate with respect to deviations of inflation, φ_π , and output, φ_y , to 0.75, 2.5, and 0.1, respectively. We assume no endogenous fiscal stabilization for the baseline model, that is, $\gamma_y^G = \gamma_Y^T = 0$. While the financing method for cyclical government spending is irrelevant when there is no endogenous fiscal stabilization, we set $\alpha = \lambda$. The processes of total factor productivity and cost-push shocks are taken from the estimated values in [Kulish, Morley and Robinson \(2017\)](#), who estimate a DSGE model with the ZLB on U.S. data. Finally, the AR(1) coefficient of the risk premium shock, ρ_q , is 0.9.

We now describe the targets and parameters for internal calibration. The first target is the difference in the cyclical income between the two household types. [Kwark and Ma \(2021\)](#) reports the relative volatility of income share using CPS data.⁸ They report that the ratio of the volatility of income of the bottom 30% to the volatility of the income of the rest is $\frac{\sigma(\log(Y_t^{\text{bottom}20}))}{\sigma(\log(Y_t^{\text{rest}}))}$
 $= \frac{\frac{2}{3} \times 2.51 + \frac{1}{3} \times 0.57}{\frac{1}{7}(0.57 + 2 \times 0.44 + 2 \times 0.37 + 1.5 \times 0.33 + 0.5 \times 1.18)} \approx 4$. We treat the model counterpart of this moment as the relative income share volatility between HtMs and Savers. The second target is long-run income distribution, which allows us to discipline the long-run income inequality between the two household types. We target the income distribution using the reported median income data from [Kaplan, Violante and Weidner \(2014\)](#). In [Kaplan, Violante and Weidner \(2014\)](#), the poor-HtM median income is \$ 20,000, the wealthy-HtM median income is \$ 50,000, and the Ricardian's median income is \$ 70,000. Since the calculated portions of poor HtMs and wealthy HtMs are 0.1 and 0.2, respectively, we compute the relative steady-state income between HtMs and Savers as $\frac{70000}{\frac{1}{3} \times 20000 + \frac{2}{3} \times 50000} = 1.75 = \mu_y$. The third target is a 1.5% per annum real interest rate in the deterministic steady state. Finally, we target a 10% ZLB hitting frequency for baseline calibration.⁹ Given these target moments, we jointly calibrate the cyclical income inequality parameter, γ_Y^H , the long-run income share parameter for HtMs, θ^H , the discount factor, β , and the standard deviation of risk premium shocks, σ_q . The calibrated parameter values are reported in Table 2.

3 Inequality and Optimal Inflation Target

In this section, we first derive the approximated welfare function for each household type and then investigate how each component of the welfare function changes with cyclical and long-run inequality. To focus on the role of cyclical and long-run inequality on welfare, we assume that the parameters governing countercyclical fiscal policy are equal to zero ($\gamma_Y^T = \gamma_Y^G = 0$). Then, we compute the optimal inflation target under the calibrated cyclical and long-run inequality.

⁸They define household income as the sum of labor income, self-employment income, and net asset income before tax.

⁹This value is quite conservative given the two recent recessions. U.S. interest rates were at the ZLB from early 2009 until late 2015 and again from Q2 2020 to Q1 2022. If the sample period spans Q3 1954 to Q4 2023, the empirical ZLB frequency would be approximately 13%.

3.1 Individual Welfare Function

We present the type-specific approximated welfare. Specifically, the second-order approximation of the expected per-period type $j \in \{S, H\}$ welfare is:¹⁰

$$EU^j \equiv \mathbf{E} \left(\frac{(C_t^j)^{1-\sigma} - 1}{1-\sigma} - \chi \frac{N_t^{1+\varphi}}{1+\varphi} \right) \quad (14)$$

$$\approx \underbrace{\Gamma^j}_{\text{steady-state component}} + \underbrace{\Gamma_y^j \mathbf{E}(\hat{y}_t) + \Gamma_\pi^j \mathbf{E}(\hat{\pi}_t) + \Gamma_{y^2}^j \mathbf{E}(\hat{y}_t^2) + \Gamma_{\pi^2}^j \mathbf{E}(\hat{\pi}_t^2) + \Gamma_{ya}^j \mathbf{COV}(\hat{y}_t, \hat{a}_t)}_{\text{business cycle component}}$$

with

$$\Gamma^j = \frac{(C^j)^{1-\sigma} - 1}{1-\sigma} - \chi \frac{N^{1+\varphi}}{1+\varphi} - \chi (Yv^p)^{1+\varphi} \frac{1}{2}(1+\varphi) \frac{\sigma_a^2}{1-\rho_a^2}$$

$$\Gamma_y^j = (C^j)^{1-\sigma} \theta_{cyj} - \chi (Yv^p)^{1+\varphi}$$

$$\Gamma_\pi^j = -\chi (Yv^p)^{1+\varphi} \frac{\varepsilon_p \bar{\theta}_p (\Pi - 1)}{(1 - \bar{\theta}_p \Pi)(1 - \bar{\theta}_p)}$$

$$\Gamma_{y^2}^j = (C^j)^{1-\sigma} \left(\theta_{cyj2} + \frac{1}{2}(1-\sigma)\theta_{cyj}^2 \right) - \chi (Yv^p)^{1+\varphi} \frac{1}{2}(1+\varphi)$$

$$\Gamma_{\pi^2}^j = -\chi (Yv^p)^{1+\varphi} \frac{1}{2} \frac{\varepsilon_p \bar{\theta}_p (\varepsilon_p (\Pi - 1) + 1)}{(1 - \bar{\theta}_p)^2}$$

$$\Gamma_{ya}^j = \chi (Yv^p)^{1+\varphi} (1 + \varphi)$$

$$\bar{\theta}_p = \theta_p \Pi^{\varepsilon_p - 1},$$

$$\hat{c}_t^j = \theta_{cyj} \hat{y}_t + \theta_{cyj2} \hat{y}_t^2 \quad j \in \{S, H\}$$

$$\theta_{cSy} = \left(\frac{C^S}{Y} \right)^{-1} \left(\frac{1 - \lambda \theta^H (1 + \gamma_Y^H)}{1 - \lambda} - \frac{\lambda \gamma_Y^T}{1 - \lambda} - \frac{1 - \alpha}{1 - \lambda} s_g \gamma_Y^G \right)$$

$$\theta_{cHy} = \left(\frac{C^H}{Y} \right)^{-1} \left(\theta^H (1 + \gamma_Y^H) + \gamma_Y^T - \frac{\alpha}{\lambda} s_g \gamma_Y^G \right)$$

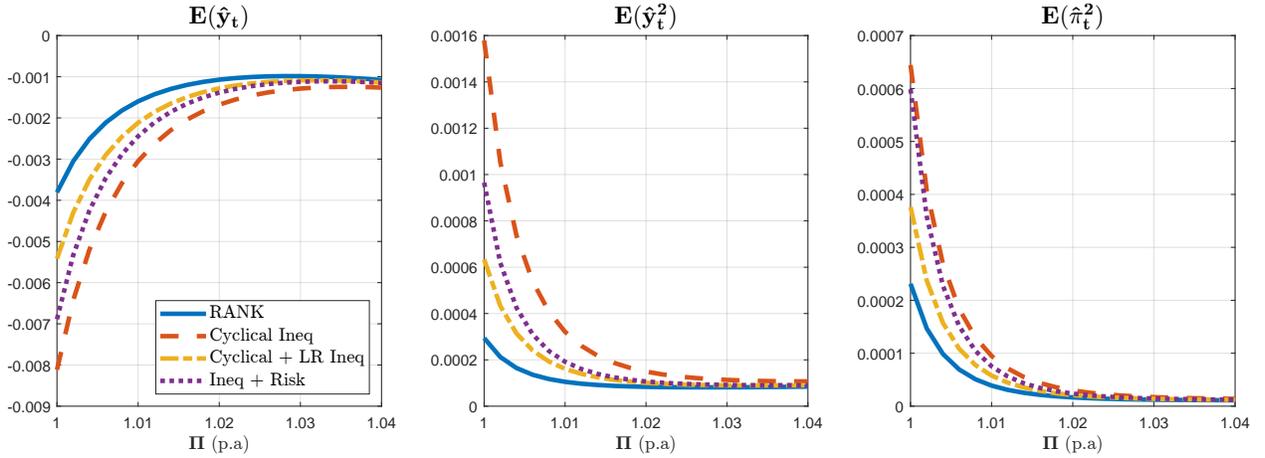
$$\theta_{cSy2} = \frac{1}{2} \left(\frac{C^S}{Y} \right)^{-1} \left(\frac{1 - \lambda \theta^H (1 + \gamma_Y^H)^2}{1 - \lambda} - \frac{\lambda \gamma_Y^T}{1 - \lambda} - \frac{1 - \alpha}{1 - \lambda} s_g (\gamma_Y^G)^2 \right. \\ \left. - \left(\frac{C^S}{Y} \right)^{-1} \left(\frac{1 - \lambda \theta^H (1 + \gamma_Y^H)}{1 - \lambda} - \frac{\lambda \gamma_Y^T}{1 - \lambda} - \frac{1 - \alpha}{1 - \lambda} s_g \gamma_Y^G \right)^2 \right)$$

$$\theta_{cHy2} = \frac{1}{2} \left(\frac{C^H}{Y} \right)^{-1} \left(\theta^H (1 + \gamma_Y^H)^2 + \gamma_Y^T - \frac{\alpha}{\lambda} s_g (\gamma_Y^G)^2 - \left(\frac{C^H}{Y} \right)^{-1} \left(\theta^H (1 + \gamma_Y^H) + \gamma_Y^T - \frac{\alpha}{\lambda} s_g \gamma_Y^G \right)^2 \right).$$

The approximate welfare function is a linear combination of several components: the constant term, the mean output, the mean inflation, the variance of output, the variance of inflation, and the covariance of output and the technology shock. Since business cycles have a negligible effect on mean inflation and the covariance terms throughout our simulation, we exclude the analysis regarding these terms and only discuss the terms that matter quantitatively for welfare.

¹⁰The full derivation is given in the Online Appendix.

Figure 2: Aggregate Moments



Notes. The figure displays $\mathbf{E}(\hat{y}_t)$, $\mathbf{E}(\hat{y}_t^2)$, and $\mathbf{E}(\hat{\pi}_t^2)$ as functions of trend inflation across four different economies. The solid line represents the RANK economy ($\gamma_Y^H = 0, \mu_y = 1, s = 1$); the dashed line represents an economy with cyclical income inequality ($\gamma_Y^H = 1.5, \mu_y = 1, s = 1$); the dash-dotted line corresponds to an economy with both cyclical and long-run income inequality ($\gamma_Y^H = 1.5, \mu_y = 1.75, s = 1$); and the dotted line depicts an economy with cyclical and long-run income inequality along with income risk ($\gamma_Y^H = 1.5, \mu_y = 1.75, s = 0.9998$).

Raising trend inflation affects welfare through two distinct channels: (1) the steady-state component (Γ^j), and (2) the business cycle component ($\Gamma_y^j \mathbf{E}(\hat{y}_t) + \Gamma y^{2j} \mathbf{E}(\hat{y}_t^2) + \Gamma \pi^{2j} \mathbf{E}(\hat{\pi}_t^2)$). Given the structure of income inequality in the economy, higher trend inflation introduces inefficiencies by increasing steady-state price dispersion across firms, as captured by changes in the steady-state component. At the same time, it reduces the welfare cost of aggregate fluctuations by lowering the frequency of hitting the ZLB, which is reflected in changes to the business cycle component. The optimal level of trend inflation balances the marginal welfare loss from greater long-run price dispersion against the marginal welfare gain from fewer ZLB episodes. Therefore, analyzing how each component responds to changes in trend inflation helps clarify the trade-offs faced by different household types. We begin by examining how trend inflation affects business cycle moments, which are common across all households, and then turn to the coefficients, which vary by household type.

3.1.1 Business Cycle Component: Aggregate Moments

Figure 2 shows how trend inflation affects three welfare-relevant business cycle moments across four economies that differ in their income inequality structures.¹¹ In all cases, higher trend inflation reduces the mean output loss and lowers aggregate volatility—formally, $\frac{\partial \mathbf{E}(\hat{y}_t)}{\partial \Pi} > 0$, $\frac{\partial \mathbf{E}(\hat{y}_t^2)}{\partial \Pi} < 0$, and $\frac{\partial \mathbf{E}(\hat{\pi}_t^2)}{\partial \Pi} < 0$.¹² Since higher trend inflation raises the steady-state nominal interest rate, it

¹¹ Figures 2 and 8 are generated from a single simulation of 100,000 periods without any burn-in for each inflation target.

¹² Due to the presence of the ZLB, output fluctuations are asymmetric: declines in output are larger than increases. This asymmetry implies that mean output falls below its deterministic counterpart, as reflected in $\mathbf{E}(\hat{y}_t) < 0$.

results in fewer ZLB episodes, thereby increasing mean output and reducing volatility in both output and inflation. For a given coefficient in the welfare function, these improvements in aggregate moments raise the welfare of all individuals.

The improvement in aggregate moments resulting from a marginal increase in the trend inflation rate depends critically on the structure of income inequality. Compared to the RANK economy (solid line), the economy with countercyclical income inequality (dashed line) exhibits larger mean output losses and greater volatility in both output and inflation. Since HtMs experience more volatile income than Savers and have a unit marginal propensity to consume, countercyclical income inequality amplifies fluctuations in aggregate demand, leading to more frequent ZLB episodes. As a result, raising trend inflation provides stronger stabilization in the presence of countercyclical income inequality than in the RANK economy—formally, $\frac{\partial^2 \mathbf{E}(\hat{y}_t)}{\partial \gamma^H \partial \Pi} > 0$, $\frac{\partial^2 \mathbf{E}(\hat{y}_t^2)}{\partial \gamma^H \partial \Pi} < 0$, and $\frac{\partial^2 \mathbf{E}(\hat{\pi}_t^2)}{\partial \gamma^H \partial \Pi} < 0$.

When long-run income inequality is added to cyclical income inequality (dash-dotted line), output losses and aggregate volatility become less pronounced. In the presence of long-run inequality, HtMs have a lower steady-state income than Savers, which limits the impact of their volatile income on aggregate output volatility. As a result, the amplitude of business cycles diminishes with greater long-run inequality, leading to a lower frequency of ZLB episodes. Consequently, a marginal increase in trend inflation yields a smaller improvement in aggregate moments compared to an economy with only cyclical income inequality—formally, $\frac{\partial^2 \mathbf{E}(\hat{y}_t)}{\partial \mu_y \partial \Pi} < 0$, $\frac{\partial^2 \mathbf{E}(\hat{y}_t^2)}{\partial \mu_y \partial \Pi} > 0$, and $\frac{\partial^2 \mathbf{E}(\hat{\pi}_t^2)}{\partial \mu_y \partial \Pi} > 0$.

Finally, when idiosyncratic income risk is introduced via type switching between Savers and HtMs, the precautionary saving motive lowers the steady-state real interest rate. This decline, in turn, reduces the steady-state nominal interest rate, increasing the frequency of ZLB episodes. As a result, mean output losses and aggregate volatilities become more pronounced (dotted line) relative to an economy without income risk. Consequently, in the presence of income risk, a marginal increase in trend inflation yields a greater improvement in business cycle moments compared to the no-risk case—formally, $\frac{\partial^2 \mathbf{E}(\hat{y}_t)}{\partial s \partial \Pi} > 0$, $\frac{\partial^2 \mathbf{E}(\hat{y}_t^2)}{\partial s \partial \Pi} < 0$, and $\frac{\partial^2 \mathbf{E}(\hat{\pi}_t^2)}{\partial s \partial \Pi} < 0$.

To sum up, holding the coefficients in the welfare function constant, the welfare gain from raising trend inflation arises from reductions in mean output losses and in the volatility of output and inflation. The extent to which a marginal increase in trend inflation improves these business cycle moments depends on the structure of income inequality. Cyclical income inequality and idiosyncratic income risk tend to increase the frequency of hitting the ZLB, thereby amplifying the marginal welfare gain from higher trend inflation. In contrast, long-run income inequality reduces the frequency of hitting the ZLB, diminishing the marginal welfare gain.

3.1.2 Business Cycle Component: Coefficients

The extent to which improvements in aggregate moments affect individual welfare depends on individual income levels and volatilities. The individual income profile, along with its interaction with the utility function, is reflected in the coefficients on aggregate moments (i.e., $\Gamma_y^j, \Gamma_{y^2}^j, \Gamma_{\pi^2}^j$). For instance, if a particular household group has a more negative coefficient on output volatility, a decline in output volatility yields greater welfare gains for that group relative to others. In this subsection, we examine how income inequality shapes the coefficients on aggregate moments.

Coefficients on mean output $E(\hat{y}_t)$ Γ_y^j is the weight of the mean output in the welfare function. This coefficient represents the extent to which the increased mean output affects individual utility. The increased mean output can decrease the utility by making households work more on average, but can increase the utility by increasing the mean consumption. Thus, the sign and magnitude of the mean output coefficient depends on the relative size of these two opposing effects. Using the fact that $\frac{C^S}{Y} = \frac{(1-s_g)\mu_y}{(1-\lambda)\mu_y + \lambda}$, $\frac{C^H}{Y} = \frac{1-s_g}{(1-\lambda)\mu_y + \lambda}$, $\theta_{cSy} = \frac{1-\lambda\theta_H(1+\gamma_Y^H)}{1-\lambda}$, and $\theta_{cHy} = \theta_H(1+\gamma_Y^H)$, this coefficient can be rewritten as:

$$\Gamma_y^S = Y^{1-\sigma} \underbrace{\left(\frac{(1-s_g)\mu_y}{(1-\lambda)\mu_y + \lambda} \right)^{-\sigma}}_{\text{utility curvature}} \underbrace{\frac{1-\lambda\theta_H(1+\gamma_Y^H)}{1-\lambda}}_{\text{consumption responsiveness}} - \underbrace{\chi(Yv^p)^{1+\varphi}}_{\text{labor disutility}}$$

$$\Gamma_y^H = Y^{1-\sigma} \left(\frac{(1-s_g)}{(1-\lambda)\mu_y + \lambda} \right)^{-\sigma} \theta_H(1+\gamma_Y^H) - \chi(Yv^p)^{1+\varphi}.$$

The coefficient can be decomposed into three terms: the term that captures the disutility from labor supply, the term that captures the curvature of the utility function, and the term that captures the consumption responsiveness to an increased output. Since Savers and HtMs supply the same amount of labor, the term that captures the disutility from labor supply, $\chi(Yv^p)^{1+\varphi}$, is identical across household types. However, the latter two terms are different across households. The following proposition characterizes how income inequality influences heterogeneous welfare exposure to the mean output change.

Proposition 1. *Assume households are risk-averse such that $\sigma > 1$, then*

1. $\Gamma_y^S \leq \Gamma_y^H$.
2. $\frac{\partial \Gamma_y^S}{\partial \gamma_Y^H} < 0$ and $\frac{\partial \Gamma_y^H}{\partial \gamma_Y^H} > 0$.
3. $\frac{\partial \Gamma_y^H}{\partial \mu_y} > 0$ and $\frac{\partial \Gamma_y^S}{\partial \mu_y} < (>)0$ if $\gamma_Y^H < (>)\frac{(\sigma-1)\lambda(1-\lambda)\mu_y}{\lambda(1-\lambda)\mu_y + \sigma\lambda^2}$.

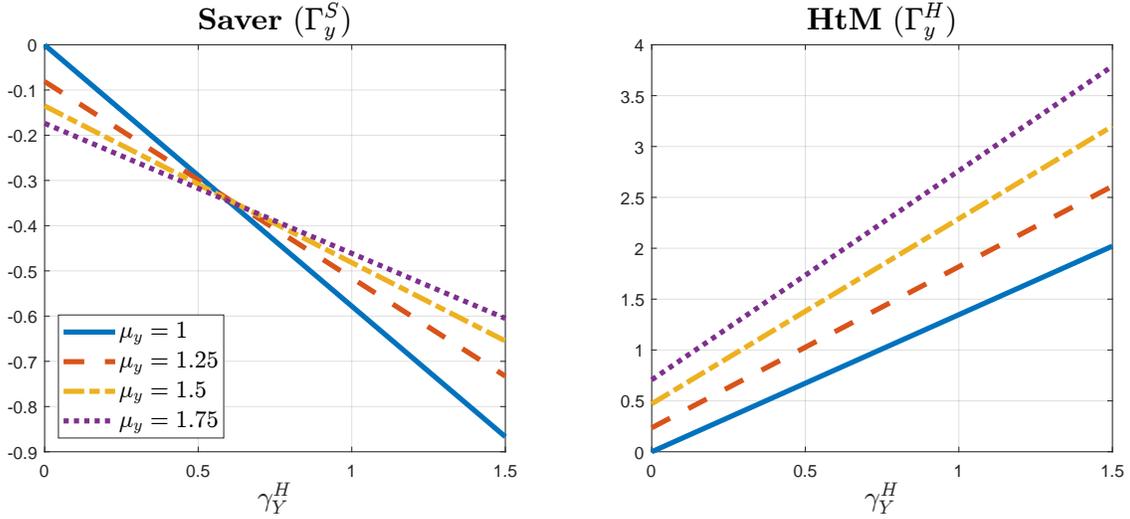
Proof. Refer to Appendix A. □

The first statement of Proposition 1 indicates that reducing the magnitude of mean output losses (or equivalently, increasing mean output) improves welfare more for HtMs than for Savers. This heterogeneous welfare effect arises from the utility curvature and differences in consumption responsiveness across household types. Regarding the utility curvature component, the relevant term is larger for HtMs than for Savers (i.e., $Y^{1-\sigma} \left(\frac{(1-s_g)\mu_y}{(1-\lambda)\mu_y+\lambda} \right)^{-\sigma} < Y^{1-\sigma} \left(\frac{(1-s_g)}{(1-\lambda)\mu_y+\lambda} \right)^{-\sigma}$), when long-run inequality exists ($\mu_y > 1$). Under long-run inequality, HtMs consume less in steady state and thus have a higher marginal utility of consumption. Consequently, a given increase in consumption generates a larger utility gain for HtMs. Turning to the consumption responsiveness component, when income inequality is sufficiently cyclical such that $1 + \gamma_Y^H > \mu_y$, the term is larger for HtMs than for Savers (i.e., $\frac{1-\lambda\theta_H(1+\gamma_Y^H)}{1-\lambda} < \theta_H(1+\gamma_Y^H)$), meaning that an increase in mean output raises consumption more for HtMs than for Savers. Taken together, in the presence of long-run and cyclical inequality, $\Gamma_y^S < \Gamma_y^H$ holds, indicating that HtMs benefit more from reducing mean output losses than Savers. In the absence of inequality (i.e., $\gamma_Y^H = 0, \mu_y = 1$), the coefficients are identical across household types.

The second statement illustrates how the degree of cyclical income inequality affects the individual welfare gains from changes in mean output. Greater countercyclical income inequality (i.e., a higher value of γ_Y^H) reduces Savers' preference for higher mean output while increasing that of HtMs. This result can be understood through the relationships $(1-\lambda)\theta_t^S + \lambda\theta_t^H = 1$ and $\theta_t^H = \theta^H \left(\frac{Y_t}{Y} \right)^{\gamma_Y^H}$. An increase in γ_Y^H raises the sensitivity of HtMs' income to aggregate output θ_t^H , which in turn reduces the sensitivity of Savers' income θ_t^S . As a result, higher mean aggregate output translates into larger increases in HtMs' mean income and consumption, and smaller increases for Savers. This explains why more countercyclical inequality raises the welfare weight on mean output for HtMs while lowering it for Savers.

The final statement describes the role of long-run income inequality in shaping the coefficients on mean output. Specifically, a higher level of long-run inequality raises the mean output coefficient for HtMs. In contrast, it lowers the coefficient for Savers when γ_Y^H is small, but raises it when γ_Y^H is large. There are two opposing channels through which long-run inequality affects the mean output coefficient for HtMs. The first is a curvature effect. Greater long-run inequality implies that HtMs have lower steady-state income and consumption relative to Savers. Given the concavity of the utility function, HtMs experience a higher marginal utility of consumption. As a result, an increase in mean output yields larger utility gains for HtMs when long-run inequality is high. The second is a quantity effect. With greater long-run inequality, HtMs' aggregate income share becomes smaller. Hence, the share of aggregate income allocated to HtMs becomes smaller when there is an additional increase in mean aggregate income, making them value higher mean output less. When $\sigma > 1$, the curvature effect dominates, explaining why a higher level of long-run inequality increases the HtMs' weight on the mean output.

Figure 3: Mean Output Coefficients



Notes. The figure displays the mean output coefficient for Savers (left) and hand-to-mouth spenders (right) as functions of cyclical (γ_Y^H) and long-run inequality (μ_y).

For Savers, the curvature effect works adversely. Greater long-run inequality implies higher steady-state income and consumption for Savers relative to HtMs. As a result, Savers face lower marginal utility of consumption, and the utility gain from an increase in mean output becomes smaller. In contrast, the quantity effect works in favor of Savers. With greater long-run inequality, Savers' share of aggregate income rises. Consequently, any additional increase in mean aggregate income leads to a larger increase in the income allocated to Savers. For a small value of γ_Y^H , the curvature effect dominates the quantity effect, causing Γ_y^S to decrease with μ_y . For a large value of γ_Y^H , the quantity effect outweighs the curvature effect, causing Γ_y^S to increase with μ_y .

Figure 3 illustrates how the mean output coefficients vary with cyclical and long-run inequality, represented by γ_Y^H and μ_y , respectively. The coefficient for HtMs is positive, indicating that they benefit from an increase in mean output, as the utility gain from higher consumption outweighs the disutility from increased labor. In contrast, the coefficient for Savers is negative, reflecting that the disutility from additional labor dominates the utility gain from consumption. Consistent with Proposition 1, the HtMs' coefficient increases with the level of cyclical inequality, while the Savers' coefficient declines. Additionally, greater long-run inequality is associated with a higher coefficient for HtMs. However, for Savers, the effect of long-run inequality depends on the level of cyclical inequality: when γ_Y^H is large, the Savers' coefficient increases with long-run inequality; when γ_Y^H is small, it decreases.

The total welfare effect of a marginal increase in trend inflation through the mean output component can be expressed as $\frac{\partial \Gamma_y^j}{\partial \Pi} E(\hat{y}_t) + \Gamma_y^j \frac{\partial E(\hat{y}_t)}{\partial \Pi}$. Since $\frac{\partial \Gamma_y^j}{\partial \Pi}$ is negligibly small, total marginal welfare effect can be summarized as $\Gamma_y^j \frac{\partial E(\hat{y}_t)}{\partial \Pi}$. Here, $\frac{\partial E(\hat{y}_t)}{\partial \Pi}$ captures the mean output change

resulting from a higher inflation target, which is positive (as shown in the left panel of Figure 2), and Γ_y^j reflects the welfare gain from a given increase in mean output. Given that the mean output coefficient is positive for HtMs, higher trend inflation leads to greater welfare gains for them. In contrast, since the coefficient is negative for Savers, higher trend inflation reduces their welfare.

Coefficients on the variance of output $\mathbf{E}(\hat{y}_t^2)$ $\Gamma_{y^2}^j$ captures the welfare exposure to the variance of output and can differ across households. Substituting out individual consumption-output ratio, the output variance coefficients become:

$$\begin{aligned} \Gamma_{y^2}^S &= \frac{1}{2} Y^{1-\sigma} \left(\frac{(1-s_g)\mu_y}{(1-\lambda)\mu_y + \lambda} \right)^{-\sigma} \left(\frac{1-\lambda\theta_H(1+\gamma_Y^H)^2}{1-\lambda} - \sigma \left(\frac{(1-s_g)\mu_y}{(1-\lambda)\mu_y + \lambda} \right)^{-1} \left(\frac{1-\lambda\theta_H(1+\gamma_Y^H)}{1-\lambda} \right)^2 \right) \\ &\quad - \frac{1}{2}(1+\varphi)\chi(Yv^p)^{1+\varphi} \\ \Gamma_{y^2}^H &= \underbrace{\frac{1}{2} Y^{1-\sigma} \left(\frac{(1-s_g)}{(1-\lambda)\mu_y + \lambda} \right)^{-\sigma}}_{\text{utility curvature}} \left(\underbrace{\theta_H(1+\gamma_Y^H)^2}_{\text{income convexity}} - \underbrace{\sigma \left(\frac{(1-s_g)}{(1-\lambda)\mu_y + \lambda} \right)^{-1} \theta_H^2(1+\gamma_Y^H)^2}_{\text{income volatility}} \right) \\ &\quad - \underbrace{\frac{1}{2}(1+\varphi)\chi(Yv^p)^{1+\varphi}}_{\text{labor supply fluctuation}}. \end{aligned}$$

The output variance coefficients can be decomposed into four distinct components. The first captures the utility curvature effect, which depends on the degree of long-run income inequality. For agents with lower steady-state income (i.e., HtMs), a given level of output fluctuation is more harmful, as they face higher marginal utility. The second component reflects income convexity, which is more pronounced for HtMs. As can be seen from $\theta_t^H = \theta^H \left(\frac{Y_t}{Y}\right)^{\gamma_Y^H}$, a greater cyclical inequality (i.e., higher γ_Y^H) makes HtMs' income cyclical θ_t^H more convex in aggregate output. Thus, for a given amount of output variation, greater cyclical inequality increases the HtMs' mean income and consumption by increasing the mean of θ_t^H , enhancing their utility. The third component captures welfare losses from individual income volatility, which are larger for HtMs. When an individual's income responds more strongly to fluctuations in output, output volatility translates into greater individual income volatility. Due to the concavity of the utility function in consumption, more volatile individual income leads to larger welfare losses. Finally, the last term represents utility losses from labor supply fluctuations, stemming from the convexity of the labor disutility function.

Proposition 2. *Assume households are risk-averse such that $\sigma > 1$ and $\lambda < \frac{1}{2}$. Then:*

1. *There exists a threshold level of risk aversion $\bar{\sigma}(\lambda, s_g)$ such that $\Gamma_{y^2}^H \leq \Gamma_{y^2}^S \leq 0$ if $\sigma \geq \bar{\sigma}(\lambda, s_g)$.*
2. $\frac{\partial \Gamma_{y^2}^H}{\partial \gamma_Y^H} < 0$ and $\frac{\partial \Gamma_{y^2}^H}{\partial \mu_y} < 0$.

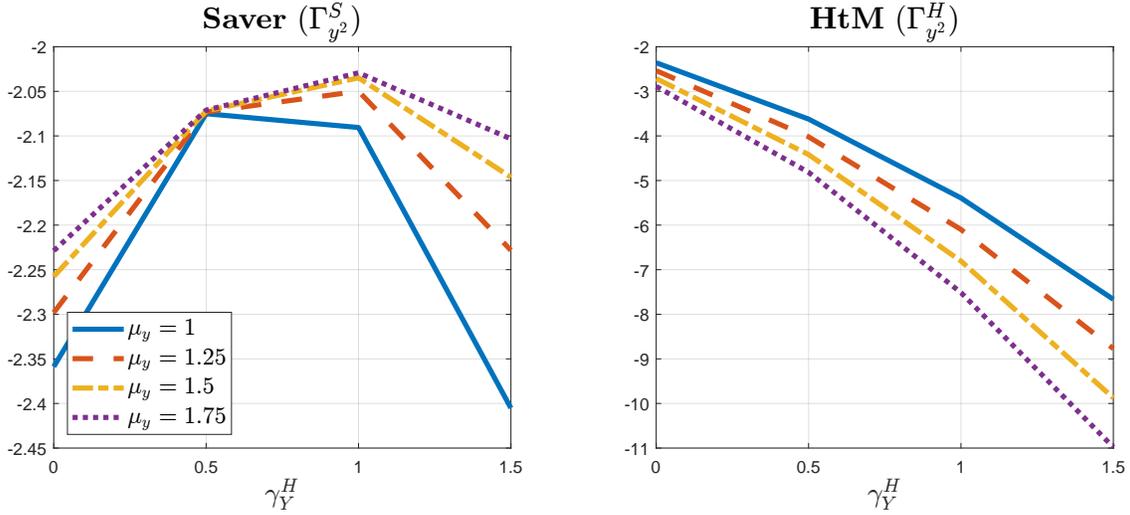
Proof. Refer to Appendix A. □

The first statement of Proposition 2 implies that output volatility is painful for the entire population, and more so for HtMs than for Savers, if the consumption utility function is sufficiently concave. Intuitively, since HtMs' steady-state consumption is lower than that of Savers in the presence of long-run inequality, HtMs' consumption fluctuates around a point at which the concavity of utility is greater than that of Savers. Thus, a given amount of output variance lowers the expected utility of HtMs more than that of Savers. In addition, since HtMs' consumption varies more sensitively in response to aggregate output changes in the presence of countercyclical inequality, concave utility implies that output variance causes the expected utility of HtMs to be lower than that of Savers.

The second statement of Proposition 2 indicates that more severe cyclical and long-run income inequality increases the welfare cost of output fluctuations for HtMs. Greater cyclical inequality (i.e., a higher γ_Y^H) makes HtMs' income more convex in aggregate output and their individual income more volatile. While the former effect (income convexity) is welfare-improving and the latter (income volatility) is welfare-reducing, the proposition states that the latter effect dominates, provided HtMs are sufficiently risk-averse. In the case of greater long-run inequality, a larger steady-state income gap causes HtMs' consumption to fluctuate around a point where the concavity of utility is higher, further lowering their expected utility in response to output volatility. On the other hand, with greater long-run inequality, the share of aggregate income allocated to HtMs during high-output periods declines, reducing their individual income volatility and thus mitigating welfare losses. Nonetheless, the proposition shows that the curvature-based effect dominates, implying that long-run income inequality decreases the output variance coefficient for HtMs.

In contrast to HtMs, the effect of income inequality on the Savers' coefficient is non-monotonic. As with HtMs, there are two opposing channels through which γ_Y^H affects the Savers' weight on the variance of output. On the one hand, more countercyclical inequality makes Savers' income and consumption less volatile in response to aggregate output changes, reducing the welfare cost of output variance. On the other hand, as shown by the expression $\theta_t^S = \frac{1}{(1-\lambda)} - \frac{\lambda\theta^H\left(\frac{Y_t}{Y}\right)^{\gamma_Y^H}}{(1-\lambda)}$, a greater cyclical inequality (i.e., higher γ_Y^H) makes Savers' income cyclical θ_t^S more concave in aggregate output. Thus, for a given amount of output variation, greater cyclical inequality decreases the Savers' mean income and consumption by decreasing the mean of θ_t^S , which is detrimental to their utility. Greater long-run income inequality also exerts two offsetting effects on the Savers' coefficient. On the positive side, higher steady-state income reduces the curvature of the utility function around Savers' consumption level, thereby diminishing the welfare cost of output fluctuations. Conversely, as Savers' share of steady-state income rises, their individual income is more exposed to variation in aggregate output, increasing the cost of output volatility. Whether cyclical and long-run inequality increases or decreases the welfare cost of output volatility depends on the relative strength of these opposing effects.

Figure 4: Output Variance Coefficients

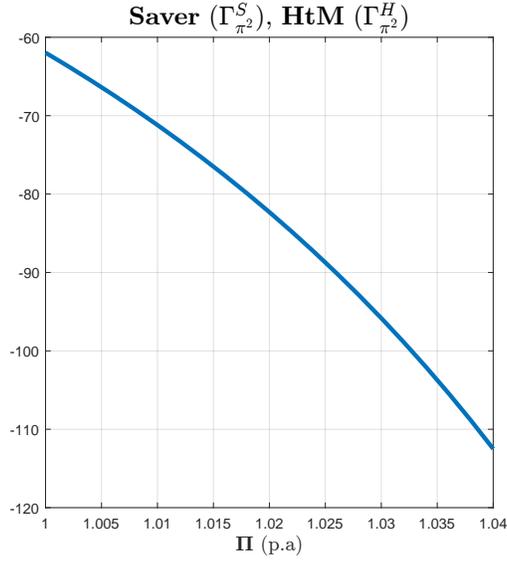


Notes. The figure displays the output variance coefficient for Savers (left) and hand-to-mouth spenders (right) as functions of cyclical (γ_Y^H) and long-run inequality (μ_y).

Figure 4 illustrates how the output variance coefficients vary with cyclical and long-run inequality, represented by γ_Y^H and μ_y , respectively. Note that the coefficients are negative for both household types, with the HtMs' coefficient being more negative for any given combination of γ_Y^H and μ_y . Compared to the RANK economy ($\gamma_Y^H = 0$, $\mu_y = 1$), HtMs perceive output volatility to be approximately five times more detrimental under the baseline calibration ($\gamma_Y^H = 1.5$, $\mu_y = 1.75$). In contrast, the Savers' coefficient increases with higher values of μ_y but exhibits a non-monotonic relationship with γ_Y^H . Around the baseline calibration, greater steady-state income inequality reduces the welfare cost of output volatility for Savers. Moreover, at low levels of γ_Y^H , countercyclical income inequality initially reduces Savers' cost of output volatility, but at higher levels of γ_Y^H , it increases the cost. The Savers' coefficient under the baseline calibration is slightly higher than in the RANK model, indicating that output volatility is somewhat less costly in an economy with income inequality. Overall, from the magnitude of the output variance coefficient under the baseline calibration, it follows that higher trend inflation enhances household welfare by reducing output volatility, with the welfare gain being larger for HtMs.

Coefficients on the variance of inflation $\mathbb{E}(\hat{\pi}_t^2)$ $\Gamma_{\pi^2}^j$ is the weight on inflation variance in the welfare function and represents the utility cost associated with a given level of inflation volatility. As seen in equation (14), this coefficient is negative and identical for both Savers and HtMs, indicating that the welfare cost of inflation variability is equally shared across households. The intuition behind the negative sign and uniform magnitude of this coefficient across household types is as follows. First, higher inflation volatility increases price dispersion across firms. For a given level of output, greater price dispersion requires more labor input, which reduces house-

Figure 5: Inflation Variance Coefficient



Notes. The figure displays the inflation variance coefficient as a function of the trend inflation rate.

hold utility. Since labor supply is identical across households, the utility cost from increased labor input is the same for all groups. Second, because the shares of labor and dividend income are identical for both Savers and HtMs, the welfare cost of increased price dispersion that arises from consumption is also the same across households. To be more specific, rewriting the household budget constraint, assuming no cyclical government spending ($TG_t^j = TG^j$) and transfer ($\mathcal{T}_t^j = 0$), we obtain:

$$C_t^j + TG^j = Y_t^j \equiv \theta_t^j Y_t = \theta_t^j ((1 - \nu)W_t N_t + D_t) = \theta_t^j (MC_t v_t^p Y_t + (1 - MC_t v_t^p) Y_t)$$

for $j \in \{S, H\}$. Higher price dispersion, v_t^p , increases aggregate labor income, $MC_t v_t^p Y_t$, and reduces aggregate dividend income, $(1 - MC_t v_t^p) Y_t$. Since group j receives a constant share θ^j of both aggregate labor income and aggregate dividends, the increase in labor income and the decrease in dividend income exactly offset each other. As a result, the rise in price dispersion has zero net effect on consumption for either household group, implying that both groups are equally exposed to inflation variability.

Contrary to the coefficients on the first and second moments of output deviations, the coefficient on inflation variance is invariant to the degree of income inequality. Instead, higher trend inflation leads to a substantial increase in the inflation variance coefficient in absolute value, as shown in Figure 5. Since higher trend inflation raises steady-state price dispersion and labor supply, additional price dispersion arising from inflation volatility requires even more labor input. This becomes increasingly costly due to the convexity of the labor disutility function. Thus, higher trend inflation affects household welfare with respect to inflation volatility through two opposing

channels. On the one hand, it reduces the frequency of ZLB episodes, lowering inflation volatility and thereby improving household welfare, as shown in Figure 2. On the other hand, it increases steady-state labor supply, amplifying the disutility from a given level of inflation volatility. The net welfare effect of a higher inflation target through the inflation volatility component depends on the relative strength of these two offsetting forces.

3.1.3 Steady-State Component

We now turn to the steady-state welfare impact of a higher inflation target, captured by Γ^j . This term represents the steady-state welfare cost that arises from steady-state price dispersion across firms due to positive trend inflation. The additional steady-state welfare costs of marginally raising the inflation target are represented by $\frac{\partial \Gamma^j}{\partial \Pi}$, which is negative for quantitatively relevant inflation rates. Intuitively, with a higher inflation target, steady-state output declines due to the inefficiencies created by the increased steady-state price dispersion. Accordingly, the steady-state welfare loss is increasing in the steady-state level of inflation for both types of agents. However, these welfare losses are not identical across agents, as stated by the following proposition.

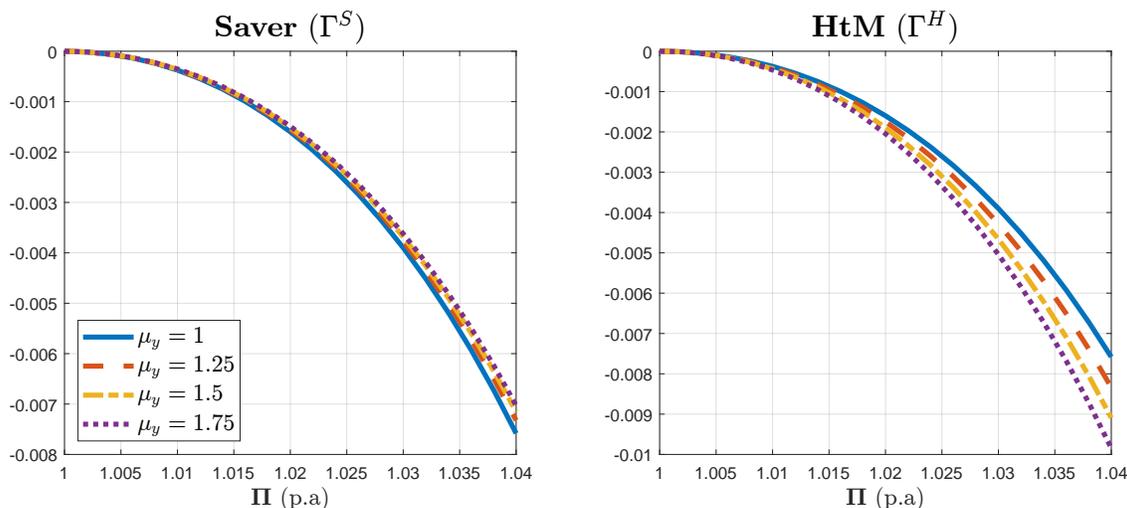
Proposition 3. *Assume that long-run inequality exists (i.e., $\mu_y > 1$) and that agents are risk-averse such that $\sigma > 1$. Then, $\left| \frac{\partial \Gamma^S}{\partial \Pi} \right| < \left| \frac{\partial \Gamma^H}{\partial \Pi} \right|$.*

Proof. Refer to Appendix A. □

Proposition 3 states that an increase in trend inflation is more costly for HtMs. Two opposing forces explain why HtMs experience a larger utility loss than Savers when trend inflation rises. The first is the heterogeneous quantity effect. In the presence of long-run inequality, Savers receive a larger share of aggregate income than HtMs. As a result, a given decline in aggregate income induced by higher trend inflation reduces the absolute income of Savers more than that of HtMs, making HtMs less painful than Savers. The second is the heterogeneous curvature effect. With long-run consumption inequality, a decline in consumption has asymmetric welfare effects. Since HtMs have lower steady-state consumption, the same reduction in consumption entails a larger utility loss for them due to the concavity of the utility function. When $\sigma > 1$, the utility function is sufficiently concave for the curvature effect to dominate the quantity effect. In this case, HtMs suffer greater welfare losses from higher trend inflation than Savers. Conversely, when $\sigma < 1$, the quantity effect dominates, and the welfare loss is larger for Savers. Finally, under log utility ($\sigma = 1$), the two effects offset each other, and the steady-state cost of trend inflation is borne equally by both groups.

Figure 6 illustrates how the steady-state component of the welfare function varies with the inflation target for different levels of long-run inequality μ_y , holding all other parameters at their baseline values. As expected, steady-state utility declines with higher trend inflation for both household types. For any given level of trend inflation, an increase in long-run inequality widens the utility gap between HtMs and Savers, indicating that HtMs bear a larger steady-state cost

Figure 6: Steady-State Utility Cost



Notes. The figure displays the steady-state cost of trend inflation for Savers (left) and hand-to-mouth spenders (right).

relative to Savers as inequality rises. This greater cost for HtMs reflects the dominance of the curvature effect over the quantity effect. Note that cyclical inequality does not affect steady-state costs, as its impact emerges only when the economy deviates from the steady state.

3.1.4 Summary: Income Inequality and Trend Inflation Trade-Off

The analysis in Section 3.1 clarifies how income inequality alters the cost–benefit trade-off of a higher inflation target. The presence of inequality affects the marginal benefit of trend inflation by changing both the aggregate moments (i.e., the mean and variance of output and inflation) and the weights assigned to these moments in the welfare function. Cyclical income inequality and idiosyncratic income risk tend to increase the frequency of hitting the ZLB, thereby amplifying mean output losses as well as the variance of output and inflation. In contrast, long-run income inequality reduces the likelihood of hitting the ZLB, thereby mitigating mean output losses and aggregate volatility. Because households have different weights to these aggregate moments—owing to differences in income levels and consumption volatility—the marginal benefit of higher trend inflation varies across household types. As Propositions 1 and 2 indicate, marginal benefit is larger for HtMs since they perceive mean output losses and macroeconomic volatility more painfully than Savers.

The presence of inequality also alters the marginal cost of higher trend inflation by affecting the steady-state component of the welfare function. As stated in Proposition 3, greater long-run income inequality increases the steady-state cost of trend inflation more for HtMs than for Savers.

Whether inequality changes the optimal trend inflation depends on the extent to which it al-

ters the marginal benefits and costs. A priori, the answer is ambiguous for the following reason. On the one hand, if the marginal benefit of higher trend inflation is substantially greater for HtMs than in the RANK economy, households in an unequal economy may prefer a higher inflation target.¹³ On the other hand, since the marginal cost is also higher for HtMs than in the RANK economy, households in an unequal economy may prefer a lower inflation target.¹⁴ If the increase in marginal benefits is fully offset by the increase in marginal costs, then the optimal inflation target in an unequal world may be indistinguishable from that in the RANK economy. To determine whether the marginal benefit or cost dominates, we quantitatively evaluate the optimal inflation targets at both the household and utilitarian levels in the next subsection.

3.2 Optimal Inflation Target

We compute the optimal inflation target for each group by maximizing their unconditional expected welfare. The unconditional expected welfare solves:

$$\begin{aligned} EV^S(\Pi | \gamma_Y^H, \mu_Y, s) &= EU^S(\Pi | \gamma_Y^H, \mu_Y, s) + s\beta EV^S(\Pi | \gamma_Y^H, \mu_Y, s) + (1-s)\beta EV^H(\Pi | \gamma_Y^H, \mu_Y, s) \\ EV^H(\Pi | \gamma_Y^H, \mu_Y, s) &= EU^H(\Pi | \gamma_Y^H, \mu_Y, s) + (1-h)\beta EV^S(\Pi | \gamma_Y^H, \mu_Y, s) + h\beta EV^H(\Pi | \gamma_Y^H, \mu_Y, s), \end{aligned}$$

where $EV^j(\Pi | \gamma_Y^H, \mu_Y, s)$ denotes the unconditional expected welfare of households in group $j \in \{S, H\}$, conditional on the gross steady-state inflation rate Π and the inequality structure (γ_Y^H, μ_Y, s) . The household-specific optimal inflation rate, denoted by $\Pi^{j,opt}(\gamma_Y^H, \mu_Y, s)$, maximizes $EV^j(\Pi | \gamma_Y^H, \mu_Y, s)$. The utilitarian (aggregate) optimal trend inflation, $\Pi^{opt}(\gamma_Y^H, \mu_Y, s)$, maximizes the weighted sum $(1-\lambda)EV^S(\Pi | \gamma_Y^H, \mu_Y, s) + \lambda EV^H(\Pi | \gamma_Y^H, \mu_Y, s)$.

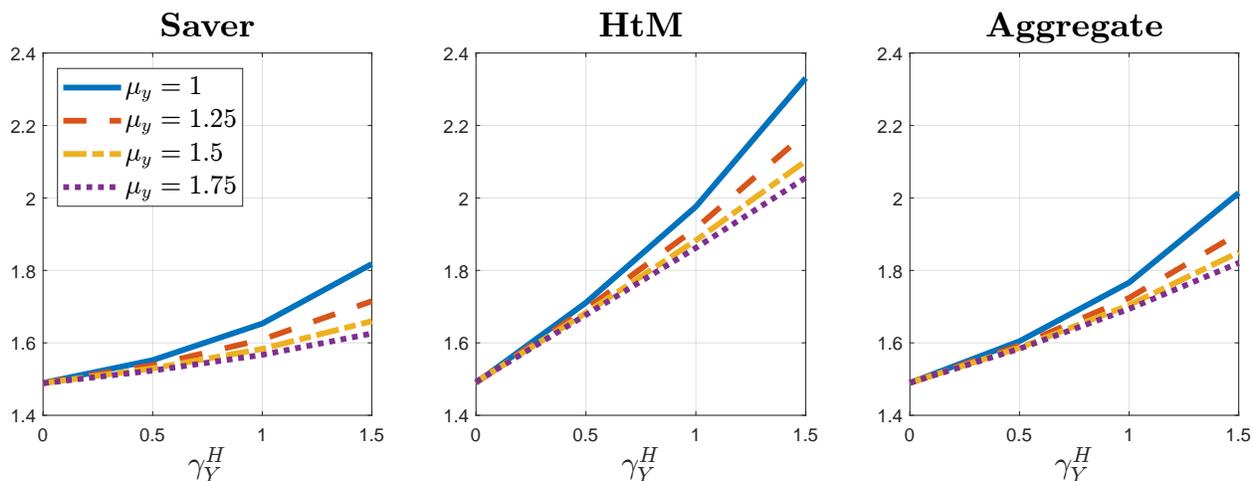
To clearly investigate the role of cyclical and long-run inequality in determining the optimal inflation target, we first assume there is no idiosyncratic income risk (i.e., $s = 1$). The left and middle panels of Figure 7 display the optimal inflation targets for Savers and HtMs as functions of cyclical and long-run inequality. The model without cyclical and long-term inequality (i.e., $\gamma_Y^H = 0, \mu_y = 1$) corresponds to a RANK model, in which the optimal inflation target is 1.49% and identical for both groups. When cyclical and long-run inequality are present (i.e., $\gamma_Y^H > 0, \mu_y > 1$), the optimal inflation targets diverge across groups, with HtMs preferring a higher target than Savers for any given combination of γ_Y^H and $\mu_y > 1$. Unsurprisingly, the aggregate optimal inflation target lies between the two household-specific targets, as shown in the right panel.

Given long-run inequality, more countercyclical income inequality (i.e., an increase in γ_Y^H) raises the optimal inflation target for all households. With more countercyclical inequality, the variance of output and inflation increases, and mean output losses become more severe. In this

¹³Recall that under the baseline calibration ($\gamma_Y^H = 1.5, \mu_y = 1.75$), HtMs perceive output volatility to be approximately five times more detrimental compared to the RANK economy ($\gamma_Y^H = 0, \mu_y = 1$), as shown Figure 4.

¹⁴Figure 6 shows that the marginal cost of trend inflation, denoted by $\left| \frac{\partial \Gamma^H}{\partial \Pi} \right|$, is higher for HtMs under the baseline calibration ($\mu_y = 1.75$) than in the RANK economy ($\mu_y = 1$).

Figure 7: Optimal Inflation Target



Notes. The figure displays the optimal inflation target as a function of cyclical and long-run inequality for savers (left) and hand-to-mouth spenders (middle), along with the aggregate optimal inflation target (right).

case, a higher inflation target becomes optimal to reduce the welfare cost associated with output and inflation volatility. Given cyclical inequality, a higher level of long-run inequality (i.e., an increase in μ_y) decreases the optimal inflation target for all households. As discussed in Section 3.1.2, greater steady-state inequality reduces the HtMs' share in aggregate consumption, thereby diminishing the contribution of their consumption to aggregate dynamics. Accordingly, the variance of output and inflation decreases, leading to a lower optimal inflation target.

We turn to the role of income risk in shaping the optimal inflation target. Columns (2) and (3) of Table 3 report the optimal inflation target with and without risk under the baseline calibration of cyclical and income inequality. Incorporating income risk further raises the optimal inflation target at both the individual and aggregate levels. As the income risk triggers the precautionary saving motive and lowers the steady-state nominal and real interest rates, the ZLB binds more frequently. With more frequent ZLB episodes, the variance of inflation and output increases, resulting in a higher optimal inflation target for all households.

Table 3 further reveals that the HtMs' optimal inflation target is higher than that of the Savers under the calibrated cyclical and long-run inequality, regardless of the incorporation of income risk. The reason is that the marginal benefit of trend inflation, which is greater for HtMs, plays a more dominant role in determining the optimal inflation target than the marginal cost. Rows 1 and 2 of Figure 8 illustrate how each welfare component varies with a rise in trend inflation for Savers and HtMs. Consider the THANK economy without income risk (dash-dotted line) as an illustrative example. The figure clearly highlights the significance of marginal benefits. Observing the leftmost panels of rows 1 and 2, the marginal cost of higher trend inflation—captured by the slope of the constant term—is somewhat greater for HtMs relative to Savers. However, as indicated

Table 3: Optimal Inflation Target

	(1) RANK	(2) THANK (no income risk)	(3) THANK (income risk)
Saver (% p.a)	1.49	1.63	1.80
HtM (% p.a)	1.49	2.06	2.17
Aggregate (% p.a)	1.49	1.82	1.95

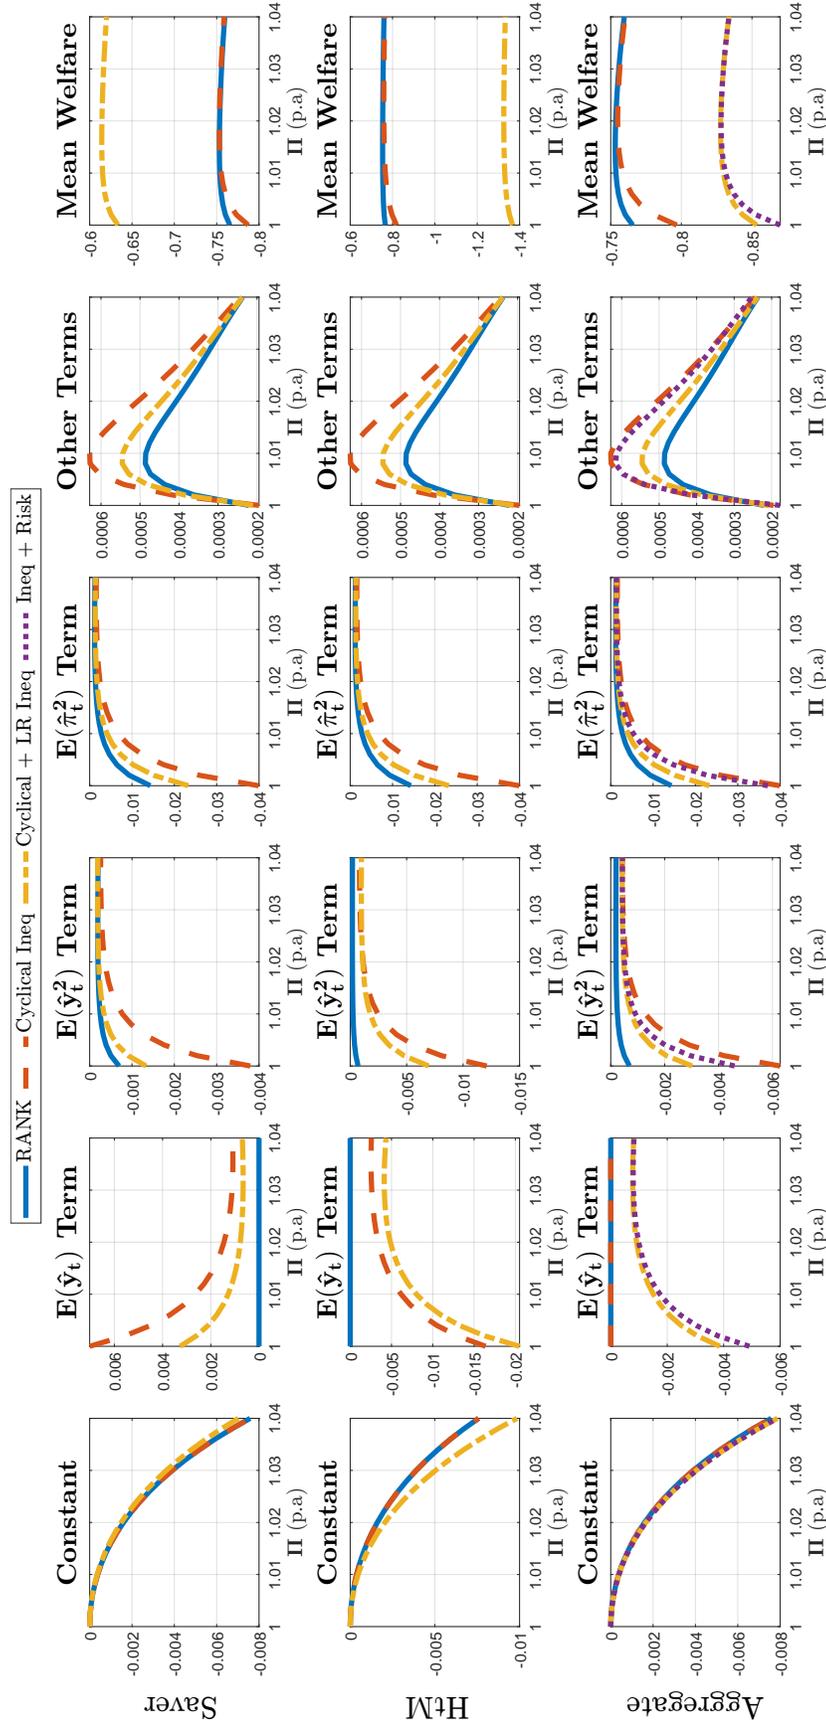
by the mean and variance of the output terms, higher trend inflation yields significantly greater welfare gains for HtMs than for Savers. The HtMs' higher optimal inflation target stems from the dominance of these marginal benefits over the marginal costs.

Turning to the aggregate optimal inflation target, Table 3 shows that the optimal inflation rate is 1.49% in RANK, 1.82% in THANK without income risk, and 1.95% in THANK with income risk. The increase in the utilitarian optimal inflation target in THANK relative to RANK is less than 50 basis points and remains close to the commonly observed 2% target in developed countries. The primary reason that the increase in the optimal inflation rate is relatively limited can be found in the third row of Figure 8. As shown in the first panel, the marginal cost of trend inflation in THANK with (dotted line) and without income risk (dash-dotted line) is comparable to that observed in a RANK economy (solid line). This occurs because the HtMs' larger marginal cost and the Savers' smaller marginal cost nearly balance out under the calibrated income inequality. Therefore, the difference in optimal inflation rates between THANK and RANK primarily stems from the marginal benefit side. Observing the mean output, output variance, and inflation variance terms shown in the remaining panels of row 3, it is evident that the marginal benefit of trend inflation—captured by the slope of the mean output, output variance, and inflation variance terms—is higher in THANK, justifying a higher inflation target. However, as can be understood from comparing THANK without income risk (dash-dotted line) against the economy without income risk and without long-run inequality (dashed line), long-run income inequality reduces the slope of output and inflation volatility, dampening the marginal welfare gains from higher trend inflation. Because HtMs' income and consumption dynamics play a smaller role in driving the aggregate dynamics under long-run inequality, the aggregate demand fluctuation is less pronounced. The presence of long-run inequality makes the increase in the aggregate optimal inflation rate in THANK relatively modest.

State-dependent pricing The Calvo pricing that we adopt assumes fixed price rigidity, which is inconsistent with empirical evidence showing a positive relationship between price flexibility and inflation.¹⁵ In this framework, steady-state price dispersion increases mechanically with trend inflation, which may overstate the marginal costs of trend inflation. One way to capture the positive relationship between the frequency of price adjustment and steady state inflation is through a state-dependent pricing model in the spirit of [Dotsey, King and Wolman \(1999\)](#). How would the

¹⁵See [Blanco et al. \(2024\)](#) and the references therein for supporting evidence.

Figure 8: Welfare Decomposition



Notes. The figure displays the steady-state component, the mean output component, the output variance component, the inflation variance component of the welfare function, and the other terms (the sum of mean inflation and the covariance terms), each conditional on the trend inflation rate. The label 'RANK' (solid line) refers to the benchmark economy without inequality or income risk ($\gamma_y^H = 0, \mu_y = 1, s = 1$). 'Cyclical Ineq' (dashed line) represents an economy with only cyclical inequality ($\gamma_y^H = 1.5, \mu_y = 1, s = 1$). 'Cyclical + LR Ineq' (dash-dotted line) corresponds to an economy with both cyclical and long-run inequality ($\gamma_y^H = 1.5, \mu_y = 1.75, s = 1$), while 'Ineq + Risk' includes both forms of inequality along with income risk ($\gamma_y^H = 1.5, \mu_y = 1.75, s = 0.9998$). All lines in the steady state component start at zero to make the slopes across lines comparable.

utilitarian optimal inflation target in THANK change under state-dependent pricing? To address this question, we first discuss how the optimal inflation rate differs under state-dependent pricing relative to Calvo pricing in RANK. We then conjecture that this difference is more pronounced in THANK because endogenous price flexibility mitigates the cost of mean output loss faced by HtMs.

In RANK with state-dependent pricing, higher trend inflation affects household welfare through three channels that are absent under Calvo pricing. First, state-dependent pricing mitigates the increase in the steady-state cost of price dispersion as trend inflation rises. From the perspective of the linear-quadratic welfare function, this effect appears as a smaller increase in the constant term. Firms face a fixed adjustment cost whenever they change prices and adjust only when the gap between the current price and the optimal price exceeds a threshold. As steady-state inflation rises, firms reach this threshold more quickly. Consequently, when the adjustment cost is relatively small, price adjustments occur more frequently than under Calvo pricing, which reduces steady-state price dispersion.

Second, state-dependent pricing implies a smaller increase in the welfare cost of a given level of inflation volatility than Calvo pricing when trend inflation rises. In the linear-quadratic welfare function, this appears as a smaller increase in the coefficient on the variance of inflation. Intuitively, because more firms adjust prices in response to an inflationary shock than under Calvo pricing, the resulting distortion in relative prices becomes less costly.

Third, higher trend inflation implies greater price flexibility under state-dependent pricing than under Calvo pricing, because firms adjust prices more frequently. This causes inflation to fall more in response to a given adverse demand shock, increasing the likelihood of hitting the ZLB. In the linear-quadratic welfare function, this effect is captured by higher inflation variance.

The first and second channels reduce the cost of raising trend inflation and put upward pressure on the optimal inflation target under state-dependent pricing relative to Calvo pricing. The third channel reduces the benefit of increasing trend inflation and puts downward pressure on the optimal inflation target. In the menu cost model of [Blanco \(2021\)](#), in which household income inequality is absent, the first channel dominates, yielding a higher optimal inflation rate than in the Calvo model.

In THANK with state-dependent pricing, the three channels described above persist. Therefore, state-dependent pricing still yields a higher optimal inflation target than Calvo pricing. However, we conjecture that the increase in the optimal inflation target is larger in THANK than in RANK because of an additional welfare term that interacts with endogenous price stickiness: the welfare cost of mean output loss faced by HtMs. This cost appears explicitly in our linear-

quadratic welfare function (14), represented by the mean output term (i.e., $\Gamma_y^H \mathbf{E}(\hat{y}_t)$), and is quantitatively nonnegligible in THANK, while it is essentially zero in RANK.

Recall from Figure 8 that the mean output term in THANK increases with trend inflation. Under state-dependent pricing, we expect this term to increase by more for the following reason. When trend inflation rises under state-dependent pricing, price stickiness declines, making the aggregate supply curve steeper. With a steeper aggregate supply curve, the output loss in response to an adverse demand shock is smaller. As a result, higher trend inflation generates welfare gains for HtMs through two channels. First, it reduces the frequency of the ZLB. Second, it mitigates output losses during ZLB episodes due to the steeper aggregate supply curve. Both channels reduce mean output losses and thereby increase the HtMs' mean output term and welfare. The second channel, which is absent under Calvo pricing, implies a larger marginal benefit of raising trend inflation, resulting in a larger increase in the optimal inflation target than in RANK.

3.3 Role of Private Debt

In the baseline THANK model, we assume zero wealth and therefore abstract from the channel through which inflation or deflation affects asset and liability positions, as documented by [Doepke and Schneider \(2006\)](#). In HANK models with wealth inequality, inflation affects aggregate consumption dynamics through redistribution among households with different nominal asset and liability positions. To explore the extent to which this channel matters for the optimal inflation target, we introduce private debt into the THANK model following the structure in [Bilbiie, Monacelli and Perotti \(2013\)](#).

Recall that Z_t^S and Z_t^H denote the amount of loans held by Savers and the amount of borrowing by HtMs, respectively, at the end of period t . Savers make loans to HtMs, who always borrow up to the limit $-\bar{B}^p$ (i.e., $Z_t^H = -\bar{B}^p$). The private bond market clearing condition is:

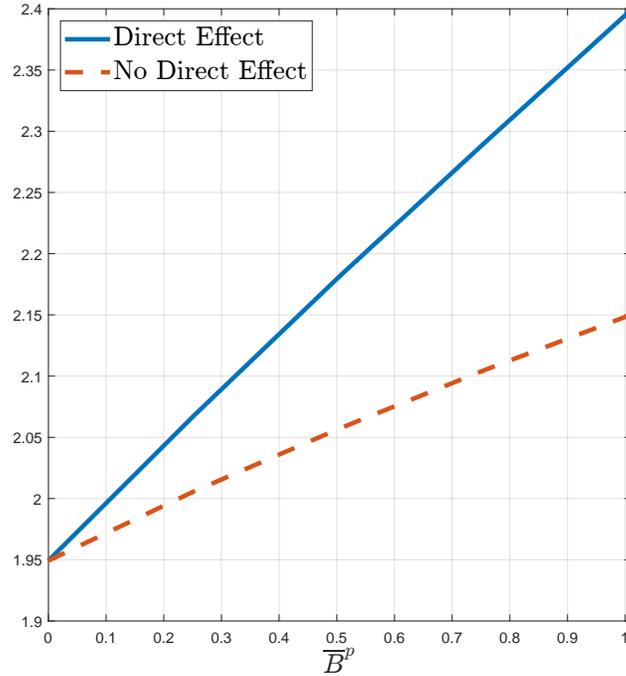
$$(1 - \lambda)Z_t^S + \lambda Z_t^H = 0,$$

which implies $Z_t^S = \frac{\lambda}{(1-\lambda)}\bar{B}^p$. Because net government bond holdings remain zero, as in the baseline model without private debt, wealth inequality in this setting is determined by the heterogeneous private bond positions. After type-switching occurs, the beginning-of-period $t + 1$ per capita bond holdings for each type become:

$$\begin{aligned} B_t^S &= s \frac{\lambda}{(1-\lambda)} \bar{B}^p - (1-s)\bar{B}^p \\ B_t^H &= (1-h) \frac{\lambda}{1-\lambda} \bar{B}^p - h\bar{B}^p, \end{aligned} \tag{15}$$

where s denotes the probability of remaining a Saver and h denotes the probability of remaining Hand-to-Mouth. For the calibration of the debt limit, we set it equal to average labor income,

Figure 9: Aggregate Optimal Trend Inflation and Private Debt



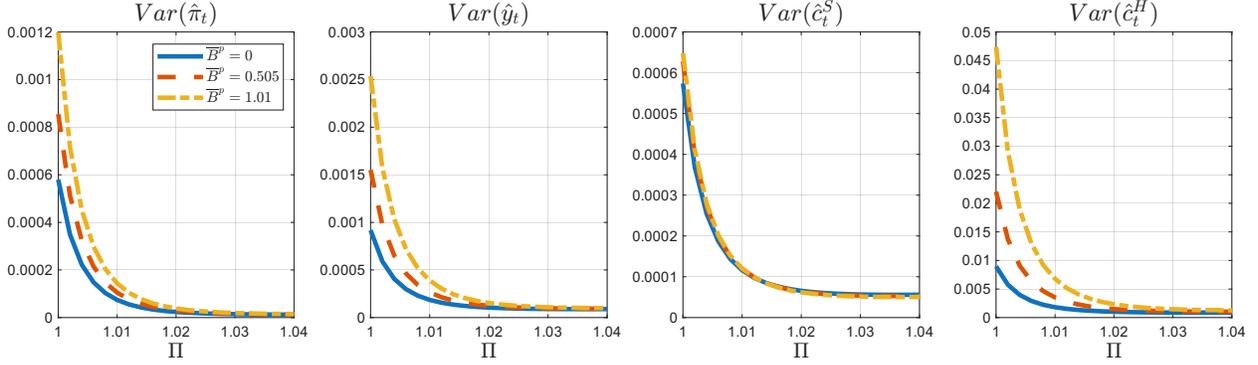
Notes. The figure displays the utilitarian optimal trend inflation as a function of the borrowing limit \bar{B}^p .

following Kaplan, Moll and Violante (2018), which gives $\bar{B}^p = 1.01$. Clearly, since $s = 0.9998$ and $h = 0.9995$, Savers hold a positive wealth position while HtMs hold a negative one (i.e., $B_t^S > 0$ and $B_t^H < 0$).

The solid line in Figure 9 shows how the utilitarian optimal inflation target varies with the degree of wealth inequality, represented by the borrowing limit.¹⁶ The utilitarian optimal trend inflation increases with wealth inequality and reaches 2.40% under the calibrated debt limit. This value exceeds the optimal trend inflation of 1.95% in the absence of wealth inequality. Two forces explain why the optimal trend inflation rises with the debt level. First, a risk-premium shock (i.e., a positive shock to q_t) increases the debt-servicing cost faced by HtMs and raises the return on loans received by Savers. Because this redistribution affects the consumption of constrained HtMs more than that of Savers, aggregate consumption and output fall. Second, when nominal policy rates are zero during ZLB episodes, deflation raises the real cost of debt (i.e., $\frac{R_{t-1}q_{t-1}}{\Pi_t}$) through Fisherian debt deflation. This increases the burden on indebted HtMs and raises the real return on loans for Savers. The heightened real cost of debt disproportionately affects debt-burdened HtMs, causing aggregate consumption and output to fall further compared with the case without wealth inequality. Both forces increase the marginal benefits of raising trend inflation.

¹⁶The utilitarian quadratic welfare function under private debt, which we optimize, is in Appendix B. Unlike in the baseline economy without wealth, it is not tractable to derive an analytical quadratic welfare function for each household type when wealth inequality is present. Thus, we derive only its utilitarian counterpart and compute the utilitarian optimal inflation target.

Figure 10: Private Debt and Second Moments



Notes. The figure displays the variance of inflation, output, Savers' consumption, and HtMs' consumption for different levels of the borrowing limit \bar{B}^p .

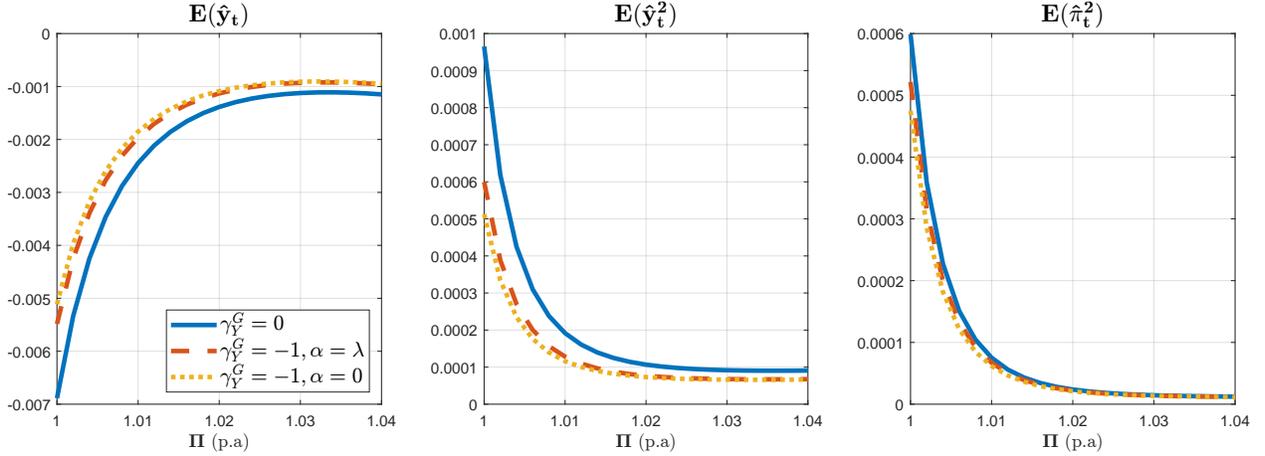
Figure 10 confirms the mechanism behind the higher optimal trend inflation induced by wealth inequality by presenting the volatility of inflation, output, and consumption. A higher debt level raises the volatility of inflation and output. This result reflects the direct effect of the risk-premium shock and Fisherian debt deflation, which increase consumption volatility for HtMs substantially, as shown in the fourth panel. Although Savers benefit in wealth from the direct effect of the risk-premium shock and the debt deflation channel during ZLB episodes, their consumption volatility is almost unchanged across debt levels. The reason is that while risk-premium shocks and deflation increase their real wealth in the liquidity trap, the fall in output caused by the more adversely affected HtMs offsets the rise in consumption from higher wealth.

Notice that the dashed line in Figure 9 shuts down the direct effect of the risk-premium shock on Savers' and HtMs' wealth by levying a lump-sum tax equal to $\frac{R_{t-1}}{\Pi_t} (q_{t-1} - 1) B_{t-1}^j$ ($j \in \{S, H\}$) on both household types. Accordingly, the dashed line captures the optimal trend inflation that is purely implied by the Fisherian debt deflation channel. Not surprisingly, when the direct effect of the risk-premium shock is absent, the aggregate optimal inflation target is lower, but it still increases with the level of the borrowing limit.

4 Fiscal Policy and Optimal Inflation Target

Until now, we have assumed the absence of endogenous fiscal stabilization (i.e., $\gamma_Y^G = \gamma_Y^T = 0$). However, during the Global Financial Crisis and the COVID-19 pandemic, fiscal authorities deployed large stimulus packages to stabilize economic downturns. When countercyclical fiscal policy is present, the optimal inflation rate is affected not only by income inequality but also by the degree of fiscal intervention. In this section, we examine the mechanism through which countercyclical fiscal policy affects the optimal inflation target and then discuss whether fiscal policy

Figure 11: Business Cycle Component: Aggregate Moments



Notes. The figure shows the first and second moments of inflation and output as functions of the level of trend inflation under three government spending regimes: an economy without cyclical government spending ($\gamma_Y^G = 0$); an economy with countercyclical fiscal spending where the tax burden is distributed equally across the household types ($\gamma_Y^G = -1, \alpha = \lambda$); and an economy where Savers bear the entire tax burden ($\gamma_Y^G = -1, \alpha = 0$).

can substitute for a higher inflation target in an environment with recurrent ZLBs.¹⁷

We focus on two fiscal measures: government spending, G_t , and targeted transfers to hand-to-mouth households, \mathcal{T}_t^H . First, government spending is characterized by two parameters (γ_Y^G, α), where $\gamma_Y^G < 0$ governs how aggressively government spending expands during recessions, and α determines tax progressivity. Specifically, $\alpha = 0$ corresponds to maximum tax progressivity, with Savers absorbing the entire cyclical tax burden from fiscal spending, while $\alpha = \lambda$ implies that the tax burden is distributed equally across the population. Next, the transfer policy is characterized by a single parameter $\gamma_Y^T < 0$, which governs the amount of transfers given to HtMs during recessions.

4.1 Government Spending and Optimal Trend Inflation

We examine the interplay between cyclical government spending and the optimal trend inflation rate by analyzing how government spending affects the cost-benefit trade-off of trend inflation. To maintain consistency with the structure of the previous section, we begin by investigating the impact of government spending on each welfare component, defined as the product of aggregate moments and their respective coefficients.

4.1.1 Business Cycle Component: Aggregate Moments

Figure 11 illustrates how a rise in trend inflation affects the first and second moments of output and inflation in three THANK economies characterized by distinct government spending policies:

¹⁷In Appendix C, we compute the jointly optimal mix of the inflation target and the countercyclicality of fiscal policy.

an economy without cyclical government spending (solid line), an economy with countercyclical fiscal spending in which the tax burden is distributed equally across the population (dashed line), and an economy in which Savers fully bear the total tax burden (dotted line). The effect of cyclical government spending on aggregate moments can be summarized in two main points: (1) countercyclical use of government spending reduces mean output losses and aggregate volatility; (2) for a given amount of government spending, more progressive taxation leads to a more reduction in mean output losses and aggregate volatility.

The first point can be observed in Figure 11, which shows that introducing countercyclical government spending reduces mean output losses and the volatility of inflation and output for a given level of trend inflation. This outcome arises because public demand offsets the contraction in private demand during recessions, thereby lowering the frequency and severity of ZLB episodes. The second point can be inferred from the observation that mean output losses and the variances of output and inflation are smaller when the tax burden is fully borne by Savers, compared to when it is equally shared across both household types. This result stems from the fact that HtMs, who exhibit much higher MPCs, are exempt from the tax burden and therefore experience a greater increase in disposable income when government spending rises, compared to a scenario where they face taxation. In conclusion, introducing countercyclical government spending mitigates the severity of recessions, thereby reducing the extent to which higher trend inflation is needed to improve aggregate moments.

4.1.2 Business Cycle Component: Coefficients

Cyclical government spending also influences households' weights on aggregate moments through the countercyclical tax burden. Under a balanced budget, the effect of government spending on the welfare function coefficients depends on how the tax burden is distributed. The following proposition illustrates the heterogeneous effects of cyclical government spending on the coefficients.

Proposition 4. *Assume no cyclical transfer ($\gamma_Y^T = 0$) and $\alpha \in [0, \lambda]$. Then:*

1. $\frac{\partial \Gamma_y^S}{\partial \gamma_Y^G} < 0$, and $\frac{\partial \Gamma_y^H}{\partial \gamma_Y^G} \leq 0$, with equality holding at $\alpha = 0$.
2. If $(1 - \lambda)\mu_y - \lambda\gamma_Y^H$, then $\frac{\partial \Gamma_{y^2}^S}{\partial \gamma_Y^G} > 0$ and $\frac{\partial \Gamma_{y^2}^H}{\partial \gamma_Y^G} \geq 0$, with equality holding at $\alpha = 0$.

Proof. Refer to Appendix A. □

The first statement of Proposition 4 means that more countercyclical government spending (i.e., more negative γ_Y^G) makes households more painful from mean output losses. Since countercyclical government spending requires countercyclical taxation, a negative mean output deviation implies a higher average tax burden compared to the steady state. Consequently, a given mean

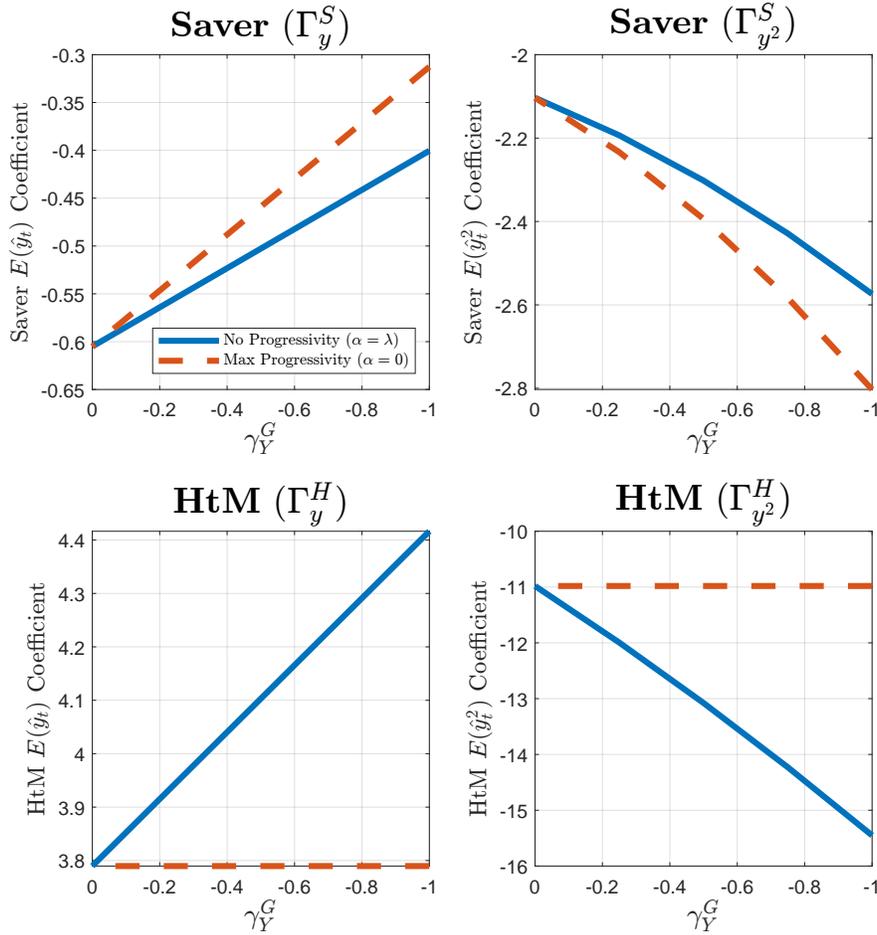
output loss under countercyclical government spending leads to a larger reduction in mean disposable income than in the absence of such policy, thereby worsening household welfare. In the extreme case where taxation is fully progressive ($\alpha = 0$), the mean output coefficient for HtMs remains unchanged, as they are entirely exempt from the cyclical tax burden. The second statement shows that, under the mild regularity condition $(1 - \lambda)\mu_y - \lambda\gamma_Y^H > 0$, more countercyclical government spending increases the welfare cost of output volatility for all households.¹⁸ Since more aggressive stabilization requires a more countercyclical tax burden, households experience greater volatility in disposable income for a given level of output volatility. When $\alpha = 0$, the coefficient for HtMs remains unchanged despite the increased cyclicity of government spending, as they do not bear any tax burden.

Figure 12 illustrates how the coefficients on the mean and variance of output respond to changes in the cyclicity of government spending, given $\gamma_Y^H = 1.5$, $\mu_y = 1.75$, and a 2% trend inflation rate. The solid line represents the coefficient when the tax burden is equally distributed across all agents. Consistent with Proposition 4, the coefficient on mean output increases, while the coefficient on output variance decreases as government spending becomes more countercyclical. Moreover, the change in coefficients for a given increase in γ_Y^G is significantly larger for HtMs than for Savers. This indicates that both mean output losses and output volatility are more detrimental to HtMs. This arises because HtMs, having lower steady-state income, face a steeper marginal utility curve. As a result, the increase in mean tax burden and the higher volatility in disposable income due to more countercyclical taxation impose greater welfare costs on HtMs than on Savers. When tax progressivity is at its maximum ($\alpha = 0$), the coefficient changes for Savers become more pronounced than under minimal progressivity ($\alpha = \lambda$), while the HtMs' coefficients become invariant to the cyclicity of government spending. This is because, under full tax progressivity, the entire cyclical tax burden falls on Savers. Consequently, they face a greater increase in the mean tax burden for a given output loss and a greater increase in income volatility for a given output variance as γ_Y^G becomes more negative. In contrast, HtMs' full exemption from taxation renders their coefficients unaffected by changes in γ_Y^G .

Drawing on the insights discussed in this section, one might wonder whether the introduction of cyclical government spending reduces the need for a higher inflation target. The answer may be positive, as endogenous fiscal stabilization mitigates the severity of the ZLB, thereby reducing mean output losses and macroeconomic volatility. However, since cyclical government spending is accompanied by a countercyclical tax burden, a given level of mean output loss and output volatility becomes more costly. This, in turn, may lead households to support a higher inflation target. Therefore, whether cyclical government spending decreases the optimal inflation target depends on whether the welfare gains from improved aggregate moments outweigh the welfare

¹⁸This regularity condition implies that γ_Y^H is not excessively large, so that Savers' consumption still responds positively to increases in output.

Figure 12: Role of Government Spending on Coefficients



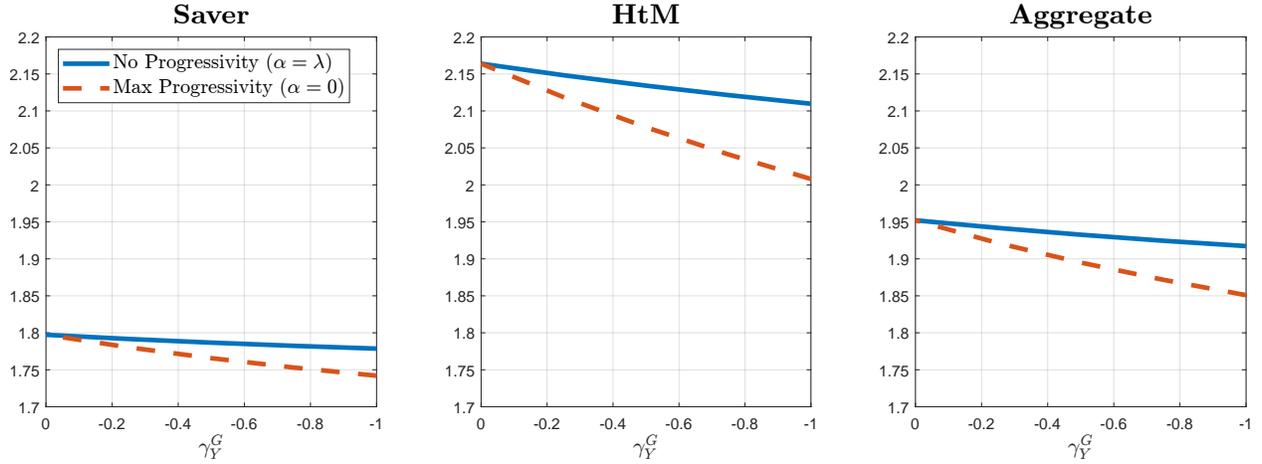
Notes. The figure displays the coefficients on the mean and variance of output as functions of the cyclicity of government spending, given $\gamma_Y^H = 1.5$, $\mu_y = 1.75$, and a 2% trend inflation rate, under two scenarios: an economy with minimal tax progressivity ($\alpha = \lambda$) and an economy with maximal tax progressivity ($\alpha = 0$).

losses caused by changes in the coefficients. We examine this quantitative question below.

4.1.3 Optimal Inflation Rate

In this section, we compute the optimal trend inflation rate under varying degrees of government spending cyclicality financed through taxation. We focus on two polar cases of tax financing: maximal tax progressivity ($\alpha = 0$) and minimal tax progressivity ($\alpha = \lambda$). Figure 13 illustrates how the cyclicity of government spending affects optimal inflation rates at both the individual and aggregate levels under maximal (dashed line) and minimal (solid line) tax progressivity. The case of $\gamma_Y^G = 0$ corresponds to the baseline calibrated THANK economy without cyclical fiscal spending, as discussed in Section 3. The findings can be summarized in two points: (1) countercyclical government spending lowers the optimal inflation rate at both the individual and aggregate levels, and (2) the extent of this reduction is more pronounced under a more progressive tax system.

Figure 13: Optimal Trend Inflation and Cyclical Government Spending



Notes. The figure displays the optimal trend inflation rate as a function of the cyclical government spending, given $\gamma_Y^H = 1.5$ and $\mu_y = 1.75$, under two scenarios: an economy with minimal tax progressivity ($\alpha = \lambda$) and an economy with maximal tax progressivity ($\alpha = 0$).

The first point is evident from the declining pattern of the optimal inflation rate with respect to cyclical government spending in Figure 13. It suggests that more countercyclical government spending reduces the marginal benefit of increasing trend inflation for all households. Since the steady-state cost of trend inflation is unaffected by the cyclical government spending, the reduction in the optimal inflation target is entirely driven by the decline in marginal benefits. Specifically, countercyclical government spending improves household welfare primarily by reducing mean output losses and macroeconomic volatility, thereby diminishing the benefit of raising the trend inflation rate.

The second point can be confirmed by comparing the two scenarios in Figure 13. The optimal inflation rate for all households is lower when Savers bear the full tax burden ($\alpha = 0$) than when the tax burden is equally shared between Savers and HtMs ($\alpha = \lambda$). Since government spending financed through more progressive taxation leads to reduced aggregate volatility and a smaller tax burden for HtMs during recessions, the marginal benefit of raising the inflation target is smaller for HtMs under a more progressive tax system. Interestingly, Savers' optimal inflation rate also decreases more under progressive taxation. Unlike HtM households, government spending financed through more progressive taxation exerts two opposing effects on the marginal benefit of trend inflation for Savers. On the one hand, reduced mean output losses and lower aggregate volatility, shown in Figure 11, diminish the marginal benefit of trend inflation. On the other hand, the higher tax burden increases the welfare cost of a given level of output volatility, reflected in a more negative $\Gamma_{y^2}^S$, thereby raising the marginal benefit of trend inflation. Under the baseline calibration, the former effect dominates the latter, leading Savers to prefer a lower inflation target.

Finally, consistent with the individual-level optimal trend inflation, the utilitarian optimal inflation rate, shown in the right panel of Figure 13, also declines with more countercyclical government spending and with more progressive taxation.

4.1.4 Role of Government Debt

We now allow government spending to be financed not only by lump-sum taxes but also by government debt and examine how the optimal inflation rate changes when government debt is introduced.

As in Bilbiie (2024), we assume that, at the end of period t , only Savers are allowed to purchase government bonds, with Z_t^S denoting the amount of bonds they purchase. In contrast, HtMs are always constrained, so the amount of bonds they purchase is $Z_t^H = 0$. Thus, the bond market clearing condition is:

$$b_t = (1 - \lambda)Z_t^S + \lambda Z_t^H = (1 - \lambda)Z_t^S, \quad (16)$$

where b_t denotes the stock of government bonds supplied at the end of period t . Given that s and h denote the probability of remaining a Saver and the probability of remaining Hand-to-Mouth, respectively, the beginning-of-period $t + 1$ per capita bond holdings for each type after type-switching are:

$$B_t^S = sZ_t^S = \frac{s}{1 - \lambda}b_t \quad (17)$$

$$B_t^H = (1 - h)Z_t^S = \frac{1 - h}{1 - \lambda}b_t = \frac{1 - s}{\lambda}b_t, \quad (18)$$

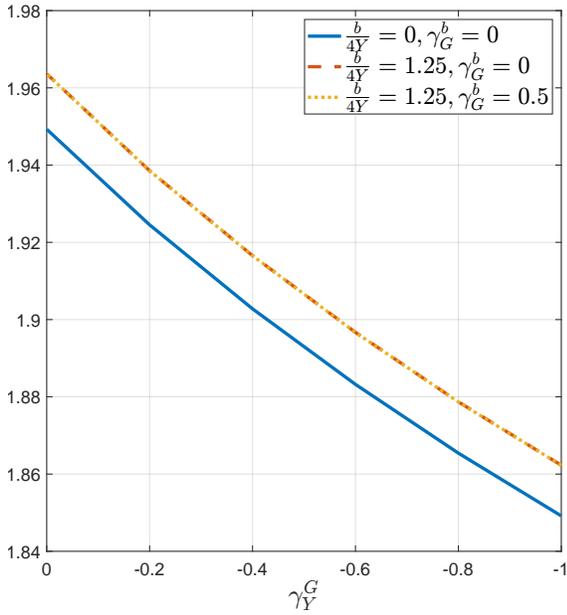
where the second equality uses the bond market clearing condition (16). These expressions represent the wealth positions of Savers and HtMs at the beginning of period $t+1$. Given that $s = 0.9998$, Savers clearly hold a larger amount of wealth than HtMs (i.e., $B_t^S > B_t^H$), implying that Savers receive a larger real interest income when nominal interest rates are zero and inflation falls. Accordingly, this wealth heterogeneity widens the consumption gap between Savers and HtMs at the ZLB.

With public debt, the government budget constraint becomes:

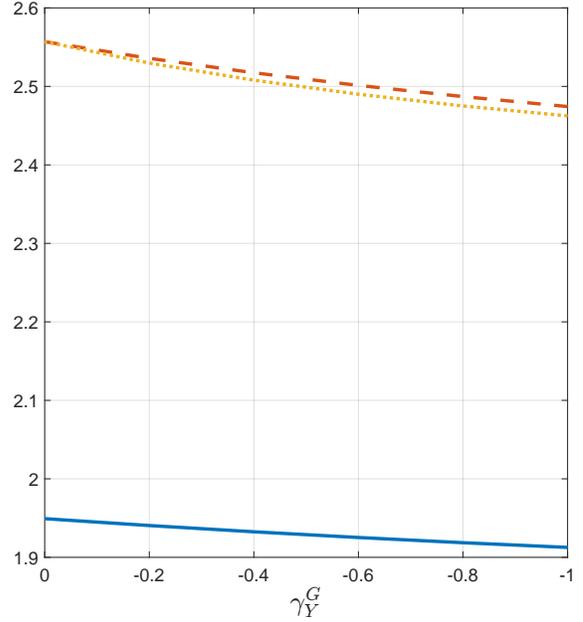
$$G_t + \frac{R_{t-1}}{\Pi_t}b_{t-1} = (1 - \lambda)TG_t^S + \lambda TG_t^H + b_t.$$

Government spending on government consumption, G_t , and the real interest payment on outstanding debt, $\frac{R_{t-1}}{\Pi_t}b_{t-1}$, are financed by lump-sum taxes collected from households, TG_t^S and TG_t^H , and by issuing new debt, b_t . Bond issuance is proportional to government consumption,

Figure 14: Aggregate Optimal Trend Inflation and Government Debt



(a) Maximal Tax Progressivity ($\alpha = 0$)



(b) Minimal Tax Progressivity ($\alpha = \lambda$)

Notes. The figure displays the utilitarian optimal trend inflation as a function of the cyclical government spending for different levels of the government debt-to-GDP ratio, with the left panel corresponding to maximal tax progressivity ($\alpha = 0$) and the right panel corresponding to minimal tax progressivity ($\alpha = \lambda$).

following:

$$b_t = b + \gamma_G^b (G_t - G),$$

where $\gamma_G^b \geq 0$ measures the sensitivity of bond issuance to cyclical government consumption. Tax allocation across household types is determined by:

$$TG_t^S - TG^S = \frac{1 - \alpha}{1 - \lambda} (G_t - G) + \frac{1 - \alpha}{1 - \lambda} \left(\frac{R_{t-1}}{\Pi_t} b_{t-1} - b_t - \left(\frac{R}{\Pi} - 1 \right) b \right),$$

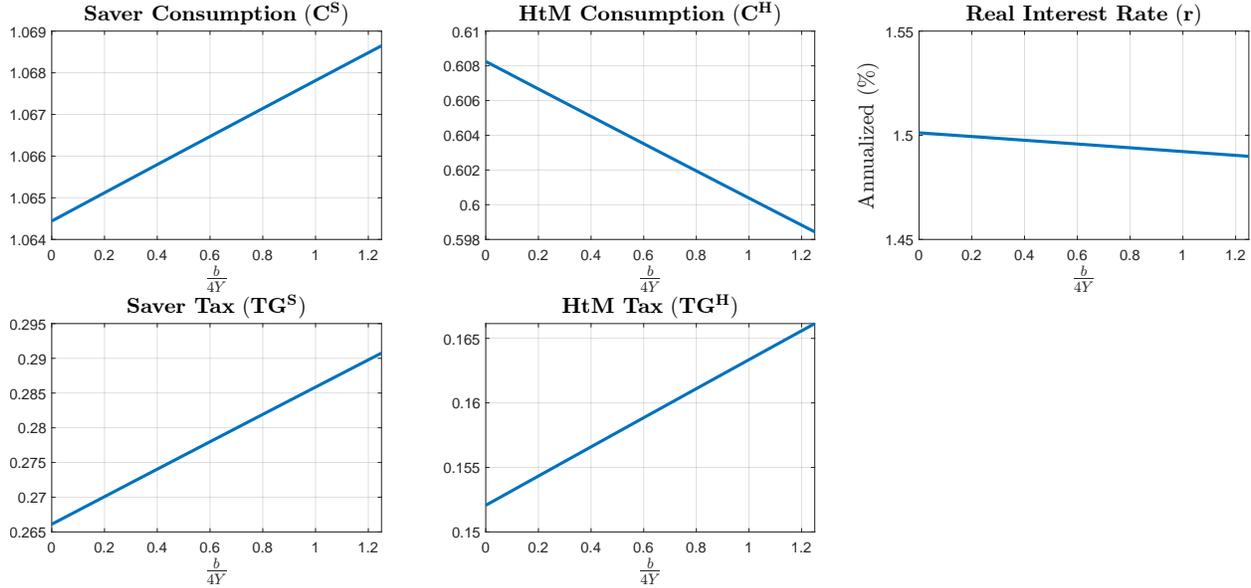
where α governs the share of the tax burden borne by Savers for both cyclical government consumption and net repayment on public debt. When $\alpha = 0$, Savers bear the entire tax burden, whereas $\alpha = \lambda$ implies an equal tax burden on Savers and HtMs. We calibrate the steady-state level of public debt to match the ratio of U.S. federal debt to annual GDP, which is 1.25 as of 2025. We then examine whether the presence of government debt affects the result that “countercyclical fiscal policy lowers the optimal inflation target” under two different tax schemes: maximal ($\alpha = 0$) and minimal ($\alpha = \lambda$) tax progressivity. The utilitarian welfare function used to compute the optimal inflation rate under positive government debt is the same as in the case with private debt and is listed in Appendix B.

Maximal tax progressivity Figure 14(a) shows how the effect of countercyclical government spending on the utilitarian optimal inflation rate depends on the level of public debt when Savers bear the full tax burden. Comparing the no-debt case (solid line) with the positive-debt case (dashed line), it is clear that, regardless of the debt level, more countercyclical government consumption lowers the optimal inflation rate. This shows that the welfare-improving mechanism through which countercyclical government consumption reduces mean output losses and macroeconomic volatility operates regardless of the level of public debt.

Notice that, even in the case of positive debt with cyclical government debt issuance (dotted line), more countercyclical government purchases lower the optimal inflation rate. This result indicates that the welfare-improving mechanism of countercyclical fiscal policy through stabilization remains robust when debt dynamics are included. Intuitively, since Savers bear the entire tax burden, issuing more government debt instead of collecting lump-sum taxes to finance countercyclical government purchases does not change the present value of taxes faced by Savers and therefore does not affect their current consumption. As a result, the source of government revenue used to fund government purchases does not affect the extent to which countercyclical fiscal policy lowers the optimal inflation target.

As a secondary finding, a comparison between the no-debt case (solid line) and the positive-debt case (dashed line) shows that increasing the steady state level of government debt raises the optimal inflation rate modestly for a given degree of government spending cyclicality. The main reason is a rise in steady-state consumption dispersion across household types, which leads to more frequent ZLB episodes for a given trend inflation rate due to stronger precautionary saving by Savers. To be specific, a higher steady-state level of government debt affects consumption inequality through two channels. First, it requires higher steady-state interest payments by the government, which create additional financing needs. As a result, steady-state lump-sum taxes, TG^S and TG^H , increase and reduce steady-state consumption for both household types. The bottom two panels of Figure 15 show the rise in lump-sum taxes as government debt increases. Second, a positive asset position generates steady-state interest income that raises consumption for both groups. For Savers, who hold substantial amounts of government bonds, the increase in interest income more than offsets the rise in taxes, leading to higher steady-state consumption. For HtMs, who hold only small amounts of government bonds, the tax burden dominates the interest income effect, resulting in lower steady-state consumption. This widening steady-state consumption gap implies that Savers experience a larger consumption loss when they become HtMs with probability $1 - s$. Accordingly, their precautionary savings motive strengthens, lowering the steady-state real interest rate. The top panels of Figure 15 show the increasing consumption gap and falling real interest rate as steady-state government debt rises. A lower steady-state real interest rate is associated with more frequent ZLB episodes, implying greater marginal benefits from raising trend inflation.

Figure 15: Government Debt and Steady-State Values

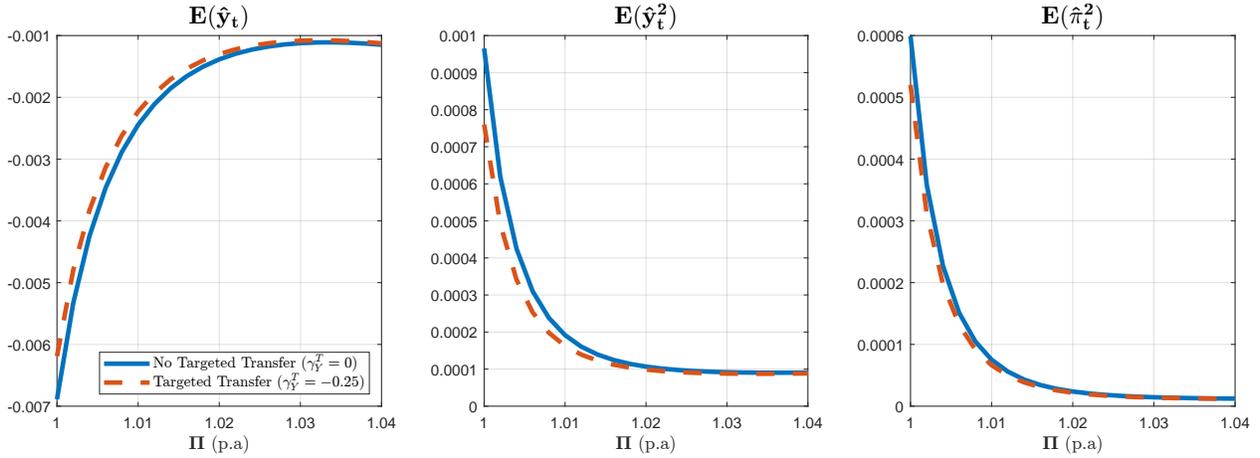


Notes. The figure displays steady-state consumption, real interest rates, and lump-sum taxes for different levels of the steady-state government debt-to-GDP ratio under maximal tax progressivity ($\alpha = 0$).

Minimal tax progressivity Now assume that HtMs bear the tax burden from government spending on consumption and net repayment on public debt (i.e., $\alpha = \lambda$). Comparing the no-debt case (solid line) with the positive-debt case (dashed line) in Figure 14(b) reveals that, even under minimal tax progressivity, countercyclical fiscal policy lowers the optimal inflation target regardless of the debt level. This indicates that, even in the presence of positive government debt, the stabilizing effect of countercyclical government purchases dominates the destabilizing effect of the tax burden on HtMs. The positive-debt case with cyclical government debt issuance (dotted line) shows that countercyclical government purchases lower the optimal inflation rate slightly more than in the constant-debt case (dashed line). This is because the decline in HtMs' current disposable income at the ZLB becomes smaller when tax financing is partially replaced with debt financing. As a result, the drop in output at the ZLB is smaller than under constant debt, leading to a lower optimal inflation rate. However, the quantitative importance of this mechanism is negligible.

In addition, as in the case where Savers bear the entire cyclical tax burden, increasing the steady-state level of government debt for a given level of government consumption cyclicity raises the optimal inflation rate, but with a larger magnitude. The primary mechanism through which positive government debt raises the optimal inflation target under equal taxation is the HtMs' increased tax obligation to finance net repayment on public debt at the ZLB, which amplifies their income volatility. When the ZLB binds, ex-post real rates rise because inflation falls. The higher real rate implies increased government debt servicing costs, which require higher lump-

Figure 16: Role of Targeted Transfers: Aggregate Moments



Notes. The figure shows the first and second moments of inflation and output as functions of the level of trend inflation under two transfer regimes: an economy without cyclical transfers ($\gamma_Y^T = 0$) and an economy with cyclical transfers ($\gamma_Y^T = -0.25$).

sum taxes, TG_t^S and TG_t^H . While Savers can partially offset the decline in temporary disposable income, constrained HtMs experience a sharp fall in consumption. As a result, aggregate output contracts more severely in an economy with a higher steady-state level of government debt, leading to a higher optimal inflation rate.

4.2 Targeted Transfer and Optimal Trend Inflation

We now turn to the second countercyclical fiscal policy: targeted transfers. Similar to cyclical government spending, transfer policies alter mean output losses, macroeconomic volatility, and the weight households assign to these aggregate moments, thereby changing the marginal benefit of trend inflation. To understand the mechanisms through which transfers alter the cost-benefit trade-off of trend inflation, we begin by analyzing how transfers affect aggregate moments.

4.2.1 Business Cycle Component: Aggregate Moment

Figure 16 compares how welfare-relevant aggregate moments change with trend inflation across two economies: one with cyclical targeted transfers (dashed line) and one without them (solid line). As shown in the figure, cyclical transfers reduce mean output losses and the volatility of both output and inflation across all levels of trend inflation. Thus, similar to the case of cyclical government spending, cyclical transfers reduce the scope for trend inflation to improve aggregate outcomes. Intuitively, transfers from low-MPC Savers to high-MPC HtMs stimulate aggregate demand and income. Due to countercyclical income inequality, HtMs' income rises more than that of Savers, amplifying the increase in aggregate demand. This, in turn, mitigates the severity of ZLB episodes.

4.2.2 Business Cycle Component: Coefficients

As with cyclical government spending, countercyclical transfers affect households' weights on mean output losses and output volatility by changing their disposable income. The following proposition highlights the effects of cyclical transfers on the mean output coefficients.

Proposition 5. *Assume no cyclical government spending ($\gamma_Y^G = 0$), then*

1. $\frac{\partial \Gamma_y^S}{\partial \gamma_Y^T} < 0$ and $\frac{\partial \Gamma_y^H}{\partial \gamma_Y^T} > 0$.
2. If $\frac{1-s_g-2\sigma(1+\gamma_Y^H)}{2\sigma((1-\lambda)\mu_y+\lambda)} < \gamma_Y^T < \frac{(\sigma-\frac{1}{2}(1-s_g))(1-\lambda)\mu_y-\lambda\gamma_Y^H}{\sigma\lambda((1-\lambda)\mu_y+\lambda)}$, $\frac{\partial \Gamma_{y^2}^S}{\partial \gamma_Y^T} > 0$ and $\frac{\partial \Gamma_{y^2}^H}{\partial \gamma_Y^T} < 0$.

Proof. Refer to Appendix A. □

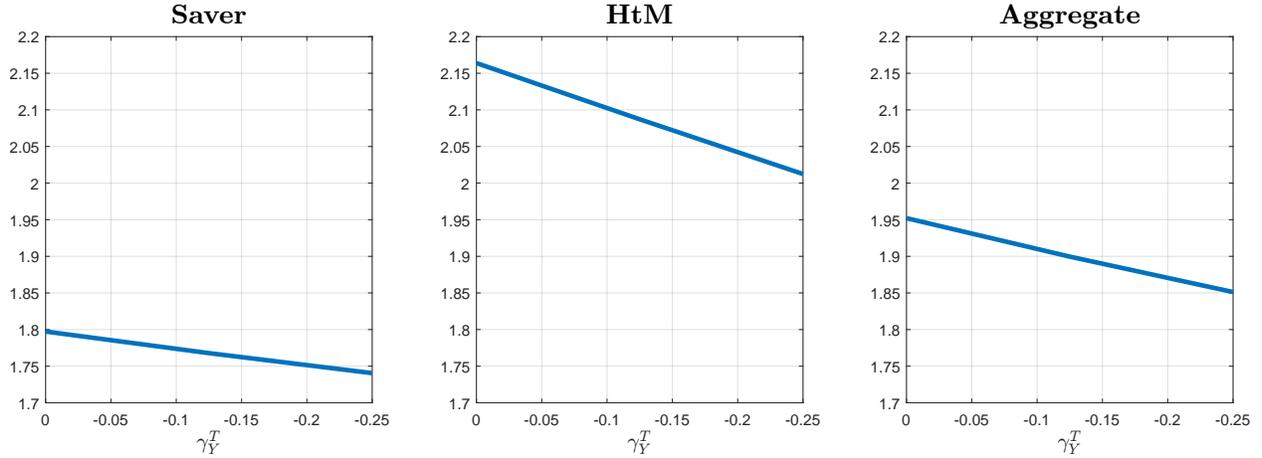
The first statement of Proposition 5 indicates that a more countercyclical transfer policy (i.e., a more negative γ_Y^T) increases the welfare loss from mean output losses for Savers while alleviating it for HtMs. With more countercyclical transfers, the mean output losses stemming from frequent ZLB episodes raise the Savers' average tax burden, thereby reducing their mean disposable income. In contrast, since transfers are directed toward HtMs, they help offset the decline in their mean disposable income, making mean output losses less costly from their perspective. The second statement implies that a more cyclical transfer policy increases the welfare cost of a given level of output volatility for Savers while reducing it for HtMs. This result is intuitive as countercyclical transfers dampen the volatility of disposable income for HtMs but amplify it for Savers, given the same level of output volatility.

Based on our analysis of the effects of transfers on aggregate moments and their coefficients, we can infer how the introduction of countercyclical transfers alters the optimal inflation target. Since countercyclical transfers reduce both the mean output losses and macroeconomic volatility, as well as HtMs' weight on these moments, such transfers would lead HtMs to favor a lower inflation target. For Savers, however, the impact of transfers on the optimal inflation target is more ambiguous. While reduced output volatility and mean output losses would incline Savers toward a lower inflation target, the increased volatility of their disposable income would push them to favor a higher inflation target. Next, we investigate whether cyclical transfers ultimately raise or lower the optimal inflation target.

4.2.3 Optimal Trend Inflation

Figure 17 shows how the optimal inflation target changes with the degree of cyclical transfers at both the individual and aggregate levels. The case $\gamma_Y^T = 0$ corresponds to the benchmark THANK economy described in Section 3, which excludes cyclical transfers. The main takeaway is that more countercyclical transfers lower the optimal inflation target across all populations, indicating that the marginal benefits of raising trend inflation decline as transfers become more countercyclical.

Figure 17: Optimal Trend Inflation and Cyclical Targeted Transfer



Notes. The figure displays the optimal trend inflation rate as a function of the cyclical target transfers, given $\gamma_Y^H = 1.5$ and $\mu_y = 1.75$.

The decreasing pattern of HtMs' optimal inflation target aligns with our conjecture, outlined in the previous subsection. Importantly, the figure also shows a decline for Savers, implying that countercyclical transfers generate welfare gains from reduced output volatility and mean output losses that outweigh the welfare costs from increased volatility in their disposable income. Thus, the net welfare effect of countercyclical transfers is positive even for Savers, suggesting that such transfers reduce the need to raise the optimal inflation target. Naturally, as the individual-level optimal inflation targets fall with more countercyclical transfers, the aggregate optimal inflation target declines as well.

5 Conclusion

This paper studies how household income inequality, income risk, and countercyclical fiscal policy shape the optimal inflation target in an economy with an occasionally binding zero lower bound (ZLB). Using a tractable heterogeneous agent New Keynesian (THANK) model, we analytically characterize the cost-benefit trade-off of trend inflation across income groups.

The main findings are summarized as follows. Cyclical income inequality and income risk raise the marginal benefits of trend inflation for both Savers and HtMs by amplifying aggregate demand fluctuations, with the marginal benefits being larger for HtMs. In contrast, long-run income inequality dampens aggregate volatility and thus reduces marginal benefits, while simultaneously increasing the marginal cost of inflation for HtMs and reducing it for Savers due to differences in marginal utility of consumption. Quantitatively, under calibrated levels of inequality and income risk, the utilitarian optimal inflation rate rises modestly relative to the RANK benchmark, reflecting the large marginal benefits for HtMs. Introducing private debt further raises the op-

timal inflation target, as Fisherian debt deflation during ZLB episodes disproportionately harms indebted HtMs and deepens output contractions. Finally, countercyclical government spending and targeted transfers reduce the marginal benefits of trend inflation by stabilizing output and inflation and mitigating mean output losses, leading to a lower optimal inflation rate. In addition, for a given degree of government spending cyclical, higher government debt further increases the optimal inflation target, as the taxes required to finance higher debt servicing costs depress HtMs' consumption more at the ZLB.

As we note in the introduction, one limitation of this paper is that we keep the Taylor rule fixed and vary only the steady-state inflation rate. If Ramsey optimal monetary policy (i.e., monetary policy under commitment) were in place instead, we would expect different results. As discussed by [Jung, Teranishi and Watanabe \(2005\)](#), under a commitment policy the central bank keeps interest rates low for an extended period, which mitigates the severity of recessions during ZLB episodes. As a result, the marginal benefits of raising the inflation target are much lower than under a fixed Taylor rule. [Coibion, Gorodnichenko and Wieland \(2012\)](#) show in their extension that the optimal steady-state inflation rate in RANK is close to zero under Ramsey optimal policy. The same implication applies to our model with income inequality. Even in the presence of inequality, the optimal inflation target would be very small, so the increase in the optimal inflation target from introducing inequality would be much smaller than what we find under a suboptimal Taylor rule.

References

- Abo-Zaid, Salem**, "Optimal monetary policy and downward nominal wage rigidity in frictional labor markets," *Journal of Economic Dynamics and Control*, 2013, 37 (1), 345–364.
- Adam, Klaus and Junyi Zhu**, "Price-Level Changes and the Redistribution of Nominal Wealth across the Euro Area," *Journal of the European Economic Association*, 2016, 14 (4), 871–906.
- Auclert, Adrien**, "Monetary Policy and the Redistribution Channel," *American Economic Review*, 2019, 109 (6), 2333–2367.
- , Matthew Rognlie, and Ludwig Straub**, "Micro Jumps, Macro Humps: Monetary Policy and Business Cycles in an Estimated HANK Model," 2020. NBER Working Paper No. 26647.
- Bayer, Christian, Benjamin Born, and Ralph Luetticke**, "Shocks, Frictions, and Inequality in US Business Cycles," *American Economic Review*, May 2024, 114 (5), 1211–47.
- Benigno, Pierpaolo and Luca Antonio Ricci**, "The Inflation-Output Trade-Off with Downward Wage Rigidities," *American Economic Review*, June 2011, 101 (4), 1436–66.
- Bhandari, Anmol, David Evans, Mikhail Golosov, and Thomas J. Sargent**, "Inequality, Business Cycles, and Monetary-Fiscal Policy," *Econometrica*, 2021, 89 (6), 2559–2599.
- Bilbiie, Florin O.**, "Limited Asset Markets Participation, Monetary Policy and (Inverted) Aggregate Demand Logic," *Journal of Economic Theory*, 2008, 140 (1), 162–196.
- , "The New Keynesian cross,"** *Journal of Monetary Economics*, 2020, 114, 90–108.
- Bilbiie, Florin O**, "Monetary Policy and Heterogeneity: An Analytical Framework*," *The Review of Economic Studies*, 06 2024, p. rdae066.
- , Diego R Känzig, and Paolo Surico**, "Capital and income inequality: An aggregate-demand complementarity," *Journal of Monetary Economics*, 2022, 126, 154–169.
- , Tommaso Monacelli, and Roberto Perotti**, "Public debt and redistribution with borrowing constraints," *The Economic Journal*, 2013, 123 (566), F64–F98.
- Blanco, Andrés**, "Optimal Inflation Target in an Economy with Menu Costs and a Zero Lower Bound," *American Economic Journal: Macroeconomics*, July 2021, 13 (3), 108–41.
- , Corina Boar, Callum J Jones, and Virgiliu Midrigan**, "The Inflation Accelerator," Working Paper 32531, National Bureau of Economic Research May 2024.
- Carreras, Marc, Olivier Coibion, Yuriy Gorodnichenko, and Johannes Wieland**, "Infrequent but Long-Lived Zero Lower Bound Episodes and the Optimal Rate of Inflation," *Annual Review of*

Economics, 2016, 8 (1), 497–520.

Cho, Daeha, “Unemployment risk, MPC heterogeneity, and business cycles,” *Quantitative Economics*, May 2023, 14 (2), 717–751.

and Eunseong Ma, “Inflation Indexation and Zero Lower Bound,” 2024. Working Paper.

, Jung Hyun Kim, Kwang Hwan Kim, and Suk Joon Kim, “Optimal Trend Inflation in an Open Economy,” *Journal of Monetary Economics*, 2025, 149, 103675.

Coibion, Olivier, Yuriy Gorodnichenko, and Johannes Wieland, “The optimal inflation rate in New Keynesian models: should central banks raise their inflation targets in light of the zero lower bound?,” *Review of Economic Studies*, 2012, 79 (4), 1371–1406.

Davila, Eduardo and Andreas Schaab, “Optimal Monetary Policy with Heterogeneous Agents: A Timeless Ramsey Approach,” 2022. NBER Working Paper No. 30961.

Doepke, Matthias and Martin Schneider, “Inflation and the Redistribution of Nominal Wealth,” *Journal of Political Economy*, 2006, 114 (6), 1069–1097.

Dotsey, Michael, Robert G. King, and Alexander L. Wolman, “State-Dependent Pricing and the General Equilibrium Dynamics of Money and Output,” *The Quarterly Journal of Economics*, 05 1999, 114 (2), 655–690.

Fagan, Gabriel and Julian Messina, “Downward wage rigidity and optimal steady-state inflation,” Working Paper Series 1048, European Central Bank Apr 2009.

Fernández-Villaverde, Jesús, Joël Marbet, Galo Nuño, and Omar Rachedi, “Inequality and the zero lower bound,” *Journal of Econometrics*, 2025, 249, 105819.

Hagedorn, Marcus, Iourii Manovskii, and Kurt Mitman, “The Fiscal Multiplier,” 2019. Working Paper, University of Oslo.

Holden, Tom D., “Computation of Solutions to Dynamic Models with Occasionally Binding Constraints,” EconStor Preprints 130143, ZBW - Leibniz Information Centre for Economics April 2016.

Holden, Tom D., “Existence and uniqueness of solutions to dynamic models with occasionally binding constraints,” *Review of Economics and Statistics*, 2023, 105 (6), 1481–1499.

Jung, Taehun, Yuki Teranishi, and Tsutomu Watanabe, “Optimal Monetary Policy at the Zero-Interest-Rate Bound,” *Journal of Money, Credit, and Banking*, 2005, 37 (5), 813–835.

Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante, “Monetary Policy According to HANK,” *American Economic Review*, 2018, 108 (3), 697–743.

, **Giovanni L Violante, and Justin Weidner**, “The wealthy hand-to-mouth,” *Brookings papers on economic activity*, 2014, (1), 77–153.

Kim, Jinill and Francisco J. Ruge-Murcia, “How Much Inflation is Necessary to Grease the Wheels?,” *Journal of Monetary Economics*, 2009, 56 (3), 365–377.

Kulish, Mariano, James Morley, and Tim Robinson, “Estimating DSGE models with zero interest rate policy,” *Journal of Monetary Economics*, 2017, 88, 35–49.

Kwark, Noh-Sun and Eunseong Ma, “Entrepreneurship and Income Distribution Dynamics: Why Is the Income Share of Top Income Earners Acyclical over the Business Cycle?,” *International Economic Review*, 2021, 62 (1), 321–356.

Leeper, Eric M., Michael Plante, and Nora Traum, “Dynamics of fiscal financing in the United States,” *Journal of Econometrics*, 2010, 156 (2), 304–321.

L’Huillier, Jean-Paul and Raphael Schoenle, “Raising the Inflation Target: What Are the Effective Gains in Policy Room?,” 2023. Working Paper.

McKay, Alisdair and Christian K. Wolf, “Optimal Policy Rules in HANK,” 2023. Working Paper.

and Ricardo Reis, “The Role of Automatic Stabilizers in the U.S. Business Cycle,” *Econometrica*, 2016, 84 (1), 141–194.

, **Emi Nakamura, and Jón Steinsson**, “The Power of Forward Guidance Revisited,” *American Economic Review*, 2016, 106 (10), 3133–3158.

Mineyama, Tomohide, “Revisiting the Optimal Inflation Rate with Downward Nominal Wage Rigidity: The Role of Heterogeneity,” *Journal of Economic Dynamics and Control*, 2022, 139, 104350.

Nakamura, Emi and Jón Steinsson, “Five Facts about Prices: A Reevaluation of Menu Cost Models,” *Quarterly Journal of Economics*, 2008, 123 (4), 1415–1464.

Schaab, Andreas, “Micro and Macro Uncertainty,” Technical Report, Working Paper 2020.

Appendices

A Proof of Propositions

Proposition 1. Assume households are risk-averse such that $\sigma > 1$, then

1. $\Gamma_y^S \leq \Gamma_y^H$.
2. $\frac{\partial \Gamma_y^S}{\partial \gamma_Y^H} < 0$ and $\frac{\partial \Gamma_y^H}{\partial \gamma_Y^H} > 0$.
3. $\frac{\partial \Gamma_y^H}{\partial \mu_y} > 0$ and $\frac{\partial \Gamma_y^S}{\partial \mu_y} < (>)0$ if $\gamma_Y^H < (>)\frac{(\sigma-1)\lambda(1-\lambda)\mu_y}{\lambda(1-\lambda)\mu_y + \sigma\lambda^2}$.

Proof. For the first statement, subtracting the two coefficients yields:

$$\Gamma_y^S - \Gamma_y^H = Y^{1-\sigma} \left(\frac{(1-s_g)}{(1-\lambda)\mu_y + \lambda} \right)^{-\sigma} \frac{(1-\lambda)(\mu_y - \mu_y^\sigma) - (\lambda + \mu_y^\sigma(1-\lambda))\gamma_y^H}{\mu_y^\sigma(1-\lambda)((1-\lambda)\mu_y + \lambda)}$$

Since $\mu_y \leq \mu_y^\sigma$ for $\sigma \geq 1$, we obtain $\Gamma_y^S - \Gamma_y^H \leq 0$ for $\mu_y \geq 1$ and $\gamma_Y^H \geq 0$. The second statement follows directly by differentiating with respect to γ_Y^H :

$$\begin{aligned} \frac{\partial \Gamma_y^S}{\partial \gamma_Y^H} &= -Y^{1-\sigma} \left(\frac{(1-s_g)\mu_y}{(1-\lambda)\mu_y + \lambda} \right)^{-\sigma} \frac{\lambda\theta_H}{1-\lambda} < 0 \\ \frac{\partial \Gamma_y^H}{\partial \gamma_Y^H} &= Y^{1-\sigma} \left(\frac{1-s_g}{(1-\lambda)\mu_y + \lambda} \right)^{-\sigma} \theta_H > 0. \end{aligned}$$

Finally, the third statement is established by differentiating with respect to μ_y :

$$\frac{\partial \Gamma_y^H}{\partial \mu_y} = Y^{1-\sigma} (1-s_g)^{-\sigma} ((1-\lambda)\mu_y + \lambda)^{\sigma-2} (1-\lambda)(1+\gamma_Y^H)(\sigma-1) > 0,$$

$$\frac{\partial \Gamma_y^S}{\partial \mu_y} = Y^{1-\sigma} (1-s_g)^{-\sigma} ((1-\lambda)\mu_y + \lambda)^{\sigma-2} \mu_y^{-\sigma} \left(\lambda(1+\gamma_Y^H) - \sigma \cdot \frac{\lambda((1-\lambda)\mu_y - \lambda\gamma_Y^H)}{\mu_y(1-\lambda)} \right),$$

$$< (>) 0 \quad \text{if} \quad \gamma_Y^H < (>) \frac{(\sigma-1)\lambda(1-\lambda)\mu_y}{\lambda(1-\lambda)\mu_y + \sigma\lambda^2}, \quad \text{for } \sigma > 1.$$

□

Proposition 2. Assume households are risk-averse such that $\sigma > 1$ and $\lambda < \frac{1}{2}$. Then:

1. There exists a threshold level of risk aversion $\bar{\sigma}(\lambda, s_g)$ such that $\Gamma_{y^2}^H \leq \Gamma_{y^2}^S \leq 0$ if $\sigma \geq \bar{\sigma}(\lambda, s_g)$.
2. $\frac{\partial \Gamma_{y^2}^H}{\partial \gamma_Y^H} < 0$ and $\frac{\partial \Gamma_{y^2}^S}{\partial \mu_y} < 0$.

Proof. For the first statement, subtracting $\Gamma_{y^2}^H$ from $\Gamma_{y^2}^S$ yields:

$$\Gamma_{y^2}^S - \Gamma_{y^2}^H = \frac{1}{2} Y^{1-\sigma} \left(\frac{(1-s_g)}{(1-\lambda)\mu_y + \lambda} \right)^{-\sigma} \left(\frac{(1-\lambda)\mu_y + \lambda - \lambda(1+\gamma_Y^H)^2}{(1-\lambda)\mu_y^\sigma((1-\lambda)\mu_y + \lambda)} - \frac{\sigma((1-\lambda)\mu_y - \lambda\gamma_Y^H)^2}{(1-\lambda)^2(1-s_g)\mu_y^{\sigma+1}((1-\lambda)\mu_y + \lambda)} \right) - \left(1 - \frac{\sigma}{1-s_g} \right) \frac{(1+\gamma_Y^H)^2}{(1-\lambda)\mu_y + \lambda}.$$

Rewriting this expression with a common denominator, we obtain:

$$\begin{aligned} \text{sign} \left(\Gamma_{y^2}^S - \Gamma_{y^2}^H \right) &= \text{sign} \left(((\sigma - (1-s_g))(1-\lambda)^2\mu_y^{\sigma+1} - ((1-s_g)\lambda(1-\lambda)\mu_y + \sigma\lambda^2)) (\gamma_Y^H)^2 \right. \\ &\quad + 2(\sigma - (1-s_g)) (\lambda(1-\lambda)\mu_y + (1-\lambda)^2\mu_y^{\sigma+1}) \gamma_Y^H \\ &\quad \left. + (\sigma - (1-s_g))(1-\lambda)^2 (\mu_y^{\sigma+1} - \mu_y^2) \right) \end{aligned}$$

Since $\mu_Y \geq 1$ and $\sigma > 1$, it follows that:

$$\begin{aligned} &((\sigma - (1-s_g))(1-\lambda)^2\mu_y^{\sigma+1} - ((1-s_g)\lambda(1-\lambda)\mu_y + \sigma\lambda^2)) (\gamma_Y^H)^2 \\ &+ 2(\sigma - (1-s_g)) (\lambda(1-\lambda)\mu_y + (1-\lambda)^2\mu_y^{\sigma+1}) \gamma_Y^H \\ &+ (\sigma - (1-s_g))(1-\lambda)^2 (\mu_y^{\sigma+1} - \mu_y^2) \\ &\geq ((\sigma - (1-s_g))(1-\lambda)^2\mu_y^2 - ((1-s_g)\lambda(1-\lambda)\mu_y + \sigma\lambda^2)) (\gamma_Y^H)^2 \\ &+ 2(\sigma - (1-s_g)) (\lambda(1-\lambda)\mu_y + (1-\lambda)^2\mu_y^2) \gamma_Y^H. \end{aligned}$$

The right-hand side is always greater than or equal to 0 if:

$$\sigma \geq \frac{(1-s_g)(1-\lambda)((1-\lambda)\mu_y^2 + \lambda\mu_y)\gamma_Y^H + 2(1-s_g)((1-\lambda)^2\mu_y^2 + \lambda(1-\lambda)\mu_y)}{((1-\lambda)^2\mu_y^2 - \lambda^2)\gamma_Y^H + 2((1-\lambda)^2\mu_y^2 + \lambda(1-\lambda)\mu_y)} \equiv \tilde{\sigma}(\gamma_Y^H, \mu_y, \lambda, s_g).$$

Since $\tilde{\sigma}(\gamma_Y^H, \mu_y, \lambda, s_g)$ is increasing in γ_Y^H , we have:

$$\lim_{\gamma_Y^H \rightarrow \infty} \tilde{\sigma}(\gamma_Y^H, \mu_y, \lambda, s_g) = \frac{(1-s_g)(1-\lambda)((1-\lambda)\mu_y^2 + \lambda\mu_y)}{((1-\lambda)^2\mu_y^2 - \lambda^2)}.$$

This limit is decreasing in μ_y and reaches a maximum at $\mu_y = 1$. Hence, if $\sigma \geq \frac{(1-s_g)(1-\lambda)}{1-2\lambda} \equiv \bar{\sigma}(\lambda, s_g)$, $\Gamma_{y^2}^S \geq \Gamma_{y^2}^H$ holds.

The next step is to show that $\Gamma_{y^2}^S < 0$. Since $-\frac{1}{2}(1+\varphi)\chi(Yv^p)^{1+\varphi} < 0$, it suffices to show the first term is negative. The sign of the first term depends on the sign of

$\frac{1-\lambda\theta_H(1+\gamma_Y^H)^2}{1-\lambda} - \sigma \left(\frac{(1-s_g)\mu_y}{(1-\lambda)\mu_y+\lambda} \right)^{-1} \left(\frac{1-\lambda\theta_H(1+\gamma_Y^H)}{1-\lambda} \right)^2$. It follows that:

$$\begin{aligned} & \text{sign} \left(\frac{(1-\lambda)\mu_y + \lambda - \lambda(1+\gamma_Y^H)^2}{(1-\lambda)((1-\lambda)\mu_y + \lambda)} - \frac{\sigma((1-\lambda)\mu_y - \lambda\gamma_Y^H)^2}{(1-\lambda)^2(1-s_g)\mu_y((1-\lambda)\mu_y + \lambda)} \right) \\ &= \text{sign} \left((1-\lambda)(1-s_g)\mu_y \left((1-\lambda)\mu_y + \lambda - \lambda(1+\gamma_Y^H)^2 \right) - \sigma((1-\lambda)\mu_y - \lambda\gamma_Y^H)^2 \right). \end{aligned}$$

Rewriting the RHS of the above equation, we obtain:

$$-\lambda((1-\lambda)(1-s_g)\mu_y + \sigma\lambda)(\gamma_Y^H)^2 + 2\lambda(1-\lambda)(\sigma + s_g - 1)\mu_y\gamma_Y^H - (1-\lambda)^2(\sigma + s_g - 1)\mu_y^2$$

This quadratic function of γ_Y^H is negative if the discriminant is negative. The discriminant is:

$$\begin{aligned} D &= 4\lambda^2(1-\lambda)^2\mu_y^2(\sigma + s_g - 1)^2 - 4\lambda(1-\lambda)^3(1-s_g)(\sigma + s_g - 1)\mu_y^3 - 4\sigma\lambda^2(1-\lambda)^2(\sigma + s_g - 1)\mu_y^2 \\ &= 4\lambda(1-\lambda)^2\mu_y^2(\sigma + s_g - 1)(-\lambda(1-s_g) - (1-\lambda)(1-s_g)\mu_y). \end{aligned}$$

Since the discriminant is negative when $s_g > 0$, the first term of $\Gamma_{y^2}^S$ is always negative.

For the second statement, differentiating $\Gamma_{y^2}^H$ with respect to γ_Y^H yields:

$$\frac{\partial \Gamma_{y^2}^H}{\partial \gamma_Y^H} = Y^{1-\sigma} \left(\frac{(1-s_g)}{(1-\lambda)\mu_y + \lambda} \right)^{-\sigma} \theta_H(1+\gamma_Y^H) \left(1 - \frac{\sigma}{1-s_g} \right).$$

Since $\frac{\sigma}{1-s_g} > 1$, this partial derivative is negative. Turning to the next derivative, we obtain:

$$\frac{\partial \Gamma_{y^2}^H}{\partial \mu_y} = \frac{1}{2} Y^{1-\sigma} \frac{(1+\gamma_Y^H)^2}{(1-s_g)^\sigma} \left(1 - \frac{\sigma}{1-s_g} \right) (\sigma - 1) ((1-\lambda)\mu_y + \lambda)^{\sigma-2} (1-\lambda).$$

For $\sigma > 1$, this partial derivative is negative. □

Proposition 3. Assume that long-run inequality exists (i.e., $\mu_y > 1$) and that agents are risk-averse such that $\sigma > 1$. Then, $\left| \frac{\partial \Gamma^S}{\partial \Pi} \right| < \left| \frac{\partial \Gamma^H}{\partial \Pi} \right|$.

Proof. Rewriting the constant term using the production schedule, we obtain:

$$\Gamma^j = \frac{(C^j)^{1-\sigma} - 1}{1-\sigma} - \chi \frac{(Yv^p)^{1+\phi}}{1+\phi} - \chi (Yv^p)^{1+\phi} \frac{1}{2} (1+\phi) \frac{\sigma_a^2}{1-\rho_a^2}, \quad j \in \{S, H\}.$$

Since the second and third terms are common to both agents, any heterogeneity in the welfare effect of trend inflation must originate from the first term. Differentiating the first term with respect to trend inflation yields:

$$\frac{\partial \left(\frac{(C^j)^{1-\sigma} - 1}{1-\sigma} \right)}{\partial \Pi} = \begin{cases} Y^{-\sigma} \left(\frac{\mu_y(1-s_g)}{(1-\lambda)\mu_y + \lambda} \right)^{1-\sigma} \frac{\partial Y}{\partial \Pi} & \text{if } j = S \\ Y^{-\sigma} \left(\frac{(1-s_g)}{(1-\lambda)\mu_y + \lambda} \right)^{1-\sigma} \frac{\partial Y}{\partial \Pi} & \text{if } j = H \end{cases}$$

Since $\mu_y > 1$ and $\sigma > 1$, it follows that $\left| \frac{\partial \left(\frac{(C^S)^{1-\sigma} - 1}{1-\sigma} \right)}{\partial \Pi} \right| < \left| \frac{\partial \left(\frac{(C^H)^{1-\sigma} - 1}{1-\sigma} \right)}{\partial \Pi} \right|$ in the region where $\frac{\partial Y}{\partial \Pi} < 0$.

Hence, we conclude that $\left| \frac{\partial \Gamma^S}{\partial \Pi} \right| < \left| \frac{\partial \Gamma^H}{\partial \Pi} \right|$. \square

Proposition 4. Assume no cyclical transfer ($\gamma_Y^T = 0$) and $\alpha \in [0, \lambda]$. Then:

1. $\frac{\partial \Gamma_y^S}{\partial \gamma_Y^G} < 0$, and $\frac{\partial \Gamma_y^H}{\partial \gamma_Y^G} \leq 0$, with equality holding at $\alpha = 0$.
2. If $(1 - \lambda)\mu_y - \lambda\gamma_Y^H > 0$, then $\frac{\partial \Gamma_{y^2}^S}{\partial \gamma_Y^G} > 0$ and $\frac{\partial \Gamma_{y^2}^H}{\partial \gamma_Y^G} \geq 0$, with equality holding at $\alpha = 0$.

Proof. The first statement follows directly from partial differentiation of Γ_y^S and Γ_y^H with respect to γ_Y^G :

$$\begin{aligned} \frac{\partial \Gamma_y^S}{\partial \gamma_Y^G} &= -Y^{1-\sigma} \left(\frac{(1-s_g)\mu_y}{(1-\lambda)\mu_y + \lambda} \right)^{-\sigma} \frac{1-\alpha}{1-\lambda} s_g < 0 \\ \frac{\partial \Gamma_y^H}{\partial \gamma_Y^G} &= -Y^{1-\sigma} \left(\frac{1-s_g}{(1-\lambda)\mu_y + \lambda} \right)^{-\sigma} \frac{\alpha}{\lambda} s_g \leq 0. \end{aligned}$$

To prove the second statement, differentiate $\Gamma_{y^2}^H$ with respect to γ_Y^G :

$$\frac{\partial \Gamma_{y^2}^H}{\partial \gamma_Y^G} = Y^{1-\sigma} \left(\frac{(1-s_g)}{(1-\lambda)\mu_y + \lambda} \right)^{-\sigma} \frac{\alpha}{\lambda} s_g \left(\sigma \frac{1 + \gamma_Y^H}{1-s_g} - \sigma \frac{((1-\lambda)\mu_y + \lambda)}{1-s_g} \frac{\alpha}{\lambda} s_g \gamma_Y^G - \gamma_Y^G \right).$$

Since $\gamma_Y^G \leq 0$, it follows that:

$$\begin{aligned} \frac{\partial \Gamma_{y^2}^H}{\partial \gamma_Y^G} &= Y^{1-\sigma} \left(\frac{(1-s_g)}{(1-\lambda)\mu_y + \lambda} \right)^{-\sigma} \frac{\alpha}{\lambda} s_g \left(\sigma \frac{1 + \gamma_Y^H}{1-s_g} - \sigma \frac{((1-\lambda)\mu_y + \lambda)}{1-s_g} \frac{\alpha}{\lambda} s_g \gamma_Y^G - \gamma_Y^G \right) \\ &\geq Y^{1-\sigma} \left(\frac{(1-s_g)}{(1-\lambda)\mu_y + \lambda} \right)^{-\sigma} \frac{\alpha}{\lambda} s_g \left(\sigma \frac{1 + \gamma_Y^H}{1-s_g} \right) \geq 0. \end{aligned}$$

Similarly, differentiating $\Gamma_{y^2}^S$ with respect to γ_Y^G and using $\gamma_Y^G \leq 0$, we get:

$$\begin{aligned} \frac{\partial \Gamma_{y^2}^S}{\partial \gamma_Y^G} &= Y^{1-\sigma} \left(\frac{(1-s_g)\mu_y}{(1-\lambda)\mu_y + \lambda} \right)^{-\sigma} \left(\frac{1-\alpha}{1-\lambda} s_g \right) \left(\sigma \frac{(1-\lambda)\mu_y + \lambda}{1-s_g} \left(\frac{(1-\lambda)\mu_y - \lambda\gamma_Y^H}{(1-\lambda)((1-\lambda)\mu_y + \lambda)} - \frac{1-\alpha}{1-\lambda} s_g \gamma_Y^G \right) - \gamma_Y^G \right) \\ &\geq Y^{1-\sigma} \left(\frac{(1-s_g)\mu_y}{(1-\lambda)\mu_y + \lambda} \right)^{-\sigma} \left(\frac{1-\alpha}{1-\lambda} s_g \right) \left(\sigma \frac{(1-\lambda)\mu_y + \lambda}{1-s_g} \left(\frac{(1-\lambda)\mu_y - \lambda\gamma_Y^H}{(1-\lambda)((1-\lambda)\mu_y + \lambda)} \right) \right). \end{aligned}$$

If $(1 - \lambda)\mu_y - \lambda\gamma_Y^H > 0$, the right-hand side of the inequality is positive. \square

Proposition 5. Assume no cyclical government spending ($\gamma_Y^G = 0$), then

1. $\frac{\partial \Gamma_y^S}{\partial \gamma_Y^T} < 0$ and $\frac{\partial \Gamma_y^H}{\partial \gamma_Y^T} > 0$.

2. If $\frac{1-s_g-2\sigma(1+\gamma_Y^H)}{2\sigma((1-\lambda)\mu_y+\lambda)} < \gamma_Y^T < \frac{(\sigma-\frac{1}{2}(1-s_g))(1-\lambda)\mu_y-\lambda\gamma_Y^H}{\sigma\lambda((1-\lambda)\mu_y+\lambda)}$, $\frac{\partial\Gamma_{y^2}^S}{\partial\gamma_Y^T} > 0$ and $\frac{\partial\Gamma_{y^2}^H}{\partial\gamma_Y^T} < 0$.

Proof. The first statement follows directly from the partial derivatives of Γ_y^S and Γ_y^H with respect to γ_Y^T :

$$\begin{aligned}\frac{\partial\Gamma_y^S}{\partial\gamma_Y^T} &= -Y^{1-\sigma} \left(\frac{(1-s_g)\mu_y}{(1-\lambda)\mu_y+\lambda} \right)^{-\sigma} \frac{\lambda}{1-\lambda} < 0 \\ \frac{\partial\Gamma_y^H}{\partial\gamma_Y^T} &= Y^{1-\sigma} \left(\frac{(1-s_g)}{(1-\lambda)\mu_y+\lambda} \right)^{-\sigma} > 0.\end{aligned}$$

For the second statement, the partial derivatives of $\Gamma_{y^2}^S$ and $\Gamma_{y^2}^H$ with respect to γ_Y^T are given by:

$$\begin{aligned}\frac{\partial\Gamma_{y^2}^S}{\partial\gamma_Y^T} &= -Y^{1-\sigma} \left(\frac{(1-s_g)\mu_y}{(1-\lambda)\mu_y+\lambda} \right)^{-\sigma} \frac{\lambda}{1-\lambda} \left(\frac{\sigma((1-\lambda)\mu_y-\lambda\gamma_Y^H)}{(1-s_g)(1-\lambda)\mu_y} - \frac{\sigma\lambda\gamma_Y^T((1-\lambda)\mu_y+\lambda)}{(1-s_g)(1-\lambda)\mu_y} - \frac{1}{2} \right) \\ \frac{\partial\Gamma_{y^2}^H}{\partial\gamma_Y^T} &= -Y^{1-\sigma} \left(\frac{(1-s_g)}{(1-\lambda)\mu_y+\lambda} \right)^{-\sigma} \left(\frac{1}{2} - \frac{\sigma(1+\gamma_Y^H)}{1-s_g} - \frac{\sigma((1-\lambda)\mu_y+\lambda)}{1-s_g} \gamma_Y^T \right).\end{aligned}$$

Since $s_g < 1$ and $\lambda > 0$, the signs of the derivatives depend on the expressions in the parentheses:

$$\begin{aligned}\text{sign} \left(\frac{\partial\Gamma_{y^2}^S}{\partial\gamma_Y^T} \right) &= \text{sign} \left(\frac{\sigma((1-\lambda)\mu_y-\lambda\gamma_Y^H)}{(1-s_g)(1-\lambda)\mu_y} - \frac{\sigma\lambda\gamma_Y^T((1-\lambda)\mu_y+\lambda)}{(1-s_g)(1-\lambda)\mu_y} - \frac{1}{2} \right) \\ \text{sign} \left(\frac{\partial\Gamma_{y^2}^H}{\partial\gamma_Y^T} \right) &= \text{sign} \left(\frac{1}{2} - \frac{\sigma(1+\gamma_Y^H)}{1-s_g} - \frac{\sigma((1-\lambda)\mu_y+\lambda)}{1-s_g} \gamma_Y^T \right).\end{aligned}$$

Under the stated range for γ_Y^T , $\text{sign} \left(\frac{\partial\Gamma_{y^2}^S}{\partial\gamma_Y^T} \right) > 0$ and $\text{sign} \left(\frac{\partial\Gamma_{y^2}^H}{\partial\gamma_Y^T} \right) < 0$, completing the proof. \square

B Quadratic Welfare Function under Wealth Inequality

$$\begin{aligned}
EU_t &= \mathbf{E} \left[(1-\lambda) \frac{(C_t^S)^{1-\sigma} - 1}{1-\sigma} + \lambda \frac{(C_t^H)^{1-\sigma} - 1}{1-\sigma} - \chi \frac{N_t^{1+\varphi}}{1+\varphi} \right] \\
&\approx (1-\lambda) \frac{(C^S)^{1-\sigma} - 1}{1-\sigma} + \lambda \frac{(C^H)^{1-\sigma} - 1}{1-\sigma} - \chi \frac{N^{1+\varphi}}{1+\varphi} \\
&+ (C^S)^{-\sigma} \lambda C^H \left(\left(\frac{C^H}{C^S} \right)^{-\sigma} - 1 \right) \mathbf{E}(\hat{c}_t^H) \\
&+ (C^S)^{-\sigma} \left(Y - G \gamma_Y^G - \frac{MCv^p}{1-\nu} \frac{C}{1-s_g} - \left(\left(\frac{C}{C^S} \right)^{-\sigma} - 1 \right) \frac{MCv^p}{1-\nu} \frac{C}{1-s_g} \right) \mathbf{E}(\hat{y}_t) \\
&- C^{-\sigma} \frac{MCv^p}{1-\nu} \frac{C}{1-s_g} \frac{\varepsilon_p \bar{\theta}_p (\Pi - 1)}{(1-\bar{\theta}_p)(1-\bar{\theta}_p \Pi)} \mathbf{E}(\hat{\pi}_t) \\
&- (C^S)^{-\sigma} (1-\lambda) C^S \frac{\sigma}{2} \mathbf{E}((\hat{c}_t^S)^2) \\
&+ (C^S)^{-\sigma} \left(\lambda \left(\left(\frac{C^H}{C^S} \right)^{-\sigma} - 1 \right) \frac{C^H}{2} - \lambda C^H \left(\frac{C^H}{C^S} \right)^{-\sigma} \frac{\sigma}{2} \right) \mathbf{E}((\hat{c}_t^H)^2) \\
&+ (C^S)^{-\sigma} \left(\frac{Y}{2} - \frac{G}{2} (\gamma_Y^G)^2 - \frac{1}{2} \frac{MCv^p}{1-\nu} \frac{C}{1-s_g} - \frac{1}{2} \frac{MCv^p}{1-\nu} \frac{C}{1-s_g} - \left(\frac{C}{C^S} \right)^{-\sigma} \frac{\varphi}{2} \frac{MCv^p}{1-\nu} \frac{C}{1-s_g} \right) \mathbf{E}(\hat{y}_t^2) \\
&+ (C^S)^{-\sigma} \left(\frac{MCv^p}{1-\nu} \frac{C}{1-s_g} + \left(\left(\frac{C}{C^S} \right)^{-\sigma} - 1 \right) \frac{MCv^p}{1-\nu} \frac{C}{1-s_g} + \left(\frac{C}{C^S} \right)^{-\sigma} \frac{MCv^p}{1-\nu} \frac{C}{1-s_g} \varphi \right) COV(\hat{y}_t, \hat{A}_t) \\
&+ (C^S)^{-\sigma} \left(-\frac{1}{2} \frac{MCv^p}{1-\nu} \frac{C}{1-s_g} - \frac{1}{2} \left(\left(\frac{C}{C^S} \right)^{-\sigma} - 1 \right) \frac{MCv^p}{1-\nu} \frac{C}{1-s_g} - \frac{\varphi}{2} \left(\frac{C}{C^S} \right)^{-\sigma} \frac{MCv^p}{1-\nu} \frac{C}{1-s_g} \right) \frac{\sigma_a^2}{1-\rho_a^2} \\
&- \frac{1}{2} C^{-\sigma} \frac{MCv^p}{1-\nu} \frac{C}{1-s_g} \frac{\varepsilon_p \bar{\theta}_p (\varepsilon_p (\Pi - 1) + 1)}{(1-\bar{\theta}_p)^2} \mathbf{E}(\hat{\pi}_t^2)
\end{aligned}$$

C Optimal Monetary and Fiscal Policy Mix

C.1 Zero Government Debt Case

In Sections 4.1 and 4.2, we showed that introducing countercyclical fiscal policy reduces the optimal inflation target at both the individual and aggregate levels. In that analysis, the degree of fiscal policy was treated as exogenous, and we searched for the inflation target that maximizes individual or utilitarian welfare. However, this exercise does not reveal whether the combination of countercyclical fiscal policy and a relatively low inflation target constitutes an optimal policy choice. In this subsection, we investigate whether such a policy combination leads to better welfare outcomes than adjusting the inflation target alone. If so, countercyclical fiscal policy can serve as a substitute for a high inflation target from the policymakers' perspective. To address this, we endogenize both the fiscal and monetary policy instruments and compute the jointly optimal mix

Table 4: Optimal Monetary and Fiscal Policy Mix

(a) Cyclical Government Spending			(b) Cyclical Targeted Transfer	
(Π, γ_Y^G)	$\alpha = \lambda$	$\alpha = 0$	(Π, γ_Y^T)	
Saver	(1.80, 0)	(1.80, 0)	Saver	(1.74, -0.25)
HtM	(2.11, -1)	(2.06, -1)	HtM	(2.01, -0.25)
Aggregate	(1.95, 0)	(1.87, -1)	Aggregate	(1.85, -0.25)

of the inflation target (Π) and fiscal policy (γ_Y^G or γ_Y^T).

We consider three scenarios in calculating the optimal policy mix:

1. Trend inflation and cyclical government spending, with tax financing equally shared between Savers and HtMs ($\alpha = \lambda$);
2. Trend inflation and cyclical government spending, financed exclusively by taxation on Savers ($\alpha = 0$);
3. Trend inflation and cyclical targeted transfers.

For practical purposes, we restrict the range of cyclical government spending (γ_Y^G) to $[-1, 0]$ and that of cyclical targeted transfers (γ_Y^T) to $[-0.25, 0]$.¹⁹

Table 4 reports the optimal policy mix at both the individual and aggregate levels in the THANK economy with income risk. The left panel presents the optimal combination of trend inflation and cyclical government spending, while the right panel presents the optimal combination of trend inflation and cyclical transfers. The key results from the table can be summarized as follows: (1) Savers prefer not to implement cyclical government spending, whereas HtMs prefer maximal countercyclical government spending coupled with a lower inflation target. (2) Only under progressive taxation, cyclical government spending can substitute for trend inflation at the aggregate level. (3) Cyclical targeted transfers serve as a substitute for trend inflation at both the household and aggregate levels.

The first result can be inferred from the first and second rows of the left table, which reveal a divergence in the optimal policy mix between Savers and HtM households. These rows show that, regardless of tax progressivity, it is optimal for Savers to adjust only the inflation target, whereas for HtMs, it is optimal to use a mix of both government spending and the inflation target. As discussed in Section 3, the marginal cost of increasing trend inflation is lower for Savers than for HtMs, and the income loss Savers experience during ZLB episodes is relatively modest compared to that of HtMs. Consequently, Savers have little incentive to favor government spending, which

¹⁹In [Leeper, Plante and Traum \(2010\)](#), the output elasticity of government spending is -0.034, while that of transfers is -0.13. As such, extremely negative values of these parameters may be viewed as implausible calibrations.

increases the volatility of their disposable income through taxation. Thus, for Savers, trend inflation provides cheaper insurance against ZLB episodes than government spending.

For HtMs, the marginal cost of increasing trend inflation is higher than it is for Savers, and their income losses during ZLB episodes are also significantly greater. As a result, they are more willing to accept cyclical government spending, which helps mitigate these income losses. Since countercyclical government spending reduces the severity of income losses for HtMs, the optimal inflation rate under a policy mix is lower than when cyclical government spending is not considered. Notably, the optimal inflation rate for HtMs is lower when they are exempt from the tax burden ($\alpha = 0$) than when the burden is shared equally ($\alpha = \lambda$). This is because countercyclical government spending provides stronger stabilization of disposable income for HtMs under maximally progressive taxation, reducing the need for a higher trend inflation rate.

The second result can be seen in the last row of the left table, which shows that whether countercyclical government spending substitutes for trend inflation at the aggregate level depends critically on the method of financing. Since only HtMs support the use of cyclical government spending alongside a lower inflation target, the aggregate optimal inflation rate will decline only if their welfare gains from such spending are sufficiently large. This occurs when taxation is maximally progressive. Under minimal progressivity, however, the tax burden on HtMs reduces their welfare gains from cyclical government spending, resulting in an aggregate optimal trend inflation rate that is effectively the same as it would be in the absence of such spending.

The third result is confirmed by the right table, which shows that the entire population favors the use of cyclical transfers alongside a relatively low inflation target. Since trend inflation imposes a steady-state welfare cost through increased price dispersion, HtMs prefer cyclical transfers, which reduce their income losses at no direct cost. Interestingly, even for Savers, a combination of cyclical transfers and a lower inflation target proves more beneficial than relying solely on higher trend inflation, despite their relatively low marginal cost of inflation. This suggests that the welfare gains from reduced mean output losses and aggregate volatility generated by countercyclical transfers are substantial enough to outweigh the associated tax burden. Given that both groups prefer the policy mix over relying on the inflation target alone, the aggregate optimal inflation rate under cyclical transfers decreases by 10 basis points compared to the economy without such transfers.

In sum, the optimal mix analysis reveals that whether countercyclical fiscal policy can partially substitute for a high inflation target depends on the degree of tax progressivity. When either cyclical government spending financed through progressive taxation or cyclical targeted transfers are introduced as policy tools alongside the inflation target, the optimal inflation rate in THANK

Table 5: Optimal Monetary and Fiscal Policy Mix under Positive Government Debt

(a) Maximal Tax Progressivity ($\alpha = 0$)			(b) Minimal Tax Progressivity ($\alpha = \lambda$)		
(Π, γ_Y^G)	$\gamma_G^b = 0$	$\gamma_G^b = 0.5$	(Π, γ_Y^G)	$\gamma_G^b = 0$	$\gamma_G^b = 0.5$
$b/4Y = 0$	(1.85, -1)	(1.85, -1)	$b/4Y = 0$	(1.95, 0)	(1.95, 0)
$b/4Y = 1.25$	(1.86, -1)	(1.86, -1)	$b/4Y = 1.25$	(2.47, -1)	(2.46, -1)

decreases relative to the case without such fiscal policies and gets closer to that under RANK.²⁰ This finding implies that when countercyclical fiscal policies are well-designed, inequality may not play a crucial role in determining the optimal inflation target for central banks.

C.2 Positive Government Debt Case

We investigate whether introducing public debt changes the jointly optimal mix of the inflation target (Π) and the countercyclicality of government consumption (γ_Y^G) under maximal and minimal tax progressivity. Table 5(a) reports the results under maximal tax progressivity ($\alpha = 0$) for different levels of steady-state debt and for different degrees of cyclicity in debt issuance. The optimal policy mix is the same across these dimensions. As observed in Figure 14(a), since the optimal inflation rates are largely insensitive to steady-state debt and to the cyclicity of debt issuance, the aggregate dynamics at the ZLB and its associated costs are also insensitive. As a result, for all combinations of steady-state debt and debt cyclicity, the optimal policy mix involves using cyclical government spending together with a lower inflation target. This mix mainly improves the welfare of HtMs, whose marginal cost of raising trend inflation is high and whose income losses during ZLB episodes are substantial.

Table 5(b) reports the optimal policy mix under minimal tax progressivity ($\alpha = \lambda$) across different levels of steady-state debt and different degrees of cyclicity in debt issuance. For a given level of steady-state debt, the optimal policy mix is the same across different degrees of debt cyclicity, reflecting the insensitivity of the optimal inflation rate to this dimension, as shown in Figure 14(b). However, in contrast to the economy without steady-state debt, the economy with positive steady-state debt favors both a much higher optimal inflation rate and the use of countercyclical fiscal policy.²¹ Such a higher optimal inflation rate reflects the result shown in Figure 14(b), which shows that more positive government debt leads to a large increase in the optimal inflation rate. Regarding fiscal policy, the use of countercyclical policy under positive government debt mitigates the welfare loss that HtMs face during ZLB episodes from the tax burden used to finance higher interest payments on public debt by increasing labor income.

²⁰In RANK, transfer policy does not affect aggregate moments or the coefficients in the approximate welfare function. Moreover, the optimal policy mix (Π, γ_Y^G) in RANK is (1.49, 0), indicating that the availability of countercyclical government spending does not reduce the optimal inflation target in RANK.

²¹Of course, the optimal inflation rate under a policy mix with positive government debt is slightly lower than when countercyclical fiscal policy is not used (2.46% vs. 2.55%).