

Optimal Monetary Policy According to HANK

by Sushant Acharya,¹ Edouard Challe² and Keshav Dogra³

¹Canadian Economic Analysis Department
Bank of Canada, Ottawa, Ontario, Canada K1A 0G9

²European University Institute and Center for Research in Economics and
Statistics (CREST)

³Federal Reserve Bank of New York

sacharya@bankofcanada.ca, edouard.challe@gmail.com,
keshav.dogra@ny.frb.org



Bank of Canada staff working papers provide a forum for staff to publish work-in-progress research independently from the Bank's Governing Council. This research may support or challenge prevailing policy orthodoxy. Therefore, the views expressed in this paper are solely those of the authors and may differ from official Bank of Canada views. No responsibility for them should be attributed to the Bank.

Acknowledgements

We are grateful to Adrien Auclert, Anmol Bhandari, Florin Bilbiie, Christopher Carroll, Russell Cooper, Clodomiro Ferreira, Greg Kaplan, Antoine Lepetit, Galo Nuño, and Gianluca Violante for helpful discussions. We also received useful comments from seminar participants at HEC Paris, UT Austin, UC3M, EUI, Université Paris-Dauphine, Université Paris 8, Banque de France, Drexel University, University of Essex and CREST as well as from conference participants at the Barcelona GSE Summer Forum (Monetary Policy and Central Banking), the NBER SI (Micro Data and Macro Models), the Salento Macro Meetings, ASSA-AEA, SED, T2M and T3M-VR. Edouard Challe acknowledges financial support from the French National Research Agency (Labex Ecodec/ANR-11-LABX-0047 and ANR-20-CE26-0018-01). The views expressed in this paper are those of the authors and do not necessarily represent those of the Bank of Canada, the Federal Reserve Bank of New York or the Federal Reserve System.

Abstract

We study optimal monetary policy in an analytically tractable Heterogeneous Agent New Keynesian model with rich cross-sectional heterogeneity. Optimal policy differs from that in a representative agent model because monetary policy can affect consumption inequality by reducing both idiosyncratic consumption risk and the inequality that arises from households' unequal exposures to aggregate shocks. Simple target criteria summarize the planner's trade-off between consumption inequality, productive efficiency and price stability. Mitigating consumption inequality requires putting some weight on stabilizing the level of output and correspondingly reducing the weights on the output gap and the price level relative to an economy without inequality.

Topics: Economic models, Monetary policy

JEL codes: E21, E30, E52, E62, E63

We study optimal monetary policy in an analytically tractable Heterogeneous Agent New Keynesian (HANK) model with rich cross-sectional dispersion in income, wealth and consumption. While the HANK literature has shown that household heterogeneity can change the positive effects of monetary policy on the economy (e.g., Kaplan et al. (2018); Auclert et al. (2018); Auclert (2019); Bilbiie (2019); Ravn and Sterk (2020)), the normative implications of HANK, and the reciprocal effects of monetary policy on inequality, have been less well studied. This is because characterizing optimal monetary policy in HANK models with substantial heterogeneity is technically difficult. While the response to this challenge has been mainly computational so far (Bhandari et al., 2021; Le Grand et al., 2021; Nuño and Thomas, 2021), we instead take an analytical route. We study a standard New Keynesian economy in which households face idiosyncratic income risk, with two key assumptions: (i) households have Constant Absolute Risk Aversion (CARA) utility; and (ii) the idiosyncratic shocks they face are normally distributed. As in Acharya and Dogra (2020), these assumptions facilitate linear aggregation and imply that the *positive* behavior of macroeconomic aggregates can be described independently of distributional considerations. But of course, from a *normative* perspective, consumption dispersion affects welfare and, hence, optimal monetary policy. Crucially, in our framework the welfare cost of consumption dispersion is summarized by a scalar sufficient statistic that evolves recursively. This makes the planner’s problem tractable, allowing us to explicitly solve for optimal monetary policy in HANK and to dissect how and why it differs from that in the Representative Agent New Keynesian (RANK) model.

Optimal monetary policy can differ in HANK and RANK because idiosyncratic consumption risk and inequality (trivially absent in RANK) reduce social welfare in HANK. Thus, while the RANK planner seeks to stabilize prices and keep output at its productively efficient level, the HANK planner has an additional objective—to reduce consumption inequality. Our analytical framework distinguishes between two broad ways in which monetary policy can affect consumption inequality. First, monetary policy may reduce *idiosyncratic consumption risk* faced by households. Second, it may reduce the inequality arising from households’ *unequal exposure* to aggregate shocks and policy. To understand how each of these forces affects optimal policy, we first abstract from unequal exposure altogether to focus exclusively on idiosyncratic consumption risk. We do so by studying a baseline economy in which households are ex ante identical—a utilitarian planner optimally sets wealth taxes to eliminate pre-existing wealth inequality, and dividends are equally distributed across households, eliminating ex ante differences in income. In this economy, monetary policy may reduce idiosyncratic consumption risk via two specific channels. First, it can reduce the level of idiosyncratic income risk that households face (the *income-risk channel*). How to achieve this naturally depends on the cyclicity of income risk: if income risk is countercyclical, monetary policy would need to raise output in order to lower risk, while the opposite is true if risk is procyclical. Second, monetary policy can facilitate households’ self-insurance and thereby reduce the pass-through from individual income shocks to consumption (the *self-insurance channel*). This is because low rates not only directly facilitate self-insurance through the bond market (by making it easier to borrow to insulate consumption from income shocks) they also indirectly facilitate self-insurance through the labor market due to their expansionary impact on current and future wages (against which households can borrow).

Our analysis yields both methodological and substantive results. Methodologically, we characterize optimal monetary policy in terms of a simple *target criterion* that summarizes the trade-offs facing the planner. Our analysis thus extends (and nests as a special case) the description of optimal monetary policy

in RANK (Galí, 2015; Woodford, 2003). In RANK, the target criterion features some weight on stabilizing the output gap and some on stabilizing the price level because the relevant trade-off in RANK is between departures from productive efficiency and price stability. In HANK, however, the planner also seeks to reduce consumption inequality through the channels just described. Our main substantive result is that, in the empirically relevant case of countercyclical income risk, this concern for inequality leads monetary policy to put some weight on stabilizing the *level* of output and to correspondingly reduce the weights on the output gap and the price level relative to RANK. In our calibrated model, the HANK target criterion features roughly equal weight on the level of output and the output gap and features a 50% smaller weight on the price level than in RANK.

The reason concern for inequality leads the HANK planner to put some weight on stabilizing output when income risk is countercyclical is that in this case, both the income risk and self-insurance channels work in the same direction and expansionary monetary policy unambiguously reduces consumption risk. Thus, in response to aggregate shocks which would warrant a contraction in output in RANK (e.g., a fall in productivity or an increase in desired markups), the HANK planner raises interest rates less aggressively than in RANK, curtailing the fall in output. While this comes at the cost of productive inefficiency and higher inflation, cushioning the fall in output is optimal since it mitigates the rise in consumption inequality. Thus, even when households are ex ante identical and are equally exposed to aggregate shocks, idiosyncratic risk can lead to substantial differences in optimal monetary policy.

Our methodological and substantive results carry through to the case where monetary policy affects inequality via unequal exposures in addition to idiosyncratic risk. We study two different sources of unequal exposures. First, we allow for ex ante wealth heterogeneity. This is done by departing from our baseline assumption of a utilitarian planner, assuming instead that the planner is non-utilitarian and consequently sets wealth taxes in a way that does not completely eliminate ex ante wealth dispersion. In the presence of such wealth inequality, a surprise interest rate hike redistributes from poor debtors to rich savers (the unhedged interest rate exposure (URE) channel described in Auclert (2019)), an effect that must be internalized by the planner when setting monetary policy. These effects are absent in our baseline model since the utilitarian planner uses fiscal policy to eliminate pre-existing wealth inequality. This provides an additional reason to avoid large surprise rate hikes. Consequently, following a decline in productivity or an increase in desired markups, the non-utilitarian planner implements an even smaller fall in output on impact relative to RANK; optimal policy can still be characterized by a target criterion, which now puts larger weight on the level of output (relative to either the output gap or the price level) at the time the aggregate shock hits.

The second source of unequal exposure we study is unequally distributed dividends: we assume that while all households supply labor, only a fraction receive dividends. This provides another reason to avoid large fluctuations in output. To the extent that wages and profits react differently to movements in output, such fluctuations redistribute between stockholders and workers. Thus, the HANK planner has another reason to avoid large falls in output in response to adverse shocks, since these would increase consumption inequality between non-stockholders and stockholders. The planner's desire to avoid such between-group inequality is captured by the presence of the present discounted value of dividends in the target criterion (in addition to output, the output gap and the price level). Overall, while cyclical idiosyncratic risk and unequal exposures introduce two distinct ways in which monetary policy can affect inequality, both imply

that optimal policy should put more weight on stabilizing output relative to achieving productive efficiency or price stability. (Again, in RANK, both sources of inequality are absent and the planner puts *no* weight on output stabilization, focusing exclusively on stabilizing the output gap and the price level.)

Finally, we illustrate the versatility of our framework by extending it in two more dimensions. First, we study how the presence of *hand-to-mouth* (HtM) households, who have high marginal propensities to consume, affects our results. The presence of HtM households does not qualitatively change our results but quantitatively magnifies them. This is because HtM households cannot self-insure using the bond market, making consumption inequality within this group higher and more sensitive to monetary policy than that within the group of unconstrained households—amplifying the differences between optimal policy in HANK and RANK. Second, we study the optimal monetary policy responses to other aggregate shocks, namely *discount-factor shocks* and shocks to the variance of idiosyncratic risk, i.e., *risk shocks*. We show that while the RANK optimal policy features *divine coincidence* (Blanchard and Galí, 2007) in response to these shocks, the HANK optimal policy does not—the HANK planner deviates from implementing productive efficiency and price stability following these shocks in order to reduce inequality, even though productive efficiency and price stability remain feasible.

We also leverage the tractability of our framework to understand some results obtained in the recent quantitative studies of optimal policy in HANK and to uncover the key assumptions driving these results. This literature has found that differences in optimal policy between HANK and RANK are primarily due to households’ unequal exposures to aggregate shocks: monetary policy compensates for missing markets which would provide insurance against aggregate risk (Bhandari et al., 2021) or redistributes from wealth-rich to wealth-poor households (Le Grand et al., 2021). In contrast, both of these studies find that idiosyncratic risk per se has little effect on optimal policy quantitatively. Our results suggest that these papers have not found an important role for idiosyncratic risk because they feature procyclical rather than countercyclical income risk. Indeed, we analytically show that with mildly procyclical income risk, even though households face idiosyncratic consumption risk, monetary policy cannot affect this risk because the income risk and self-insurance channels cancel each other out. Thus, absent unequal exposures to aggregate shocks, optimal monetary policy is identical in HANK and RANK when income risk is mildly procyclical. If we introduced pre-existing wealth inequality or unequal stock ownership into this economy, we would again find differences in optimal policy between HANK and RANK, but these would be exclusively driven by households’ unequal exposure to aggregate shocks. However, the conclusion that idiosyncratic consumption risk does not matter for optimal monetary policy only applies if income risk is procyclical. In the empirically relevant case of countercyclical income risk, monetary policy differs between HANK and RANK both because it mitigates consumption risk and because it either provides insurance against aggregate shocks or redistributes wealth.

Finally, our results relate to the ongoing debate about whether and how central banks should address distributional concerns. Our analysis suggests that a monetary policymaker concerned with inequality does not necessarily need to incorporate an explicit measure of inequality in their reaction function. Instead, distributional concerns can be addressed by stabilizing the *level* of output, in addition to the output gap and the price level. Stabilizing output around a high level can itself lower inequality, both by reducing idiosyncratic risk and by preventing aggregate shocks from adversely impacting more-vulnerable groups.

Related literature The papers most closely related to ours are [Bhandari et al. \(2021\)](#) and [Le Grand et al. \(2021\)](#), who also study optimal monetary policy in HANK models with rich cross-sectional household heterogeneity. One difference between our paper and theirs is methodological: these papers propose numerical algorithms to compute optimal monetary policy, while we study a HANK economy which permits analytical solutions.¹ We see the two approaches as complementary: the first allows more flexibility in the structure of preferences and idiosyncratic shocks, while the second better isolates the channels through which monetary policy optimally affects consumption inequality. Indeed, as discussed above, our analytic approach suggests that the reason these papers do not find a significant role for idiosyncratic risk per se is that they feature procyclical rather than countercyclical income risk. In this regard, an important substantive difference between our paper and theirs is that our model can feature countercyclical income risk. This allows us to study another conceptually distinct way in which monetary policy affects inequality in HANK, which leads optimal policy to deviate from RANK even when the planner does not need to provide insurance against aggregate shocks.

[Nuño and Thomas \(2021\)](#) study how URE and unexpected inflation (Fisher channel) affect optimal monetary policy in the presence of heterogeneity. Unlike us, they study a small open economy in which monetary policy cannot affect real interest rates and output. Thus, the output-inflation trade-off central to New Keynesian models is absent from their setting. While we purposely abstract from the Fisher channel by assuming that households trade real (i.e., inflation-indexed) bonds, an earlier version of this paper did study this channel; its effect on optimal policy is similar to the URE channel discussed in [Section 5.1](#).

Several authors study optimal monetary policy in New Keynesian economies with limited household heterogeneity ([Bilbiie, 2008](#); [Bilbiie and Ragot, 2021](#); [Bilbiie, 2019](#); [Hansen et al., 2020](#); [Challe, 2020](#)).² Most of these papers achieve tractability by imposing the *zero-liquidity limit* (households cannot borrow and government debt is in zero net supply).³ This rules out the self-insurance channel because in equilibrium households do not borrow or lend, spending all their income on consumption.⁴ We discuss how the zero-liquidity assumption affects optimal policy in [Section 6.2](#). More generally, our paper belongs to the literature studying the transmission and optimality of various policies in HANK. Besides the work on conventional monetary policy, this includes studies of unconventional monetary policy ([McKay et al., 2016](#); [Acharya and Dogra, 2020](#); [Bilbiie, 2019](#)), social insurance ([McKay and Reis, 2016, 2021](#); [Kekre, 2019](#)), and fiscal policy ([Auclert et al., 2018](#); [Bilbiie, 2020](#)).

Our analysis suggests that optimal monetary policy differs between HANK and RANK because monetary policy can affect consumption inequality—in particular, when income risk is countercyclical or acyclical, expansionary policy *reduces* consumption inequality. While few papers explicitly study the effect of monetary policy on consumption inequality; this implication is broadly consistent with the available evidence for the US and the UK ([Coibion et al., 2017](#); [Mumtaz and Theophilopoulou, 2017](#)).

The rest of the paper proceeds as follows. [Section 1](#) presents our baseline model. [Section 2](#) characterizes the decentralized equilibrium. [Section 3](#) sets up the planning problem, while [Section 4](#) characterizes

¹[Caballero \(1990\)](#), [Calvet \(2001\)](#), [Wang \(2003\)](#), [Angeletos and Calvet \(2006\)](#) exploit CARA preferences in real economies; [Acharya and Dogra \(2020\)](#) shows that these assumptions are helpful in understanding positive properties of HANK economies.

²See [Nisticò \(2016\)](#), who generalizes the Two-Agent New Keynesian (TANK) model of [Galí et al. \(2007\)](#) and [Bilbiie \(2008\)](#) to the case of stochastic asset-market participation, and [Debortoli and Galí \(2018\)](#) on the comparison between the TANK model and a HANK model with homogeneous borrowing-constrained households and heterogeneous unconstrained households.

³[Bilbiie and Ragot \(2021\)](#) is an exception as it allows agents to hold money in positive amounts for self-insurance purposes.

⁴Strictly speaking, even in economies with zero-liquidity, households may still be able to partially self-insure by adjusting labor supply in response to shocks. What is absent in these models is self-insurance using asset markets.

optimal monetary policy in our baseline economy with idiosyncratic consumption risk. Section 5 studies how unequal exposures to aggregate shocks and policy affect optimal policy. Section 6 extends the model further by introducing hand-to-mouth households and demand shocks. Section 7 concludes.

1 Environment

1.1 Households

We study a Bewley-Huggett economy in which households face uninsurable idiosyncratic shocks to their disutility from supplying labor. We abstract from aggregate risk but allow for a one-time unanticipated aggregate shock at date 0, after which agents have perfect foresight of aggregate variables. Our economy features a perpetual youth structure à la Blanchard-Yaari in which each individual faces a constant survival probability ϑ in any period. We introduce this demographic structure to ensure that the model features a stationary wealth distribution.⁵ Population is fixed and normalized to 1; the size of the cohort born at any date t is $1 - \vartheta$ and the date t size of a cohort born at $s < t$ is $(1 - \vartheta)\vartheta^{t-s}$. The date s problem of an individual i born at date s is

$$\begin{aligned} \max_{\{c_t^s(i), \ell_t^s(i), a_t^s(i)\}} \quad & \mathbb{E}_s \sum_{t=s}^{\infty} (\beta\vartheta)^{t-s} u\left(c_t^s(i), \ell_t^s(i); \xi_t^s(i)\right) \\ \text{s.t.} \quad & c_t^s(i) + q_t a_{t+1}^s(i) = (1 - \tau^w)\tilde{w}_t \ell_t^s(i) + (1 - \tau_t^a)a_t^s(i) + D_t^s(i) - T_t \\ & a_s^s(i) = \mathcal{T}_s \end{aligned} \tag{1}$$

$$a_s^s(i) = \mathcal{T}_s \tag{2}$$

Agents have CARA preferences over both c and (disutility of) labor $\ell - \xi$:

$$u\left(c_t^s(i), \ell_t^s(i); \xi_t^s(i)\right) = -\frac{1}{\gamma} e^{-\gamma c_t^s(i)} - \rho e^{\frac{1}{\rho}\left(\ell_t^s(i) - \xi_t^s(i)\right)} \tag{3}$$

Each agent i saves in riskless real actuarial bonds, issued by financial intermediaries (described below), which have price q_t at date t and have a pre-tax payoff of one unit of the consumption good at $t + 1$ if the agent survives. The government levies a tax τ_t^a at date t on bond holdings $a_t^s(i)$. Unlike many HANK models, our model does not feature a hard borrowing constraint, but we introduce hand-to-mouth households in Section 6. Individuals born at date t receive a transfer \mathcal{T}_t from the government. In addition, all individuals alive at date t pay lump-sum taxes T_t to the government and receive dividends $D_t^s(i)$ from firms. In the baseline model, all households receive an equal share of total dividends; i.e., $D_t^s(i) = D_t$; Section 5.2 considers the case with unequal distribution of dividends.

Given the pre-tax wage \tilde{w}_t and tax rate τ^w , a household supplies labor $\ell_t^s(i)$ at the post-tax real wage $w_t = (1 - \tau^w)\tilde{w}_t$. Households face uninsurable shocks $\xi_t^s(i) \sim N(\bar{\xi}, \sigma_\xi^2)$ to the disutility of supplying labor. $\xi_t^s(i)$ is independent across time and individuals. A larger realization of $\xi_t^s(i)$ reduces disutility and, given wages, increases the household's labor supply. Equivalently, one may think of $\xi_t^s(i)$ as a shock to the household's endowment of time available to supply labor.⁶ To see this, define leisure as $l_t^s(i) = \xi_t^s(i) - \ell_t^s(i)$.

⁵As we discuss in Section 2.1, if we had infinitely lived agents, our model would not feature a stationary wealth distribution.

⁶We thank Gianluca Violante for suggesting this interpretation.

Then one can rewrite utility (3) as $-e^{-\gamma c_t^s(i)}/\gamma - \rho e^{-l_t^s(i)/\rho}$ and the budget constraint as

$$c_t^s(i) + w_t l_t^s(i) + q_t a_{t+1}^s(i) = w_t \xi_t^s(i) + (1 - \tau_t^a) a_t^s(i) + D_t^s(i) - T_t \quad (4)$$

The LHS of (4) denotes the purchases of consumption, leisure and bonds by the household while the RHS denotes the *notional cash on-hand*—the value of the household’s time endowment along with savings net of transfers. Henceforth, we will simply refer to this as *cash-on-hand*. We allow for the possibility that the variance of ξ, σ_t^2 varies endogenously with the level of economic activity, as we discuss later.

1.2 Financial intermediaries

Competitive financial intermediaries trade actuarial bonds with households and hold government debt. Intermediaries only repay households that survive between t and $t + 1$. An intermediary solves

$$\max_{a_{t+1}, B_{t+1}} -\vartheta a_{t+1} + B_{t+1} \text{ s.t. } -q_t a_{t+1} + \Pi_{t+1} \frac{B_{t+1}}{1 + i_t} \leq 0,$$

where B_t denotes government debt, a_t denotes net claims held by households, $R_t = \frac{1+i_t}{\Pi_{t+1}}$ is the real return on government debt, i_t is the nominal interest rate set by the monetary authority and Π_{t+1} denotes inflation between t and $t + 1$. Zero profits require that the intermediary trades bonds with households at price $q_t = \vartheta/R_t$ and that $\vartheta a_{t+1} = B_{t+1}$.

1.3 Final goods producers

A representative competitive final goods firm transforms the differentiated intermediate goods $y_t^j, j \in [0, 1]$ into the final good y according to the CES aggregator $y_t = \left[\int_0^1 y_t(j)^{\frac{1}{\lambda_t}} dj \right]^{\lambda_t}$, where $\frac{\lambda_t}{\lambda_t - 1}$ is the elasticity of substitution between varieties. We allow λ_t (the flexible-price markup) to vary over time in order to introduce “cost-push” shocks. As is standard, the final good producer’s demand for variety j is

$$y_t(j) = \left(\frac{P_t(j)}{P_t} \right)^{-\frac{\lambda_t}{\lambda_t - 1}} y_t \quad (5)$$

1.4 Intermediate goods producers

There is a continuum of monopolistically competitive intermediate goods firms indexed by $j \in [0, 1]$. Each firm faces a quadratic cost $\frac{\Psi}{2} \left(\frac{P_t(j)}{P_{t-1}(j)} - 1 \right)^2$ of changing the price of the variety it produces (Rotemberg, 1982). If firm j hires $n_t(j)$ units of labor, it can only sell to the final goods firm the quantity

$$y_t(j) = z_t n_t(j) - \frac{\Psi}{2} \left(\frac{P_t(j)}{P_{t-1}(j)} - 1 \right)^2, \quad (6)$$

where z_t denotes the level of aggregate productivity at date t . The fiscal authority subsidizes the wage bill of firms at a constant rate τ , so that firm j solves

$$\max_{\{P_t^j, n_t^j, y_t^j\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t \left\{ \frac{P_t(j)}{P_t} y_t(j) - (1 - \tau) \tilde{w}_t n_t(j) \right\},$$

subject to (5) and (6). This yields the standard Phillips curve

$$(\Pi_t - 1) \Pi_t = \frac{\lambda_t}{\Psi(\lambda_t - 1)} \left[1 - \frac{z_t}{(1 - \tau) \lambda_t \tilde{w}_t} \right] + \beta \left(\frac{z_t \tilde{w}_{t+1}}{z_{t+1} \tilde{w}_t} \right) (\Pi_{t+1} - 1) \Pi_{t+1} \quad (7)$$

1.5 Government

The monetary authority sets the interest rate on nominal government debt. At date t , the fiscal authority gives lump-sum transfers \mathcal{T}_t to newborns. The wage bill subsidy is assumed to be equal to $\tau = 1 - \lambda^{-1}$, where λ denotes the steady state markup, eliminating the distortion from monopolistic competition in steady state. These expenditures are financed by issuing debt, taxing bond holdings at a rate τ_t^a and labor income taxes at a rate τ^w . The government budget constraint is

$$\frac{B_{t+1}}{R_t} + T_t + \tilde{w}_t(1 - \vartheta) \sum_{s=-\infty}^t \vartheta^{t-s} \int_i [\tau^w \ell_t^s(i) di + \tau_t^a a_t^s(i)] di = (1 - \vartheta) \mathcal{T}_t + \tau w_t \int_0^1 n_t(j) dj + B_t \quad (8)$$

We further assume that $\mathcal{T}_t = \frac{B_t}{\vartheta}$. This implies that each cohort has the same average wealth, ensuring that the economy features Ricardian equivalence, i.e., the path of government debt is irrelevant for all real allocations.⁷ This allows us to abstract from intergenerational redistribution motives that the Ramsey planner might otherwise have. Consequently, we set $B_t = 0$ for all t without loss of generality.

1.6 Market clearing

In equilibrium, the markets for the final good, labor and assets must clear:

$$\begin{aligned} y_t = c_t &\equiv (1 - \vartheta) \sum_{s=-\infty}^t \vartheta^{t-s} \int_i c_t^s(i) di \\ \int_0^1 n_t(j) dj &= (1 - \vartheta) \sum_{s=-\infty}^t \vartheta^{t-s} \int_i \ell_t^s(i) di \\ 0 = \frac{B_{t+1}}{\vartheta P_t} &= (1 - \vartheta) \sum_{s=-\infty}^t \vartheta^{t-s} \int_i a_{t+1}^s(i) di \end{aligned}$$

1.7 Aggregate shocks

We abstract from aggregate risk but allow for one-time unanticipated aggregate shocks at date 0 to the level of aggregate productivity z_0 and firms' desired markup λ_0 , which decay geometrically: $\ln z_t = \varrho_z^t \ln z_0$, $\ln \left(\frac{\lambda_t}{\lambda} \right) = \varrho_\lambda^t \ln \left(\frac{\lambda_0}{\lambda} \right)$. We discuss additional shocks in Section 6.

2 Characterizing equilibria

As in Acharya and Dogra (2020), CARA utility and normally distributed shocks imply that the model aggregates linearly and the wealth distribution does not directly affect aggregate dynamics. Next, we

⁷See Appendix I for details. In Section 6.1, when we add hand-to-mouth households to the model, Ricardian equivalence no longer holds and our maintained assumption of zero public debt does affect allocations.

describe household decisions. In what follows, we assume that the wealth tax $\tau_t^a = 0$ for all $t > 0$. This is without loss of generality since only the after-tax bond return $R_t(1 - \tau_{t+1}^a)$ affects households' decisions.⁸

Proposition 1. *In equilibrium, the date $t \geq s$ consumption and labor supply decisions of a household i born at date s are*

$$c_t^s(i) = C_t + \mu_t x_t^s(i) \quad (9)$$

$$\ell_t^s(i) = \rho \ln w_t - \gamma \rho c_t^s(i) + \xi_t^s(i), \quad (10)$$

where $x_t^s(i) = (1 - \tau_t^a)a_t^s(i) + w_t(\xi_t^s(i) - \bar{\xi})$ is demeaned cash-on-hand, C_t denotes aggregate consumption and μ_t is the marginal propensity to consume (MPC) out of cash-on-hand. These evolve according to

$$C_t = -\frac{1}{\gamma} \ln \beta R_t + C_{t+1} - \frac{\gamma \mu_{t+1}^2 w_{t+1}^2 \sigma_{t+1}^2}{2} \quad (11)$$

$$\mu_t^{-1} = 1 + \gamma \rho w_t + \frac{\vartheta}{R_t} \mu_{t+1}^{-1} \quad (12)$$

Proof. See Appendix A. □

To understand how market incompleteness affects consumption and labor supply, it is useful to compare (9) and (10) to their counterparts under complete markets. Under complete markets, households are fully insured against disutility shocks; i.e., marginal utility of consumption $e^{-\gamma c_t^s(i)}$ and the marginal disutility of labor $e^{\frac{1}{\rho}(\ell_t^s(i) - \xi_t^s(i))}$ are equalized across all states, implying $\frac{\partial c_t^s(i)}{\partial \xi_t^s(i)} = 0$ and $\frac{\partial \ell_t^s(i)}{\partial \xi_t^s(i)} = 1$: a household with a temporarily higher disutility from working ($\xi_t^s(i) < \bar{\xi}$) can reduce hours without a fall in consumption. Instead, when markets are incomplete (9) and (10) imply that

$$\frac{\partial c_t^s(i)}{\partial \xi_t^s(i)} = \mu_t w_t > 0 \quad \text{and} \quad \frac{\partial \ell_t^s(i)}{\partial \xi_t^s(i)} = 1 - \gamma \rho \mu_t w_t < 1$$

A household with $\xi_t^s(i) < \bar{\xi}$ would like to work less, but reducing hours by as much as under complete markets would cause consumption to drop too much. Thus, the household works longer hours than under complete markets while simultaneously borrowing to mitigate the fall in consumption. However, credit and labor markets provide partial but not full insurance: consumption still falls after an adverse shock.

Households' ability to self-insure using credit and labor markets depends on the future path of interest rates and wages and is measured by the MPC out of cash-on-hand μ_t . Proposition 1 states that μ_t is the same across individuals;⁹ (12) describes its evolution. Iterating this forward yields

$$\mu_t = \left[\sum_{s=0}^{\infty} Q_{t+s|t} (1 + \gamma \rho w_{t+s}) \right]^{-1}$$

μ_t , which measures the pass-through from cash-on-hand to consumption, is increasing in current and future

⁸More specifically, changing the wealth tax at date $t > 0$ from 0 to some $\tau_t^a \neq 0$ does not change the set of allocations that can be implemented by the planner. Starting from an allocation with $\tau_t^a = 0$ where the pre-tax return on bonds between dates $t-1$ and t is R_{t-1} , if the tax rate is changed to $\tau_t^a \neq 0$, changing the pre-tax interest rate to $R_{t-1}/(1 - \tau_t^a)$ keeps the post-tax interest rate and all other prices and allocations unchanged. This is because households have perfect foresight of aggregate variables and only the post-tax real interest rate matters for their decisions.

⁹We introduce MPC heterogeneity in Section 6.

interest rates and decreasing in current and future wages. Lower interest rates reduce the cost of borrowing, making it easier for a household with $\xi_t^s(i) < \bar{\xi}$ to mitigate the decline in consumption by borrowing, and hence reducing μ_t . Similarly, higher future wages reduce the disutility of working more hours in the future since even a small increase in hours worked suffices to repay the same debt, again reducing μ_t .¹⁰

While the sensitivity of household consumption to idiosyncratic income shocks (μ_t) depends on the factors we have just described, average consumption in the economy \mathcal{C}_t depends on interest rates relative to impatience and on households' precautionary motive, as shown in (11). Absent idiosyncratic risk, $\sigma_t = 0$ in (11) and we revert to the RANK Euler equation; higher real interest rates relative to household impatience raise consumption growth. The last term in (11) reflects precautionary savings. Given (9), the conditional variance of date $t + 1$ consumption of household i is $\mathbb{V}_t(c_{t+1}^s(i)) = \mu_{t+1}^2 w_{t+1}^2 \sigma_{t+1}^2$. To the extent that consumption risk is positive and households are prudent ($\gamma > 0$), households save more than in a riskless economy for the same interest rate; i.e., they choose a steeper path of consumption growth. The variance of consumption, in turn, depends on both the variance of cash-on-hand $\mathbb{V}_t(x_{t+1}^s(i)) = w_{t+1}^2 \sigma_{t+1}^2$, and the pass-through of cash-on-hand risk into consumption risk, measured by the (squared) MPC μ_{t+1}^2 .

Determination of y_t In symmetric equilibrium, aggregating (6) across firms, we have $y_t = z_t n_t - \frac{\Psi}{2} (\Pi_t - 1)^2$. Aggregating labor supply (10) and using goods and labor market clearing, we have

$$n_t = \rho \ln w_t - \gamma \rho y_t + \bar{\xi} \quad (13)$$

Combining these two equations, we have

$$y_t = z_t \frac{\rho \ln w_t + \bar{\xi}}{1 + \gamma \rho z_t} - \frac{\Psi}{2} (\Pi_t - 1)^2, \quad (14)$$

where $\frac{\Psi}{2} (\Pi_t - 1)^2$ denotes the resource cost of inflation—deviations of inflation from zero reduce output.

Deriving the aggregate IS equation Imposing goods market clearing in (11) yields the aggregate IS equation which describes the relation between output today and tomorrow:

$$y_t = y_{t+1} - \frac{1}{\gamma} \ln \beta \zeta_t \left(\frac{1 + i_t}{\Pi_{t+1}} \right) - \frac{\gamma \mu_{t+1}^2 w_{t+1}^2 \sigma_{t+1}^2}{2} \quad (15)$$

Time varying σ_t Following McKay and Reis (2021), we allow the variance of ξ to vary endogenously with aggregate output so that the model generates cyclical changes in the distribution of earnings risk. If σ_t were constant, the variance of earnings $w_t^2 \sigma^2$ would inherit the cyclicity of wages; i.e., it would be *procyclical*. In contrast, the empirical literature (Storesletten et al., 2004; Nakajima and Smirnyagin, 2019) generally finds that earnings risk is countercyclical. We assume $\sigma_t^2 w_t^2 = \sigma^2 w^2 \exp\{2\phi(y_t - y)\}$, where y denotes steady state output and $\phi = \frac{\partial \ln \mathbb{V}(x)}{\partial y}$ is the semi-elasticity of the variance of cash-on-hand $\mathbb{V}_t(x)$ w.r.t output. This allows $\mathbb{V}_t(x)$ to be either increasing in y_t (*procyclical* risk) when $\phi > 0$, decreasing in y_t

¹⁰While Acharya and Dogra (2020) already discuss how the MPC responds to future real interest rates, the path of wages has no effect on the MPC in their paper because their environment features inelastic labor supply. In this model, however, since households can choose how much labor to supply, they use this additional margin to self-insure.

(*countercyclical* risk) when $\phi < 0$, or independent of y_t (*acyclical* risk) when $\phi = 0$.¹¹

2.1 Steady state

We now characterize the zero inflation steady state.¹² We normalize the level of productivity $z = 1$ in steady state. Since $\tau = 1 - \lambda^{-1}$, imposing $\Pi_t = \Pi_{t+1} = 1$ in (7) requires that $\tilde{w} = 1$ and hence $w = 1 - \tau^w$. Then, the steady state output is $y = \frac{\rho \ln w + \xi}{1 + \gamma \rho}$. Imposing steady state in (12) and (15) yields

$$R = \beta^{-1} e^{-\frac{\Lambda}{2}} \quad \text{and} \quad \mu = \frac{1 - \tilde{\beta}}{1 + \gamma \rho w},$$

where $\Lambda = \gamma^2 \mu^2 w^2 \sigma^2$ denotes the consumption risk faced by households in steady state (scaled by the coefficient of prudence) and $\tilde{\beta} = \vartheta/R$ is the steady state price of an actuarial bond. Observe that the presence of uninsurable risk ($\Lambda > 0$) implies that the equilibrium real interest rate $R < \beta^{-1}$. Furthermore, the steady state distribution of cash-on-hand x in the population is given by

$$F(x) = (1 - \vartheta) \sum_{s=0}^{\infty} \vartheta^s \Phi \left(\frac{x}{w \sigma \sqrt{s+1}} \right), \quad (16)$$

where $\Phi(\cdot)$ is the cdf of the standard normal distribution. This follows since in steady state, conditional on survival, x is a random walk with no drift whose innovations have variance $w^2 \sigma^2$. If we had infinitely lived agents ($\vartheta \rightarrow 1$), the sum in (16) would diverge and a stationary distribution would not exist.

2.2 Linearized economy

The dynamics of the economy, given a path of interest rates, can be described by the IS equation (15), the MPC recursion (12), the definition of GDP (14) and the Phillips curve (7). These equations summarize the private sector's optimality conditions and define the implementability constraints faced by the planner. Before describing the planner's objective function, it is useful to compare the dynamics of this HANK economy to its RANK counterpart. It is easiest to compare the linearized equations describing aggregate dynamics in the neighborhood of the zero inflation steady state, which are¹³

$$\hat{y}_t = \Theta \hat{y}_{t+1} - \frac{1}{\gamma} (\hat{i}_t - \pi_{t+1}) - \frac{\Lambda}{\gamma} \hat{\mu}_{t+1} \quad (17)$$

$$\hat{\mu}_t = -(1 - \tilde{\beta}) \frac{\gamma \rho w}{1 + \gamma \rho w} \frac{\hat{w}_t}{w} + \tilde{\beta} (\hat{\mu}_{t+1} + \hat{i}_t - \pi_{t+1}) \quad (18)$$

$$\hat{y}_t = \frac{\rho}{1 + \gamma \rho} \frac{\hat{w}_t}{w} + \frac{y}{1 + \gamma \rho} \hat{z}_t \quad (19)$$

$$\pi_t = \beta \pi_{t+1} + \kappa \left(\hat{y}_t - \frac{y + \rho}{1 + \gamma \rho} \hat{z}_t \right) + \frac{\kappa \rho}{1 + \gamma \rho} \hat{\lambda}_t, \quad (20)$$

¹¹More generally, models with labor supply decisions tend to feature procyclical risk while search models tend to feature countercyclical risk. Our assumption that σ_t depends on y_t is a tractable way to incorporate countercyclical risk without incorporating a search model. This also allows us to keep our analysis close to the standard NK model.

¹²We show in Section 3.2 that zero inflation is optimal in steady state.

¹³We linearize y_t and w_t in levels while all other variables are log-linearized. \hat{i}_t is defined as $\ln[(1 + i_t)/R]$.

where $\Theta = 1 - \frac{\Lambda\phi}{\gamma}$ and $\kappa = \left(\frac{1+\gamma\rho}{\rho\Psi}\right) \frac{\lambda}{\lambda-1}$. In RANK, there is no idiosyncratic risk; i.e., $\sigma^2 = 0$, which implies $\Theta = 1$ and $\Lambda = 0$, so that (17) becomes the standard RANK IS curve. Idiosyncratic risk changes the IS equation in two ways. First, as discussed in Acharya and Dogra (2020), countercyclical income risk $\phi < 0$ implies $\Theta > 1$, procyclical income $\phi > 0$ implies $\Theta < 1$ and acyclical income risk implies $\Theta = 1$, reflecting how desired precautionary savings vary with aggregate income and hence the level of income risk. Second, the pass-through, $\hat{\mu}_{t+1}$, also enters the IS curve as it affects desired precautionary savings. In contrast, idiosyncratic risk does not affect the linearized Phillips curve (20), which is the same as in RANK.

3 Setting up the planning problem

3.1 Social welfare function

In our baseline model, we consider a utilitarian planner who attaches equal weight to the lifetime utility of each household i born at date $s \leq 0$, and β^t to the lifetime utility of any household born at a date $t > 0$. In Section 5.1, we relax this assumption and consider more general Pareto weights. The planner's objective can be written as $\sum_{t=0}^{\infty} \beta^t \mathbb{U}_t$, where \mathbb{U}_t , is simply the average utility of all surviving agents:¹⁴

$$\mathbb{U}_t = (1 - \vartheta) \sum_{s=-\infty}^t \vartheta^{t-s} \int u(c_t^s(i), \ell_t^s(i); \xi_t^s(i)) di$$

Given the structure of our economy, this can be decomposed into two parts:

Proposition 2. *The period t felicity function \mathbb{U}_t can be written as*

$$\mathbb{U}_t = u(c_t, n_t; \bar{\xi}) \times \Sigma_t, \quad \text{where} \quad \Sigma_t = (1 - \vartheta) \sum_{s=-\infty}^t \vartheta^{t-s} e^{\frac{1}{2}\gamma^2\sigma_c^2(s,t)}$$

and $\sigma_c^2(s, t)$ denotes the date t variance of consumption among individuals born at date $s \leq t$.

Proof. See Appendix B. □

Intuitively, $u(c_t, n_t; \bar{\xi})$ is the notional flow utility of a representative agent who consumes aggregate consumption c_t , supplies aggregate labor n_t , and faces the mean labor disutility $\bar{\xi}$. Σ_t is the welfare cost of consumption inequality; it is increasing in the variance of consumption, indicating that higher consumption inequality lowers social welfare. Absent risk, there would be no consumption dispersion and hence $\Sigma_t = 1$ at all dates. However, in the presence of risk, $\Sigma_t > 1$, reducing welfare relative to RANK. Recall that

¹⁴Note that the planner discounts felicity \mathbb{U}_t at the same rate as the households themselves. Consider a change in allocations which reduces the date t felicity of cohort s by du_t and increases their date $t+1$ felicity by du_{t+1} , while keeping the felicity at all other dates and for all other agents the same. A cohort s individual will be indifferent regarding this change if $du_t = \beta\vartheta du_{t+1}$. From the planner's perspective this changes aggregate welfare by $-\vartheta^{s-t} du_t + \beta\vartheta^{s+1-t} du_{t+1}$. Thus, the planner will be indifferent about this change if and only if the individuals themselves are indifferent. As discussed by Calvo and Obstfeld (1988), assuming that the planner and the households share the same rate of time preference ensures that social preferences are time-consistent, so that the first-best intertemporal allocation of consumption across cohorts does not change over time.

$u(\cdot) < 0$ and so a higher Σ_t reduces welfare. Appendix B.2 shows that Σ_t evolves according to

$$\ln \Sigma_t = \frac{1}{2} \gamma^2 \mu_t^2 w_t^2 \sigma_t^2 + \ln [1 - \vartheta + \vartheta \Sigma_{t-1}] \quad (21)$$

$$\text{with } \ln \Sigma_0 = \frac{1}{2} \gamma^2 \mu_0^2 w_0^2 \sigma_0^2 + \underbrace{\ln \left[\frac{1 - \vartheta}{1 - \vartheta e^{-\frac{\Lambda}{2} (1 - \tau_0^a)^2 \left(\frac{\mu_0}{\mu}\right)^2}} \right]}_{\text{welfare cost of pre-existing wealth inequality}} \quad (22)$$

Intuitively, higher cash-on-hand risk $w_t^2 \sigma_t^2$ and a higher pass-through μ_t both tend to increase consumption inequality. In addition, consumption inequality inherits the slow moving dynamics of wealth inequality, as can be seen from the presence of Σ_{t-1} in (21).¹⁵

The instruments available to the planner are the sequence of nominal interest rates $\{i_t\}_{t=0}^{\infty}$, which are set optimally in response to shocks, and a date 0 wealth tax τ_0^a and a labor income tax τ^w , which are set optimally absent aggregate shocks but cannot be adjusted in response to shocks.¹⁶ Formally, the timing is as follows. First, the planner chooses sequences $\{w_t, \Pi_t, \mu_t, \Sigma_t, i_t, n_t\}_{t=0}^{\infty}$, together with the date 0 wealth tax τ_0^a and the constant labor income tax τ^w , to maximize $\sum_{t=0}^{\infty} \beta^t u(y_t, n_t; \bar{\xi}) \Sigma_t$, absent aggregate shocks and given an initial wealth distribution. The constraints faced by the planner are the aggregate Euler equation (15), aggregate labor supply (13), the evolution of μ_t (12), the Phillips curve (7), the evolution of Σ_t (21) and the relationship between GDP and wages (14). This Ramsey plan converges to some steady state wealth distribution with a corresponding Σ . Throughout, we always assume that the initial (pre-tax) wealth distribution at the beginning of date 0 corresponds to the steady state of this Ramsey plan.

When studying the monetary policy response to aggregate shocks, the timing is as follows. The economy is initially in the steady state of the Ramsey plan just described, then the fiscal authority imposes the date 0 wealth tax τ_0^a which would be optimal absent aggregate shocks. Next, an unanticipated aggregate shock occurs and the Ramsey planner chooses the sequence of nominal interest rates to maximize social welfare. Formally, the planner chooses sequences $\{w_t, \Pi_t, \mu_t, \Sigma_t, i_t, n_t\}_{t=0}^{\infty}$ to maximize $\sum_{t=0}^{\infty} \beta^t u(y_t, n_t; \bar{\xi}) \Sigma_t$ subject to the constraints (15), (13), (12), (7), (21) and (14). However, the planner cannot adjust taxes in response to the shock; τ_0^a and τ^w are fixed at the level which would be optimal absent aggregate shocks.¹⁷ In the RANK version of our economy, $\sigma = 0$ and (21) is replaced by $\Sigma_t = 1$ for all t . Appendix D presents the Lagrangian associated with this problem along with the first-order necessary conditions for optimality. We begin by describing the optimal choice of fiscal instruments.

¹⁵Note that within-cohort consumption dispersion $\sigma_c^2(t, s)$ rises without bounds as the cohort ages (i.e., as $t - s \rightarrow \infty$), due to the cumulated effect of idiosyncratic shocks on the distribution of cash-on-hand. However, since every cohort gradually shrinks in size, while newborn cohorts have little consumption dispersion (i.e., $\sigma_c^2(t, t) = \mu_t^2 w_t^2 \sigma_t^2$), Σ_t does not necessarily diverge. In fact, provided that the survival rate $\vartheta < e^{-\Lambda/2}$, Σ_t is stationary.

¹⁶Our results would be identical if we allowed the planner to set time varying income and wealth taxes τ_t^w and τ_t^a at all dates, while maintaining the assumption that these cannot be adjusted in response to shocks. This is because the planner would find it strictly optimal to set the same labor income tax rate $\tau_t^w = \tau^w$ at all dates, while the planner is indifferent among all levels of τ_t^a for $t > 0$ (in particular, $\tau_t^a = 0$) since, as mentioned in Section 2, only τ_0^a affects outcomes.

¹⁷Since the shock vanishes in the long run, the steady state of this Ramsey plan with measure 0 aggregate shocks is identical to the steady state of the Ramsey plan with no aggregate shocks.

3.2 Optimal choice of fiscal instruments

Date 0 wealth tax τ_0^a We allow the planner to set a date 0 wealth tax in order to focus on the role of monetary policy in providing insurance, rather than redistribution between borrowers and lenders on average. To understand why, first consider the case in which the planner does not have access to the wealth tax ($\tau_0^a = 0$).

Comparing (22) to (21) shows that the relation between μ_0 and Σ_0 is different than the relation between μ_t and Σ_t at all other dates. Intuitively, at the beginning of date 0, the distribution of wealth is at its steady state level: some households have positive net wealth and some are debtors. Since savers and debtors have different unhedged interest rate exposures (UREs) in the sense of Auclert (2019), an unanticipated change in interest rates affects consumption inequality. Suppose that at date 0, the central bank chooses a policy path that implements a transitory drop in the real interest rate. The lower interest rates benefit debtors, reducing their interest payments and allowing them to increase their consumption. By the same token, a lower path of rates reduces the interest income of savers, causing them to reduce consumption. In other words, lower rates reduce the MPC out of wealth ($\mu_0 \downarrow$), which reduces consumption inequality and hence Σ_0 . To see this, use the fact that $\Sigma_{-1} = \Sigma = \frac{(1-\vartheta)e^{\frac{\Lambda}{2}}}{1-\vartheta e^{\frac{\Lambda}{2}}}$ and $\mu = \mathbb{E}_{-1}\mu_0$ to rearrange (22) as follows:

$$\ln \Sigma_0 = \frac{1}{2}\gamma^2\mu_0^2w_0^2\sigma_0^2 + \ln [1 - \vartheta + \vartheta\Sigma_{-1}] + \underbrace{\ln \left[\frac{1 - \vartheta e^{\frac{\Lambda}{2}}}{1 - \vartheta e^{-\frac{\Lambda}{2}}(1-\tau_0^a)^2\left(\frac{\mu_0}{\mathbb{E}_{-1}\mu_0}\right)^2} \right]}_{\text{effect of date 0 surprise/URE}}$$

While the first two terms on the RHS of the above equation are the same as that in (21), the third term is new. This reflects the fact that an *anticipated* cut in rates would not reduce inequality as much as this unanticipated cut. If wealthy agents at date -1 had anticipated lower rates at date 0, they would have saved more in order to insure a higher level of consumption at date 0. Equally, the poor debtors would have borrowed more at date -1, knowing that their debt would be less costly to repay. For this reason, what reduces Σ_0 through this channel is not a fall in μ_0 per se but a fall in μ_0 relative to its expected value $\mathbb{E}_{-1}\mu_0$, as can be seen from the last term in (22). To be clear, anticipated cuts in rates *do* reduce inequality, as discussed earlier: the lower μ_t reduces Σ_t in equation (21). But there is an additional effect that comes from a surprise fall in interest rates. In our environment, since we do not have aggregate risk (only unanticipated shocks at date 0), the fact that the Ramsey planner is only allowed to reoptimize at date 0 implies that this additional affect of an unanticipated change in μ can only occur at date 0.

Absent wealth taxes, the utilitarian planner would exploit the channel just described to *redistribute* consumption between borrowers and lenders at date 0, making the conduct of optimal monetary policy different at date 0 than at all subsequent dates.¹⁸ However, the planner also has another instrument which can be used to redistribute from lenders to borrowers, namely the wealth tax. This instrument is less flexible than monetary policy since it cannot be set in a state contingent way. However, it is optimal for the utilitarian planner to set this tax at a level ($\tau_0^a = 1$) which completely eliminates pre-existing wealth inequality, setting the second term in (22) to zero.¹⁹ This not only eliminates the incentive of monetary

¹⁸A previous version of this paper studied the case in which the planner does not have recourse to wealth taxes and monetary policy exploits the URE channel at date 0 to engineer consumption redistribution.

¹⁹See Appendix D.1 for a proof.

policy to deliver a surprise rate cut absent shocks, it also implies that households are equally exposed to aggregate shocks at date 0. Consequently, all consumption inequality going forwards is the result of uninsurable idiosyncratic risk, not unequal exposures to aggregate shocks ex ante, and any differences between HANK and RANK arise purely due to idiosyncratic risk. In particular, since wealth is equalized across households at date 0, the URE channel is not operative and the relation between μ_t and Σ_t is the same at date 0 as at all other dates $t > 0$.²⁰

Labor income tax We also allow the planner to set a set a proportional labor income tax τ^w . This tax is set optimally absent aggregate shocks implying that the planner does not need to use monetary policy to affect inequality on average and optimally sets zero inflation in steady state.²¹ However, the tax cannot be adjusted in response to aggregate shocks, reflecting the idea that fiscal policy is slow to adjust.²² Thus, monetary policy still has a role in dealing with changes in inequality in response to aggregate shocks.

Setting $\tau^w \neq 0$ introduces productive inefficiency by driving a wedge between post-tax wages and the marginal product of labor ($z = 1$) in steady state. Equivalently, it drives a wedge between steady state output y and its productively efficient level $y^e = \frac{\bar{\xi}}{1+\gamma\rho}$, which we will normalize to one by setting $\bar{\xi} = 1 + \gamma\rho$. The planner trades off this cost against the benefits of reducing consumption inequality. Appendix D.2 shows that this trade-off can be summarized by the following optimality condition:

$$\underbrace{\Omega}_{\text{benefit from reduction of inequality}} \equiv \underbrace{\frac{\Lambda}{(1-\beta)(1-\Lambda)}}_{\text{reduction of inequality due to reduced pass-through from income to consumption risk } (\Omega^c)} + \underbrace{\frac{\Theta - 1}{(1-\beta)(1-\Lambda)}}_{\text{reduction of inequality due to reduced income risk}} = \underbrace{\frac{w - 1}{1 + \gamma\rho w}}_{\text{cost of deviating from productive efficiency}}, \quad (23)$$

which implies that the optimal income tax is $\tau^w = 1 - \frac{1+\Omega}{1-\gamma\rho\Omega}$. Ω summarizes the benefit from a reduction in consumption inequality due to higher economic activity. In RANK ($\Lambda = 0, \Theta = 1$), there is no inequality and thus no benefit from reducing inequality ($\Omega = 0$), so that $w = 1$ or $\tau^w = 0$ is optimal. In the presence of risk, higher output (implemented via lower τ^w) affects consumption inequality through both a *self-insurance* channel and an *income-risk channel*. (23) states that at an optimum, the marginal benefit of lower inequality due to higher output through both these channels, Ω , equals the marginal cost of distorting productive efficiency which is proportional to the gap between wages and the marginal product of labor.

Consider first the self-insurance channel. With acyclical income risk ($\Theta = 1$) the level of economic activity does not affect household income risk. Thus, increasing output above y^e does not reduce income inequality (second term on the LHS of (23) is zero). However, higher output and wages still make it easier to self-insure through the labor market and reduce the pass-through from income shocks into consumption, measured by the first term of the LHS, reducing consumption inequality. Thus, even with acyclical risk $\Omega = \Omega^c \equiv \frac{\Lambda}{(1-\beta)(1-\Lambda)} > 0$, the planner subsidizes labor ($\tau^w < 0$) to raise steady state output above y^e .

Next, consider the income-risk channel. With countercyclical income risk ($\Theta > 1$), pushing output

²⁰This result is special to the case of the utilitarian planner. In Section 5.1, we show that a non-utilitarian planner optimally sets the wealth tax at a level which does not completely eliminate pre-existing wealth inequality. Thus, while the wealth tax removes the incentive for monetary policy to create a surprise rate cut on average, optimal policy does exploit the URE channel in a state contingent way in response to aggregate shocks, making optimal monetary policy different at dates 0 and $t > 0$.

²¹See Appendix D.1 for a proof.

²²We do not let τ^w be state contingent since this would allow the planner to ignore the Phillips curve as a constraint at all dates (Correia et al., 2008) and implement any desired path of output while keeping inflation at zero.

above y^e lowers income risk, reducing consumption inequality even for a fixed μ . In addition, higher output reduces μ , further reducing consumption inequality. Thus, the benefit from higher output is even larger than if $\Theta = 1$ —both LHS components in (23) are positive and Ω is larger ($\Omega > \Omega^c$). Consequently, the planner subsidizes labor income even more, pushing steady state output further above y^e .

With procyclical income risk ($\Theta < 1$), the effect of higher output on consumption inequality is ambiguous. Higher output still facilitates self-insurance ($\Lambda > 0$), but now increases income risk ($\Theta - 1 < 0$). For sufficiently procyclical risk, the second effect dominates, $\Omega < 0$ and the optimal steady state output is below y^e , implemented with a tax $\tau^w > 0$. For mildly procyclical risk, the self-insurance channel dominates and $\Omega > 0$ with $\tau^w < 0$. The two channels perfectly offset each other if $1 - \Theta = \Lambda$, implying $\Omega = 0$; higher output has no first-order effect on consumption inequality and the planner does not distort productive efficiency in steady state, setting $\tau^w = 0$. $\Omega = 0$ will be a useful benchmark in what follows.

3.3 How does monetary policy affect inequality?

Besides the steady state relation between output and inequality, which influences the planner's choice of τ^w , there is also a short-run relationship between output and inequality, which governs how monetary policy optimally responds to aggregate shocks. The evolution of consumption inequality depends on consumption risk $\mu_t^2 \sigma_t^2$, which in turn depends on both income risk σ_t^2 and pass-through μ_t^2 . This can be seen by linearizing (21) and using our assumptions about σ_t :

$$\frac{\widehat{\Sigma}_t}{\Sigma} = \underbrace{-\gamma(\Theta - 1)\widehat{y}_t + \Lambda\widehat{\mu}_t}_{\text{consumption risk}} + \left(\frac{\widetilde{\beta}}{\beta}\right) \frac{\widehat{\Sigma}_{t-1}}{\Sigma} \quad (24)$$

Thus, one way monetary policy can reduce consumption inequality is by affecting the level of output; e.g., with countercyclical income risk ($\Theta > 1$) higher output reduces income risk. Another way is to reduce the pass-through from income to consumption (μ_t). But the planner cannot vary y_t and μ_t independently since they only have one instrument—the nominal interest rate. Lower rates reduce the pass-through but increase output. If income risk is countercyclical $\Theta > 1$, this too reduces consumption risk. However, if income risk is procyclical $\Theta < 1$, then it increases income risk, leaving the overall effect on consumption risk unclear. Using the IS equation (17), the μ recursion (18) and the definition of GDP (19), (24) can be rewritten as²³

$$\frac{\widehat{\Sigma}_t}{\Sigma} = -\gamma(1 - \widetilde{\beta})\Omega\widehat{y}_t + \Lambda \frac{(1 + \Omega)(1 - \widetilde{\beta})}{1 - \widetilde{\beta}\varrho_z(1 - \Lambda)} \frac{\gamma y}{1 + \gamma\rho} \widehat{z}_t + \left(\frac{\widetilde{\beta}}{\beta}\right) \frac{\widehat{\Sigma}_{t-1}}{\Sigma} \quad (25)$$

Just as Ω in (23) measures how much higher output reduces Σ in steady state, (25) shows that it also determines how much a transitory increase in output reduces $\widehat{\Sigma}_t$, holding productivity and other shocks fixed. In this sense $-\gamma(1 - \widetilde{\beta})\Omega$ can be thought of as the *cyclicality of consumption risk*, which generally differs from the cyclicality of *income risk*. When income risk is countercyclical ($\Theta > 1$), expansionary monetary policy reduces consumption inequality both by reducing the pass-through and by increasing output and reducing income risk. Thus, consumption risk is also countercyclical: $\partial(\widehat{\Sigma}_t/\Sigma)/\partial\widehat{y}_t = -\gamma(1 - \widetilde{\beta})\Omega < 0$. Even when income risk is acyclical ($\Theta = 1$), expansionary policy still reduces the pass-through and in-

²³See Appendix E.1 for a derivation.

equality; i.e., consumption risk is still countercyclical, $-\gamma(1 - \tilde{\beta})\Omega < 0$. When income risk is strongly procyclical ($\Theta \ll 1 \Rightarrow \Omega < 0$), consumption risk is also procyclical: expansionary monetary policy increases inequality as the benefit of a lower pass-through is outweighed by the increase in income risk, $\partial(\tilde{\Sigma}_t/\Sigma)/\partial\hat{y}_t = -\gamma(1 - \tilde{\beta})\Omega > 0$. Finally, when $\Omega = 0$, consumption risk is acyclical, and monetary policy has no first-order effect on inequality: higher output increases income risk but this is exactly balanced out by the lower pass-through, implying that monetary policy cannot affect consumption risk.

3.4 Productive efficiency and the output gap

In addition to reducing consumption inequality, the HANK planner still has the standard objectives of productive efficiency—which would be attained by setting output at its productively efficient level y_t^e —and price stability—which would be attained by setting output at its *natural level* y_t^n . y_t^e is only affected by changes in productivity:

$$y_t^e = z_t \frac{\rho \ln z_t + \bar{\xi}}{1 + \gamma \rho z_t}$$

In contrast, y_t^n —the level of output that would prevail under flexible prices—depends on both productivity and markup shocks:

$$y_t^n = y_t^e + \frac{\rho z_t}{1 + \gamma \rho z_t} \left[\ln(1 - \tau^w) - \ln\left(\frac{\lambda_t}{\lambda}\right) \right]$$

y_t^n is a useful benchmark because setting $y_t = y_t^n$ at all dates would implement zero inflation. This can be seen by rewriting the linearized Phillips curve:

$$\pi_t = \beta \pi_{t+1} + \kappa (\hat{y}_t - \hat{y}_t^n), \quad (26)$$

where $\hat{y}_t^n = \frac{\rho+y}{1+\gamma\rho} \hat{z}_t - \frac{\rho}{1+\gamma\rho} \hat{\lambda}_t$ describes the linearized dynamics of y_t^n .

Note that y_t^e and y_t^n will generally not coincide even absent markup shocks. In steady state, output equals its natural level since the planner implements zero inflation. Unless $\Omega = 0$, this level differs from the productively efficient level y^e since the planner sets a non-zero income tax τ^w , driving a permanent wedge between y_t^n and y_t^e . Away from steady state, markup shocks also drive a temporary wedge between y_t^n and y_t^e as in RANK.

The welfare relevant measure of productive inefficiency in our model is the *labor wedge*; i.e., the ratio between household's marginal rate of substitution between consumption and leisure and the marginal productivity of labor, and is simply given by $\varpi_t = w_t/z_t$. Up to first order, the log-deviation of the labor wedge from its steady state value can be expressed as $\hat{\varpi}_t = \frac{1+\gamma\rho}{\rho} \left(\hat{y}_t - \frac{y+\rho}{1+\gamma\rho} \hat{z}_t \right)$. Note that this is proportional to $\hat{y}_t - \frac{y+\rho}{1+\gamma\rho} \hat{z}_t$, which we refer to as the *output gap*; in what follows, we will use this to measure the deviation of productive inefficiency from its steady state level.²⁴ Absent markup shocks, this output gap is given by $\hat{y}_t - \hat{y}_t^n$ and absent productivity shocks, it is simply \hat{y}_t .

3.5 Calibration

While our results are primarily analytical, when plotting the IRFs we parameterize the model as follows. We choose $\bar{\xi}$ to normalize steady state output y to 1 in the HANK economy with $\Omega = 0$ (equivalently,

²⁴To be clear, a zero output gap ($\hat{y}_t - \frac{y+\rho}{1+\gamma\rho} \hat{z}_t = 0$) does not imply that the economy is productively efficient; it merely implies that the degree of productive inefficiency is at its steady state level.

in RANK). We calibrate the model to an annual frequency and target $r = 4\%$. We choose the standard deviation of $\xi_t^s(i)$, σ , so that the standard deviation of income in steady state equals 0.5.²⁵ This is in line with [Guvenen et al. \(2014\)](#) who use administrative data and find that the standard deviation of 1 year log earnings growth rate to be slightly above 0.5. We set the slope of the Phillips curve $\kappa = 0.1$, and the elasticity of substitution between varieties $\frac{\lambda}{\lambda-1}$ to 10, implying a 10% steady state markup, $\lambda = 1.1$. We set $\gamma = 2$, which implies that the coefficient of relative risk aversion, $-\frac{cu''(c)}{u'(c)}$, is 2 in RANK. Similarly, we set $\rho = 1/3$, implying that the Frisch elasticity in RANK is 1/3, which is within the range of estimates from the micro literature. We set the persistence of productivity and markup shocks to $\rho_z = 0.95^4$ and $\rho_\lambda = 0.9^4$, respectively ([Bayer et al., 2020](#)). When plotting impulse responses, we show the response to a one standard deviation shock; we set the standard deviation of productivity and markup shocks $\sigma_z = 0.012$ and $\sigma_\lambda = 0.034$ following [Bayer et al. \(2020\)](#). We set $\vartheta = 0.85$, similar to [Nisticò \(2016\)](#) and [Farhi and Werning \(2019\)](#). Finally, when we consider the case with countercyclical risk, we set $\phi = -5.3$.²⁶

4 Dynamics under optimal monetary policy

We are now in a position to characterize how monetary policy optimally responds to aggregate shocks. As is common in the NK literature, we characterize optimal policy by linearizing the first-order conditions arising from the planner's Lagrangian (presented in Appendix E.2). As in RANK, optimal policy can be described in terms of a *target criterion* that characterizes the optimal path of output and inflation.

Proposition 3 (Target criterion). *Optimal monetary policy sets nominal interest rates so that the following target criterion holds at all dates $t \geq 0$:*

$$\underbrace{\left(1 - \delta(\Omega)\right)\widehat{y}_t}_{\substack{\text{output stabilization} \\ \text{due to inequality} \\ \text{concerns}}} + \underbrace{\delta(\Omega)\left(\widehat{y}_t - \frac{y + \rho}{1 + \gamma\rho}\widehat{z}_t\right)}_{\substack{\text{output gap stabilization} \\ \text{due to productive-efficiency} \\ \text{concerns}}} + \underbrace{\frac{1 + \gamma\rho}{\Upsilon(\Omega)}\frac{\lambda}{\lambda - 1}\widehat{\pi}_t}_{\text{price stability}} = 0 \quad (27)$$

$\Upsilon(\Omega)$ and $\delta(\Omega)$ are defined in Appendix E.3 and satisfy $\Upsilon(0) = \delta(0) = 1$. When income risk is acyclical or countercyclical ($\Theta \geq 1 \Rightarrow \Omega \geq \Omega^c$), $\Upsilon(\Omega) > 1$ and $\delta(\Omega) \in (0, 1)$.

Proof. See Appendix E.3. □

Along with the linearized Phillips curve (26), (27) fully characterizes the optimal dynamics of output \widehat{y}_t and inflation π_t . To understand this target criterion, it is useful to compare it to its RANK counterpart.

Corollary 1. *Absent idiosyncratic risk ($\sigma = 0 \Rightarrow \Omega = 0$), optimal monetary policy in RANK sets the*

²⁵The standard deviation of income is given by $(1 - \gamma\rho\mu w)\omega\sigma$.

²⁶This is broadly consistent with [Storesletten et al. \(2004\)](#), who find that the standard deviation of the persistent shock to (log) household income increases from 0.12 to 0.21 as the economy moves from peak to trough. If the difference between growth in expansions and recessions is roughly 0.03, this implies that $y\frac{d\sigma_y}{dy} = \frac{0.12-0.21}{0.03} = -3$. Using $\sigma_{y,t} = (1 - \gamma\rho\mu_t w_t)\omega\sigma e^{\phi(y_t - y)}$ and the equilibrium relationship between μ_t , w_t and y_t , we have:

$$\phi = \frac{y}{\sigma_y} \frac{d\sigma_y}{dy} + \frac{\gamma\mu w}{1 - \gamma\rho\mu w} \left[1 - \gamma\rho(1 - \tilde{\beta})\Omega\right] y$$

Given our calibration, $\phi = -5.3$ implies $y\frac{d\sigma_y}{dy} = -2.99$, roughly in line with the implications of [Storesletten et al. \(2004\)](#).

nominal rate so that the following target criterion holds at all dates $t \geq 0$:

$$\left(\hat{y}_t - \frac{1 + \rho}{1 + \gamma\rho} \hat{z}_t \right) + (1 + \gamma\rho) \frac{\lambda}{\lambda - 1} \hat{p}_t = 0 \quad (28)$$

(28) is a special case of (27): absent idiosyncratic income risk $\sigma = 0$, we have $\Lambda = 0$ and $\Theta = 1$, implying $\Omega = 0$, $y = y^e = 1$ and $\Upsilon(0) = \delta(0) = 1$. The RANK planner has two objectives: productive efficiency, which would be attained by a zero output gap (first term in brackets in (28)), and price stability, which would be attained by setting $\hat{p}_t = 0$ (last term in (28)). Besides these two objectives, the HANK planner has a third objective: reducing consumption inequality. When income risk is acyclical or countercyclical ($\Omega \geq \Omega^c > 0$), this motive leads to two key differences between the HANK and RANK target criteria. First, the HANK planner puts some weight on stabilizing the *level* of output around its steady state level (which is higher than in RANK), rather than purely trying to minimize the output *gap*—(27) has a weight $1 - \delta(\Omega)$ on deviations of output from steady state, and a weight $\delta(\Omega)$ on the output gap. In our calibration with countercyclical income, $\delta = 0.6$, implying that the planner puts roughly equal weight on output and output gap stabilization. As we will see, this implies that the HANK planner does not close the output gap in response to productivity shocks. Second, the planner puts less weight on price stability, relative to either output or the output gap, compared to the RANK planner: the weight on \hat{p}_t is scaled down by a factor $\Upsilon(\Omega) > 1$. In our calibration, $\Upsilon = 1.76$, implying that the weight on price stability is almost halved relative to RANK. Thus, the HANK planner will tolerate higher fluctuations in inflation and smaller output fluctuations. Interestingly, while the difference between optimal policy in HANK and RANK stems from the HANK planner’s concern for inequality, it is not necessary to include a measure of inequality such as $\hat{\Sigma}_t$ in the target criterion; this concern for inequality is captured by the higher weight on output stabilization and the lower weight on price stability. Intuitively, this is because changes in output induced by monetary policy affect inequality; as described in Section 3.3: when income risk is acyclical or countercyclical, higher output induced by expansionary policy reduces consumption inequality.

Importantly, optimal policy does not differ from RANK simply because inequality exists but because monetary policy can affect inequality. When $\sigma > 0$ and risk is mildly procyclical $\Theta = 1 - \Lambda < 1$, we have $\Omega = 0$ and monetary policy cannot affect the evolution of consumption inequality (25). In this case, $\delta = \Upsilon = 1$ and the target criterion is the same as in RANK (28). Consequently, optimal monetary policy implements the same path of output and inflation in RANK and HANK with $\Omega = 0$, even if the nominal rate path required to implement this sequence is different in both economies.

Lemma 1. *In HANK with $\Omega = 0$, optimal monetary policy implements the same sequence $\{\hat{y}_t, \pi_t\}$ as in RANK and the target criterion is described by (28).*

Outside of this special case, monetary policy can affect inequality and optimal policy in HANK differs from that in RANK. We first consider the optimal response to productivity shocks in HANK when $\Omega \neq 0$.

4.1 Productivity shocks

We first assume there are no markup shocks, so that the output gap appearing in the target criterion is simply $\hat{y}_t - \hat{y}_t^n$. Given our maintained assumption that the subsidy τ eliminates steady state monopolistic distortions, this implies that RANK features *divine coincidence*; i.e., it is both feasible and optimal to

implement zero inflation while keeping output at its efficient level in response to productivity shocks: (28) implies that $\hat{y}_t = \hat{y}_t^n$ and $\hat{p}_t = 0$ under optimal policy. Figure 1 plots the optimal response to a date 0 productivity shock in RANK (red dashed line) and HANK (blue solid line).²⁷ The red dashed lines in panels a and b show that the RANK planner responds to a fall in productivity, which decreases $\hat{y}_t^n = \frac{y+\rho}{1+\gamma\rho} \hat{z}_t < 0$ by tracking this level, $\hat{y}_t = \hat{y}_t^n < 0$, resulting in zero inflation and achieving both productive efficiency and price stability.

Since the HANK planner has an additional objective—reducing inequality—while they *could* implement $\hat{y}_t = \hat{y}_t^n$ and $\pi_t = 0$, they would not do so whenever $\Omega \neq 0$. With acyclical or countercyclical income risk, optimal policy responds to a fall in productivity by preventing output \hat{y}_t from initially falling as much as the flexible price level of output \hat{y}_t^n . This entails positive inflation initially. In contrast, the planner commits to mildly negative output gaps ($\hat{y}_t < \hat{y}_t^n < 0$) in the future, which in turn entail mild deflation in the future. This is formalized in the following proposition.

Proposition 4. *Under optimal policy with acyclical or countercyclical income risk, following a fall in productivity ($\hat{z}_0 < 0$), at date 0, \hat{y}_0 falls less than \hat{y}_0^n and there is inflation, $\pi_0 > 0$. In addition, there exists $T > 0$ such that for all $t \in (T, \infty)$, $\pi_t < 0$ and $\hat{y}_t < \hat{y}_t^n$. Following an increase in productivity all these signs are reversed; i.e., π_t and $\hat{y}_t - \hat{y}_t^n$ are negative at date 0 and positive for all $t \in (T, \infty)$ for some $T > 0$.*

Proof. See Appendix F. □

To understand why optimal monetary policy cushions the fall in output following a fall in productivity, it is useful to reiterate why monetary policy does *not* raise \hat{y}_t above $\hat{y}_t^n = 0$ absent aggregate shocks. With acyclical or countercyclical income risk, increasing \hat{y}_t has a first-order benefit, even absent shocks, as it reduces consumption inequality. But this benefit is exactly offset by the first-order cost of raising output further above its productively efficient level. Recall that with acyclical or countercyclical income risk, the planner subsidizes labor supply, pushing wages w above the marginal product of labor z , resulting in a non-zero labor wedge $\ln(w/z) > 0$, and raising output above the productively efficient level in steady state.

Now suppose that following a negative productivity shock, monetary policy continued to set $\hat{y}_t = \hat{y}_t^n < 0 \forall t \geq 0$ (also implying $\pi_t = 0 \forall t \geq 0$). The fall in \hat{y}_t would raise consumption inequality as shown by the black dotted line in panel c of Figure 1. This raises the first-order benefit of marginally increasing output above \hat{y}_t^n to curtail the rise in inequality. Meanwhile, at $\hat{y}_t = \hat{y}_t^n$ the cost of marginally increasing output above \hat{y}_t^n , measured by the labor wedge $\varpi_t \approx \varpi + \frac{1+\gamma\rho}{\rho} (\hat{y}_t - \hat{y}_t^n)$, remains unchanged. Since the benefits of increasing \hat{y}_t above \hat{y}_t^n increases while the cost of doing so remains unchanged, the planner sets $\hat{y}_t > \hat{y}_t^n < 0$. Output still falls on impact, but by less than the flexible-price level of output y_t^n (blue curve in panel a). This trade-off is also reflected in the target criterion (27), which can be rewritten as

$$\left(\hat{y}_t - \delta(\Omega) \hat{y}_t^n \right) + \frac{1+\gamma\rho}{\Upsilon(\Omega)} \frac{\lambda}{\lambda-1} \hat{p}_t = 0$$

Intuitively, rather than tracking \hat{y}_t^n one-for-one, which would stabilize the labor wedge, the planner seeks to minimize the gap between \hat{y}_t and $\delta(\Omega) \hat{y}_t^n$ (where $\delta(\Omega) \in (0, 1)$). This reflects a compromise between the

²⁷Here and in the figures that follow, unless otherwise specified, we plot the log-deviations from steady state of all variables except the output gap in panel a, where we plot $(\hat{y}_t - \hat{y}_t^n)/y$.

planner's goal of reducing inequality, which calls for stabilizing output around its high steady state level, and fostering productive efficiency.

To implement the milder fall in output, the planner commits to a lower path of nominal rates (blue curve in panel e) relative to RANK (red dashed curve in panel e). This leads to a smaller increase in the pass-through from income to consumption risk (blue curve in panel f) than would occur if monetary policy set $\hat{y}_t = \hat{y}_t^n$ and $\pi = 0$ (black dashed curve in panel f). Given the higher path of \hat{y}_t and lower path of $\hat{\mu}_t$, while inequality still increases (blue curve in panel c), it is lower than it would have been had the planner implemented $\hat{y}_t = \hat{y}_t^n$ and $\pi_t = 0$ for all $t \geq 0$ (black dotted curve in panel c).

Implementing a higher \hat{y}_t than \hat{y}_t^n results in inflation early on. The planner tolerates this higher inflation in order to cushion the fall in output, as can be seen from the lower weight on price stability in (27) (since $\Upsilon(\Omega) > 1$). Nonetheless, to mitigate this rise in inflation, the planner commits to \hat{y}_t slightly below \hat{y}_t^n in the future. This commitment lowers inflation not just in the future but also at date 0 because of the forward-looking nature of the Phillips curve.




Figure 1: **Optimal policy in response to productivity shocks** in HANK with $\Omega > 0$ (blue solid curves) and RANK (red dashed curves). Black dotted curves denote outcomes in HANK under the non-optimal policy which sets $\hat{y}_t = \hat{y}_t^n$ and $\pi_t = 0$ for all $t \geq 0$. All panels plot log-deviations from steady state except panel a which plots $(\hat{y}_t - \hat{y}_t^n)/y$.

4.2 Markup shocks

We now discuss the optimal response to markup shocks, abstracting from productivity shocks ($\hat{z}_t = 0$). This implies that the output gap in the target criterion (27) is simply $\hat{y}_t - \frac{y+\rho}{1+\gamma\rho}\hat{z}_t = \hat{y}_t$. Recall that even in RANK, cost-push shocks such as an increase in firms' desired markups break divine coincidence. Monetary policy can no longer maintain zero inflation while keeping output at its productively efficient level since markup shocks drive a wedge between the productively efficient level y_t^e (which remains unchanged) and the level of output consistent with zero inflation y_t^n (which declines, $\hat{y}_t^n = -\frac{\rho}{1+\gamma\rho}\hat{\lambda}_t < 0$). Keeping $\pi_t = 0$ by

setting $\hat{y}_t = \hat{y}_t^n < 0$ is not optimal as this would entail too large a fall in output relative to its efficient level $\hat{y}_t^e = 0$. Conversely, keeping output at its efficient level $\hat{y}_t = \hat{y}_t^e = 0$ is not optimal as this would entail too much inflation. Consequently, the RANK planner responds to a positive markup shock by permitting some fall in output (red dashed curve in panel a in Figure 2) and some increase in inflation (red dashed curve in panel b). Monetary policy also commits to keep \hat{y}_t below \hat{y}_t^n in the future, resulting in mild deflation. Given the forward-looking Phillips curve, this further mitigates the initial increase in inflation.

In HANK with acyclical or countercyclical income risk, inflation remains costly and so optimal policy still does not perfectly stabilize output ($\hat{y}_t = 0$) following a positive markup shock. However, the welfare effects of a fall in output are different from RANK in two respects. First, since output is above its productively efficient level in steady state $y > y^e$, a fall in output *improves* productive efficiency. Second, a fall in output increases consumption inequality, reducing welfare. Proposition 5 shows that the second effect always dominates: optimal policy in HANK with acyclical or countercyclical risk allows a larger increase in inflation and a smaller fall in output than in RANK. This can also be seen by specializing the target criterion (27) to the case with only markup shocks:

$$\hat{y}_t + \frac{1 + \gamma\rho}{\Upsilon(\Omega)} \frac{\lambda}{\lambda - 1} \hat{p}_t = 0$$

The HANK target criterion has a higher weight on output stabilization (relative to inflation) compared to RANK: $\Upsilon = 1$ in RANK but $\Upsilon > 1$ in HANK with countercyclical risk.²⁸

Proposition 5. *Under optimal policy with acyclical or countercyclical income risk, following an increase in firms' desired markup ($\hat{\lambda}_0 > 0$), at date 0, \hat{y}_0 falls (but by less than \hat{y}_0^n) and $\pi_0 > 0$. In HANK with acyclical or countercyclical income risk, the fall in output is smaller than in RANK, and the increase in inflation is larger. In addition, there exists $T > 0$ such that for all $t \in (T, \infty)$, $\pi_t < 0$ and $\hat{y}_t - \hat{y}_t^n < 0$. Following a fall in desired markups all the signs are reversed.*

Proof. See Appendix F. □

Figure 2 plots the IRFs following a positive markup shock under optimal policy in RANK (red dashed curves) and HANK with acyclical or countercyclical income risk (blue solid curves). The RANK planner already permits some fall in output and increase in inflation on impact; the HANK planner allows even higher inflation to mitigate the fall in output. Allowing output to fall as much as in RANK is undesirable as it would result in higher inequality (black dotted curve which lies above the blue solid curve in panel c). To implement the smaller decline in output, the HANK planner commits to a shallower path of nominal rates (panel e), which also translates into a smaller increase in pass-through $\hat{\mu}_t$ (panel f). As in RANK, the HANK planner commits to modest deflation in the future to mitigate the initial rise in inflation.

4.3 Implementing optimal policy using an interest rate rule

Equation (27) describes optimal monetary policy in terms of a *targeting rule* rather than an *instrument rule*. Following Galí (2015), it is easy to construct an interest rate rule which uniquely implements optimal

²⁸Since $\hat{z}_t = 0$, $\delta(\Omega)$, which governs the relative weight on output stabilization and output gap stabilization, plays no role here: stabilizing output to reduce inequality is equivalent to stabilizing the output gap for productive-efficiency reasons.




Figure 2: **Optimal policy in response to markup shocks** in HANK with $\Omega > 0$ (blue solid curves) and RANK (red dashed curves). Black dotted lines denote outcomes in HANK under non-optimal policy which implements the same path of \hat{y}_t and π_t as in RANK. All panels plot log-deviations from steady state.

allocations and inflation. One such interest rate rule is

$$i_t = i_t^* + \phi \pi_t + \phi_x \Delta x_t, \quad (29)$$

where $\phi_x = \phi \frac{\Upsilon(\Omega)}{1+\gamma\rho} \left(\frac{\lambda-1}{\lambda}\right)$ and $i_t^* \equiv \Phi_x x_{t-1} + \Phi_z \hat{z}_t + \Phi_\lambda \hat{\lambda}_t$ denotes the equilibrium nominal interest rate under optimal policy, and $x_t \equiv \hat{y}_t - \delta(\Omega) \frac{y+\rho}{1+\gamma\rho} \hat{z}_t$ denotes the gap between output and its desired level, $\delta(\Omega) \frac{y+\rho}{1+\gamma\rho} \hat{z}_t$.²⁹ Appendix F.4 shows that for a sufficiently large ϕ , this rule implements the optimal allocations as a unique equilibrium. As in RANK, nominal rates still respond to changes in inflation and the gap between output and its desired level. However, with acyclical or countercyclical income risk (implying $\Upsilon > 1$), this rule reacts more strongly to changes in Δx_t , relative to π_t , compared to the corresponding rule in RANK where $\Upsilon = 1$ (ϕ_x is increasing in Υ). Also, the desired level of output is not simply its efficient level as in RANK: the desired level moves less than one-for-one with $\frac{y+\rho}{1+\gamma\rho} \hat{z}_t$. Note that (29) does not require the monetary policymaker to change nominal rates in response to changes in some measure of inequality; the concern for inequality is captured by a larger coefficient on Δx_t relative to π_t and by the fact that $\delta < 1$.

5 Unequal exposure versus idiosyncratic consumption risk

An important factor mediating the effect of monetary policy on consumption inequality is that different households—borrowers vs lenders, stockholders vs non-stockholders—are unequally exposed to aggregate shocks and policy. Our baseline model abstracts from such unequal exposures to focus on the role of

²⁹Appendix F.4 shows that i_t^* can be expressed as a function of all shocks and x_{t-1} . $\delta(\Omega) \frac{y+\rho}{1+\gamma\rho} \hat{z}_t$ is the level of output that the planner would implement if they had access to a set of time varying production subsidies which they would use to implement $\hat{p}_t = 0$ for all $t \geq 0$. With $\hat{p}_t = 0$, (27) implies that the planner would set $x_t = 0$ or $\hat{y}_t = \delta(\Omega) \frac{y+\rho}{1+\gamma\rho} \hat{z}_t$ at all dates.

idiosyncratic risk. We now introduce two sources of unequal exposures and study how they affect optimal policy: (i) initial wealth dispersion, which arises because a non-utilitarian planner optimally refrains from eliminating initial wealth inequality, and (ii) unequally distributed dividends.

5.1 The non-utilitarian planner and the URE channel

As described in Section 3, a utilitarian planner optimally uses the wealth tax to eliminate initial wealth inequality, implying that, trivially, monetary policy does not need to redistribute from savers to borrowers. As we show next, a non-utilitarian (NU) planner chooses not to eliminate pre-existing wealth inequality; thus, monetary policy *does* exploit the URE channel to redistribute consumption between savers and borrowers at date 0—not on average, but in response to aggregate shocks.

The NU planner maximizes the Pareto weighted sum of households' lifetime utilities, assigning different Pareto weights to households that have different observable characteristics at the beginning of date 0. In our model, the relevant individual state is household wealth, and so we allow the NU planner to assign Pareto weights $e^{\gamma\alpha a_0^s(i)}$ to households with wealth $a_0^s(i)$ at date 0.³⁰ $\alpha \geq 0$ indexes the planner's tolerance for pre-existing wealth inequality. When $\alpha = 0$, the planner is utilitarian and puts equal weights on all individuals who are alive at date 0. The larger α , the higher the relative weight on individuals with higher wealth at date 0. Given Pareto weights, the planner's period t felicity function is

$$\mathbb{U}_t = \underbrace{(1 - \vartheta) \sum_{s=-\infty}^0 \vartheta^{t-s} \int e^{\gamma\alpha a_0^s(i)} u(c_t^s(i), \ell_t^s(i); \xi_t^s(i)) di}_{\text{utility of individuals born before date 0}} + \underbrace{(1 - \vartheta) \sum_{s=1}^t \int \vartheta^{t-s} u(c_t^s(i), \ell_t^s(i); \xi_t^s(i)) di}_{\text{utility of individuals born after date 0}}$$

As in the baseline, \mathbb{U}_t can still be decomposed into the flow utility of a notional representative agent and the evolution of the welfare cost of consumption inequality Σ_t , which is now defined as

$$\Sigma_t = (1 - \vartheta) \sum_{s=-\infty}^0 \vartheta^{t-s} \int e^{\gamma\alpha a_0^s(i)} e^{-\gamma(c_t^s(i) - c_t)} di + (1 - \vartheta) \sum_{s=1}^t \int \vartheta^{t-s} e^{-\gamma(c_t^s(i) - c_t)} di$$

Unlike in the baseline, Σ_t is not unambiguously increasing in consumption inequality: the planner does not regard all consumption inequality arising from differences in pre-existing wealth inequality at date 0 as undesirable. However, they still regard as undesirable all inequality resulting from idiosyncratic shocks from date 0 onwards. This is reflected in the fact that while (21) is unchanged, (22) is now given by

$$\ln \Sigma_0 = \frac{1}{2} \gamma^2 \mu_0^2 w_0^2 \sigma_0^2 + \underbrace{\ln \left[\frac{1 - \vartheta}{1 - \vartheta e^{\frac{\Delta}{2} \left(\frac{\alpha - (1 - \tau_0^a) \mu_0}{\mu} \right)^2}} \right]}_{\text{welfare cost of date 0 wealth inequality}} \quad (30)$$

For $\alpha > 0$, completely eliminating wealth inequality (setting $\tau_0^a = 1$) no longer sets the second term on the RHS to zero. If the planner has access to a state contingent wealth tax (which can be changed in response to the shock), they would still set this term to zero, eliminating any undesirable wealth inequality, but

³⁰Since all households in the same cohort born at date $s > 0$ are ex ante identical, the planner assigns them the same Pareto weight. For the reasons described in footnote 14, this weight is β^s .

the level of wealth tax that accomplishes this is now $1 - \tau_0^a = \frac{\alpha}{\mu_0}$. Intuitively, whatever degree of date 0 redistribution from savers to borrowers is desired, the wealth tax can be used to deliver this, allowing monetary policy to focus on its other objectives: price stability, productive efficiency and consumption insurance. It follows that with a state contingent wealth tax, the optimal plan chosen by a planner with $\alpha > 0$ is the same as that chosen by the utilitarian planner. This is formalized in the proposition below.

Proposition 6 (State contingent τ_0^a). *If the planner has access to a state contingent τ_0^a , the dynamics of \hat{y}_t, π_t and $\hat{\Sigma}_t$ are the same for an NU planner ($\alpha > 0$) as for the utilitarian planner ($\alpha = 0$).*

Proof. See Appendix D.4. □

However, our maintained assumption is that fiscal policy *cannot* respond to shocks; the planner can only use the wealth tax to deliver the desired level of redistribution absent aggregate shocks. The wealth tax that accomplishes this is $1 - \tau_0^{a*} = \alpha/\mu$.³¹ Absent shocks, this tax sets the second term on the RHS of (30) to zero, reducing pre-existing wealth inequality to the planner's desired level. However, in response to shocks, the welfare cost of pre-existing inequality is not necessarily zero and is given by $\ln \left[\frac{1-\vartheta}{1-\vartheta e^{\frac{\lambda}{2} \left(\frac{\alpha}{\mu}\right)^2 \left(\frac{\mu-\mu_0}{\mu}\right)^2}} \right]$. Since some wealth inequality remains, a surprise change in interest rates still redistributes between savers and borrowers (URE channel), unlike in the $\alpha = 0$ case, where the wealth tax eliminates pre-existing wealth inequality. A surprise rate hike reduces output and wages and increases μ_0 above its steady state level. Recall that μ_0 is not just the pass-through from income to consumption risk but is also the MPC out of wealth. Since some pre-existing wealth inequality remains, a higher MPC out of wealth increases the consumption dispersion between borrowers and savers, relative to the planner's desired level, raising Σ_0 . Conversely a surprise rate cut reduces the MPC, lowering the consumption gap between savers and borrowers.³² Thus the effect of μ_t on Σ_t is different at date 0 than at subsequent dates. Consequently, monetary policy now takes into account the URE channel when responding to a shock. As in our baseline economy, optimal policy can be characterized in terms of a target criterion; the effect of the URE channel is reflected in the fact that the date 0 target criterion is different than at all other dates.

Proposition 7. *The NU planner sets nominal rates so that the following target criterion holds for $t = 0$:*

$$\left(1 - \delta_0(\Omega)\right)\hat{y}_0 + \delta_0(\Omega)\left(\hat{y}_0 - \frac{y + \rho}{1 + \gamma\rho}\hat{z}_0\right) + \frac{1 + \gamma\rho}{\Upsilon_0(\Omega)}\frac{\lambda}{\lambda - 1}\hat{p}_0 = 0 \quad (31)$$

and for $t > 0$, the target criterion is the same as (27) in Proposition 3:

$$\left(1 - \delta(\Omega)\right)\hat{y}_t + \delta(\Omega)\left(\hat{y}_t - \frac{y + \rho}{1 + \gamma\rho}\hat{z}_t\right) + \frac{1 + \gamma\rho}{\Upsilon(\Omega)}\frac{\lambda}{\lambda - 1}\hat{p}_t = 0, \quad (32)$$

with $\Upsilon_0(\Omega) > \Upsilon(\Omega)$ and $\delta_0(\Omega) < \delta(\Omega)$. Furthermore, $\Upsilon_0(\Omega)$ is increasing in α , while $\delta_0(\Omega)$ is decreasing in α for $\alpha > 0$.

Proof. See Appendix E.3. □

³¹Here we are also using the fact that $\mathbb{E}_{-1}\mu_0 = \mu$ since the economy is assumed to be in steady state prior to date 0 and the date 0 shock is unanticipated. See Appendix D.1 for a derivation.

³²Since the NU planner regards the consumption gap which would obtain absent shocks as optimal given the wealth tax, this fall is just as undesirable as an increase in the consumption gap. This is why Σ_0 is an increasing function of $(\mu_0 - \mu)^2$.

At dates $t > 0$, the target criterion is the same as in our baseline economy with the utilitarian planner; the NU planner’s preference for wealth redistribution and the extent of initial inequality do not affect the trade-off between price stability, productive efficiency and consumption insurance, as in Section 4. However, at $t = 0$, the NU planner puts more weight on stabilizing output, relative to either the price level or the output gap, than at subsequent dates: $\Upsilon_0(\Omega) > \Upsilon(\Omega)$ and $1 - \delta_0(\Omega) > 1 - \delta(\Omega)$. This difference is larger, the larger is the planner’s tolerance for pre-existing wealth inequality α (and hence the potential strength of the URE channel). At date 0, the NU planner has an additional motive to keep the MPC out of wealth μ_0 close to its steady state level, since doing so keeps consumption differences between borrowers and savers close to their desired level. Given the relationship between μ and y , this can be accomplished by keeping y_0 closer to its steady state level, even if this comes at the cost of higher inflation or productive inefficiency. Thus, while the effect of monetary policy on inequality via the URE channel is conceptually distinct from its effect via idiosyncratic risk, the implications for optimal policy are the same: monetary policy should put even more weight on stabilizing output relative to productive efficiency and price stability.

Productivity shocks Figure 3 shows the optimal response to a negative productivity shock. Blue lines depict the utilitarian baseline, black lines with circle markers depict the NU planner with $\alpha = 0.5\mu$ (who optimally sets $\tau_0^{a*} = 50\%$) and the magenta dotted curves depict the planner with $\alpha = \mu$ (who optimally sets $\tau_0^{a*} = 0\%$). Recall that the utilitarian planner already cushions the decline in output (blue line in panel a) relative to its natural level y_t^n . The NU planners implement even smaller declines in date 0 output in order to prevent the MPC out of wealth from rising sharply (panel d). This is accomplished by cutting rates aggressively (panel c) at date 0 and raising them in subsequent periods. Higher date 0 output comes at the cost of higher date 0 inflation (panel b); to mitigate this, the planner commits to lower output and inflation after date 0 compared to the utilitarian planner. Naturally, all the differences from the utilitarian benchmark are most pronounced for the NU planner with $\alpha = \mu$ who tolerates the most pre-existing inequality and sets $\tau_0^{a*} = 0$ (magenta dotted lines).³³

Incorporating the URE channel (with $\alpha = \mu$) makes the magnitudes of key responses in our economy broadly similar to those in the quantitative analysis of Bhandari et al. (2021), for whom the URE channel is a key factor affecting the optimal response to productivity shocks (see their Figure IV). Our RANK planner allows output to fall 0.9 percentage points on impact, while the HANK planner (with $\alpha = \mu$) only allows output to fall by a third as much; similarly in Bhandari et al. (2021), the RANK planner allows output to fall 0.5-1 percentage points, while their HANK planner only allows output to fall half as much. Similarly, while the RANK planners implement $\pi_t = 0$ in both our economy in Bhandari et al. (2021), both HANK planners allow inflation to increase on impact (around 0.05 percentage points in our economy and 0.25 percentage points in Bhandari et al. (2021)), followed by mild deflation. Finally, the HANK planners in both our paper and theirs cut nominal rates by about 1 percentage point, whereas the RANK planners raise rates.

Since our contribution is analytic and not quantitative in nature, our goal is not to precisely quantify the relative contribution of idiosyncratic risk and unequal exposures to the difference in optimal policy between HANK and RANK. Nonetheless, our results do suggest that idiosyncratic risk can account for a significant share of this difference, even under a calibration allowing for relatively large unequal exposures,

³³We do not plot Σ in this experiment because the definition of Σ is different for each planner as it depends on α . Thus, the paths of Σ in economies with different α are not comparable.

when income risk is countercyclical. We assess the relative contributions of consumption risk versus unequal exposure to the difference in optimal policy between HANK and RANK by computing the difference in the date 0 response of output and inflation across those two economies when only consumption risk generates inequality (as in Section 4), as a fraction of the same difference when both consumption risk and unequal exposure generate inequality (as in this section).³⁴ By this metric, even with a zero wealth tax (i.e., assuming $\alpha = \mu$), implying the largest possible role for unequal exposures, idiosyncratic risk accounts for around 30% of the difference in the response of output, and around 40% of the difference in the response of inflation, between HANK and RANK.³⁵ With a positive wealth tax ($\alpha < \mu$), the relative contribution of idiosyncratic risk would be greater.




Figure 3: **Optimal policy in response to a negative productivity shock with the non-utilitarian planner** Blue lines depict the utilitarian baseline, black lines with circle markers depict the NU planner with $\alpha = 0.5\mu$, magenta dotted lines depict the planner with $\alpha = \mu$ and the red dashed curves depict RANK. All panels plot log-deviations from steady state.

Markup shocks Figure 4 shows the optimal response to a positive markup shock. The utilitarian planner in the HANK economy already implements a smaller decline in output and a larger increase in inflation compared to optimal policy in RANK. These differences are amplified in the case with an NU planner, who seeks to mitigate the increase in MPC out of wealth μ_0 (panel d) and does so by cutting rates more aggressively at date 0 (panel c) and implementing higher output (black line with circle markers and magenta dotted line in panel a). As with the productivity shock, the planner seeks to mitigate the effect this has on date 0 inflation by committing to lower output and inflation in subsequent periods.

5.2 Unequal distribution of profits

In our baseline model, all households receive an equal share of firm profits. However, the literature has emphasized that the distribution of profits is a key determinant of positive (Broer et al., 2019) and normative (Bhandari et al., 2021) properties of HANK economies. We now relax the assumption of equally distributed dividends by assuming that a fraction $\eta^d < 1$ in each cohort s receives an equal share of dividends (“stockholders”), while the remaining fraction of $1 - \eta^d$ households receives no dividends (“non-stockholders”). Both groups supply labor and face the same distribution of idiosyncratic shocks $\xi_t^s(i)$.

³⁴In other words, idiosyncratic risk share (output) = $\frac{\hat{y}_0^{\text{HANK, no URE}} - \hat{y}_0^{\text{RANK}}}{\hat{y}_0^{\text{HANK, URE}} - \hat{y}_0^{\text{RANK}}}$, and similarly for inflation.

³⁵Graphically, this measures the difference between the blue and red dashed lines in Figure 3 as a fraction of the difference of the magenta dotted and the red dashed lines.




Figure 4: **Optimal policy in response to a positive markup shock with the non-utilitarian planner** Blue lines depict the utilitarian baseline, black lines with circle markers depict the NU planner with $\alpha = 0.5\mu$, magenta dotted lines depict the planner with $\alpha = \mu$ and the red dashed curves depict RANK. All panels plot log-deviations from steady state.

Appendix G presents the problem of a utilitarian planner in the economy with unequal distribution of dividends. In addition to the instruments available to the planner in the benchmark economy, we allow the planner to levy a lump-sum tax J on stockholders and make a lump-sum transfer $\frac{\eta^d}{1-\eta^d}J$ to each non-stockholder. This transfer is fixed at a level which equalizes the average consumption of the two groups in steady state, $J = \frac{1-\eta^d}{\eta^d}D$, where D denotes steady state dividends. However, the transfer cannot be adjusted to keep the average consumption of the two groups equal in response to aggregate shocks.

The CARA-normal structure of our economy still implies that households' consumption is an affine function of cash-on-hand. However, the time varying intercept of the consumption function is different for the two groups. The date t consumption of a stockholder i who was born at date $s \leq t$ is $c_t^s(i; d) = C_t^d + \mu_t x_t^s(i; d)$, while that of a non-stockholder is $c_t^s(i; nd) = C_t^{nd} + \mu_t x_t^s(i; nd)$, where

$$C_t^d = y_t + \left(\frac{1-\eta}{\eta}\right) \mu_t \mathcal{V}_t \quad , \quad C_t^{nd} = y_t - \mu_t \mathcal{V}_t \quad (33)$$

and \mathcal{V}_t denotes the present discounted value of dividends relative to their steady state value D :

$$\mathcal{V}_t = (D_t - D) + \vartheta R_t^{-1} \mathcal{V}_{t+1} \quad (34)$$

The consumption of the two groups is equalized in steady state since $\mathcal{V} = 0$. Shocks that raise current or future dividends tend to increase the consumption of stockholders for a given aggregate income and reduce the consumption of non-stockholders. Linearizing (34),

$$\widehat{\mathcal{V}}_t = \mathcal{D}_y \widehat{y}_t + \frac{y(y+\rho)}{\rho\lambda} \widehat{z}_t + \widetilde{\beta} \widehat{\mathcal{V}}_{t+1}, \quad \text{where} \quad \mathcal{D}_y = \frac{1}{\lambda} \left[\lambda - 1 - \left(\frac{1+\gamma\rho}{\rho} \right) y \right]$$

\mathcal{D}_y denotes the effect of higher output on profits, holding productivity constant. The sign of \mathcal{D}_y is theoretically ambiguous: with sticky prices, higher output, without an increase in productivity, raises revenues but also increases marginal costs. Which force dominates depends on the elasticity of labor supply, which determines how responsive wages are to an increase in hours worked, and on the steady state markup λ .³⁶

The unequal distribution of profits introduces between-group inequality—difference between the average

³⁶Notice that assuming productivity is at its steady state level and with zero inflation, the aggregate labor supply can be written as $n_t = \rho \ln w_t - \gamma \rho \frac{n_t}{\xi}$. Consequently, $\frac{d \ln n_t}{d \ln w_t} = \frac{\rho}{1+\gamma\rho} \frac{1}{y}$ denotes the aggregate elasticity of labor supply.

consumption of stockholders and non-stockholders—in addition to the within-group inequality arising from uninsurable idiosyncratic risk. Thus, shocks and policy now affect the welfare relevant measure of inequality Σ_t in two ways. First, as before, innovations to within-group consumption risk, $0.5\gamma^2\mu_t^2w_t^2\sigma_t^2$ increase inequality. Secondly, between-group consumption inequality increases Σ_t for a given level of risk:³⁷

$$\ln \Sigma_t = \frac{\gamma^2\mu_t^2w_t^2\sigma_t^2}{2} + \ln [(1 - \vartheta) \mathbb{B}_t + \vartheta\Sigma_{t-1}], \quad (35)$$

where $\mathbb{B}_t = \eta e^{-\gamma(c_t^d - y_t)} + (1 - \eta) e^{-\gamma(c_t^{nd} - y_t)}$ captures between-group differences in average consumption. Again, the implications of this source of between-group inequality can be seen by inspecting the target criterion characterizing optimal monetary policy:

Proposition 8. *Optimal policy sets nominal rates so that the following target criterion holds for $t = 0$:*

$$\Upsilon(\Omega) x_0 + (1 + \gamma\rho) \frac{\lambda}{\lambda - 1} \hat{p}_0 + \mathbb{K}(\eta^d) \mathcal{D}_y \hat{\mathcal{V}}_0 = 0 \quad (36)$$

and for $t > 0$:

$$\Upsilon(\Omega) \left(x_t - \beta^{-1} \tilde{\beta} x_{t-1} \right) + (1 + \gamma\rho) \frac{\lambda}{\lambda - 1} \left(\hat{p}_t - \beta^{-1} \tilde{\beta} \hat{p}_{t-1} \right) + \mathbb{K}(\eta^d) \mathcal{D}_y \left(1 - \beta^{-1} \tilde{\beta} \right) \hat{\mathcal{V}}_t = 0, \quad (37)$$

where $x_t = (1 - \delta(\Omega)) \hat{y}_t + \delta(\Omega) \left(\hat{y}_t - \frac{y+\rho}{1+\gamma\rho} \hat{z}_t \right)$ and we have used the fact that $x_{-1} = \pi_{-1} = 0$. \mathcal{D}_y is the effect of higher output on dividends and $\mathbb{K}(\eta^d) \geq 0$ with $\mathbb{K}(1) = 0$ and $\mathbb{K}'(\eta^d) < 0$, i.e., more-concentrated wealth (lower η^d) increases \mathbb{K} . $\Upsilon(\Omega)$ is defined as in Proposition 3.

Proof. See Appendix G. □

When $\eta^d = 1$, the target criterion (36)-(37) is identical to the one described in Proposition 3. When $\eta^d < 1$, monetary policy tries to stabilize the present discounted value of dividends $\hat{\mathcal{V}}_t$, in addition to output, the output gap and the price level. This is because fluctuations in $\hat{\mathcal{V}}_t$ generate between-group consumption inequality: higher $\hat{\mathcal{V}}_t$ widens the average consumption gap between stockholders and non-stockholders.

The weight on stabilizing dividends depends on the concentration of stockholdings. The higher the concentration, the greater the difference in per-capita consumption induced by a given change in dividends. The sign of the coefficient on stabilizing dividends depends on the effect of higher output on dividends \mathcal{D}_y . If $\mathcal{D}_y < 0$, the planner prefers higher output (even compared to the baseline) in response to a shock that raises $\hat{\mathcal{V}}$. A higher $\hat{\mathcal{V}}$ increases the relative consumption of stockholders. Raising output in response to the shock tends to reduce dividends, mitigating the rise in \mathcal{V}_t and the average consumption gap. If instead $\mathcal{D}_y > 0$, the planner prefers lower output when \mathcal{V} is higher, because now *lower* output reduces dividends. In either case, between-group inequality provides an additional motive to avoid large fluctuations in output as these tend to benefit one group relative to another. As discussed in Section 5.1, while between-group inequality is conceptually different from inequality that arises due to idiosyncratic risk, both these sources of inequality increase the weight that optimal monetary policy puts on output stabilization. (Again, in RANK, both sources of inequality are absent and optimal policy puts no weight on output stabilization, focusing exclusively on productive efficiency and price stability.)

³⁷See Appendix G.1 for the derivation of the Σ_t recursion in this case.

This motive for stabilizing dividends is particularly strong at date 0 since a change in $\widehat{\mathcal{V}}_0$ generates consumption gaps between all stockholders and non-stockholders who are alive at date 0. In contrast, a change in $\widehat{\mathcal{V}}_t$ for $t > 0$ only generates consumption gaps among agents born at date t ; the effect of higher dividends at $t > 0$ on the consumption of stockholders alive at date $s < t$ is already captured in $\widehat{\mathcal{V}}_s$ since stockholders are forward-looking and can borrow at date s against higher date t dividend income. Thus, the target criterion is not time invariant: the planner puts more weight on stabilizing \mathcal{V}_t at $t = 0$ than at all subsequent dates.

While the target criterion characterizes the optimal response to all shocks, we focus on markup shocks to save space and to relate our results to those of [Bhandari et al. \(2021\)](#), for whom the unequal distribution of dividends is a particularly important determinant of the optimal response to markup shocks. [Figure 5](#) shows the optimal response to a positive markup shock in our baseline calibration with $\mathcal{D}_y < 0$. This shock increases the present value of dividends (panel d), driving the average consumption of stockholders $\widehat{\mathcal{C}}_t^d$ above that of non-stockholders $\widehat{\mathcal{C}}_t^d$ (panel e). This effect is more severe the more concentrated stockholdings are: the magenta dotted curve shows a case with more concentration, $\eta^d = 0.1$, the black line with circle markers depicts less concentration, $\eta^d = 0.5$, and the blue curve denotes the baseline with an equal distribution of dividends. To control the rise in between-group inequality, the planner implements higher output relative to the baseline with $\eta^d = 1$ (panel a), raising wages while curtailing the increase in dividends. This difference relative to the baseline is largest at date 0 when stabilizing \mathcal{V}_0 has the largest impact on between-group inequality—in fact when $\eta^d = 0.1$ (magenta dotted line), policy *increases* output by around 0.2 percentage points in response to a positive markup shock. This is in line with the numerical results of [Bhandari et al. \(2021\)](#), whose HANK planner raises output by 0-0.1 percentage points.³⁸ Similarly, both HANK planners allow inflation to increase on impact (by 0.1 percentage points in our economy, and 0.2 percentage points in [Bhandari et al. \(2021\)](#)), followed by mild deflation—in stark contrast to the RANK plan, which features a fall in output and a smaller increase in inflation on impact.

Importantly, the motive to stabilize \mathcal{V} would be present even in an economy with no idiosyncratic risk ($\sigma_t = 0$) as long as $\eta^d < 1$. This would imply $\Upsilon = 1$ but $\mathbb{K}(\eta^d) \neq 0$. Even absent idiosyncratic risk, unequal dividends would leave households imperfectly insured against aggregate risk, allowing monetary policy to improve welfare by substituting for these missing markets as in [Bhandari et al. \(2021\)](#). This highlights that reducing idiosyncratic risk and providing insurance against “aggregate risk” as in [Bhandari et al. \(2021\)](#) are two *distinct* motives that cause optimal policy in HANK to differ from that in RANK. As in [Section 5.1](#), we can decompose the difference between HANK and RANK into the contributions of idiosyncratic risk and unequal exposures. Even with very unequally distributed dividends ($\eta^d = 0.1$), idiosyncratic risk still accounts for 17% of the difference in the date 0 response of output and 13% of the difference in inflation.

³⁸Unequally distributed dividends do not give the planner an incentive to use monetary policy to redistribute from stockholders to non-stockholders absent shocks, because the lump-sum tax available to the planner does exactly this. [Bhandari et al. \(2021\)](#) take a similar approach: they introduce an unequal distribution of dividends, calibrate the tax rate on dividends in line with US data, and calibrate Pareto weights so that absent shocks, this dividend tax is optimal. A utilitarian planner who cannot use the lump-sum tax would also use monetary policy to redistribute from stockholders to non-stockholders even absent shocks.




Figure 5: **Optimal policy in response to markup shocks** The red dashed lines depict dynamics in RANK. All other curves depict dynamics in HANK with $\Omega > 0$ in response to an increase in firms' desired markups. The blue curves depict the case with equal distribution of profits, the black lines with circle markers depict the case in which 50% of households get dividends and the dotted-magenta curve depicts the case in which only 10% of households get dividends. All panels plot log-deviations from steady state except panels d and e which plot $100 * \widehat{V}_t/y$ and $100 * (\widehat{C}_t^d - \widehat{C}_t^{nd})/y$, respectively.

5.3 The role of the cyclicity of income risk in the quantitative literature

Our analytical framework complements the quantitative literature on optimal policy in HANK, in particular, [Bhandari et al. \(2021\)](#) and [Le Grand et al. \(2021\)](#). The main differences between their findings and ours arise because their models feature *procyclical* income risk, whereas our baseline has countercyclical income risk. In these papers, household income equals labor income, which bears idiosyncratic risk, plus non-labor income (dividends net of transfers), which has no idiosyncratic risk. Households' return to labor is the product of the real wage, which is *procyclical* and common across households, and their idiosyncratic skill whose variance is exogenous. Thus, higher wages increase the variance of labor income (in levels; the variance of log labor income is acyclical) and of total income (both logs and levels). We instead incorporate countercyclical risk ([Storesletten et al., 2004](#)) by allowing the variance of idiosyncratic shocks to vary with output.

Given this difference, the main factor that causes optimal policy in our economy to differ from RANK—the fact that more-accommodative policy reduces idiosyncratic consumption risk—is absent in [Bhandari et al. \(2021\)](#) and [Le Grand et al. \(2021\)](#). To the extent that they find differences between optimal policy in HANK and RANK, these are due to other channels such as the URE channel or unequally distributed dividends, as discussed above. In [Bhandari et al. \(2021\)](#), nearly all the differences from RANK stem from the planner's desire to provide insurance against *aggregate* shocks. Thus, the planner cuts interest rates in response to negative productivity shocks to transfer resources from savers to borrowers (cf. Section 5.1), and implements higher output following a positive markup shock to transfer from stockholders to non-stockholders (cf. Section 5.2). In contrast, when [Bhandari et al. \(2021\)](#) shut off idiosyncratic risk, optimal policy is essentially unchanged, indicating that facilitating insurance against idiosyncratic risk is not an important goal of optimal monetary policy in their economy.

Our analysis suggests that the reason idiosyncratic risk per se does not lead to differences from RANK in these papers is that they feature procyclical rather than countercyclical income risk. Indeed, in our baseline with $\Omega = 0$ (mildly procyclical income risk), optimal policy would be identical in HANK and

RANK (cf. Lemma 1).³⁹ But with acyclical or countercyclical income risk, optimal policy differs from RANK even when there is no initial wealth inequality, no URE channel and equally distributed dividends, because accommodative monetary policy can reduce idiosyncratic consumption risk.

6 Some extensions

Our versatile framework can be extended in many directions to study how additional channels and shocks affect optimal monetary policy in HANK. To illustrate this, we now extend our analysis to include (i) MPC heterogeneity and (ii) *demand* shocks (i.e., shocks that do not affect the flexible-price level of output).

6.1 Hand-to-mouth agents

Our baseline model deliberately abstracts from MPC heterogeneity and shows that even absent such heterogeneity, optimal policy sharply differs from RANK. We now study how MPC heterogeneity, a feature of quantitative HANK models that has received much attention since Kaplan et al. (2018), affects optimal monetary policy. We do so by introducing a fraction η^h of hand-to-mouth (HtM) households that cannot trade bonds and consume their after-tax income. These households are otherwise identical to the remaining $1 - \eta^h$ unconstrained households that trade bonds as in the baseline—in particular, both groups draw idiosyncratic shocks from the same distribution and receive the same dividends and transfers per capita.

While the MPC of unconstrained households μ_t is still described by (12), the MPC of constrained households is $\tilde{\mu}_t = (1 + \gamma\rho w_t)^{-1}$. These households can still self-insure to some extent by adjusting hours worked, implying that $\tilde{\mu}_t < 1$. However, since they cannot insure using the bond market, their MPC is higher than those of the unconstrained households; i.e., $\tilde{\mu}_t > \mu_t$.

Appendix H shows that the presence of HtM households does not change the dynamics of aggregate variables, conditional on a given a path of interest rates. These dynamics are still given by (17)-(20). In equilibrium, since HtM households consume their income, and aggregate consumption equals aggregate income, the average consumption of unconstrained households must equal aggregate income as in our baseline.⁴⁰ However, introducing HtM households does affect social welfare and, therefore, optimal policy. While the period t felicity function of the utilitarian planner can still be written as $\mathbb{U}_t = u(c_t, n_t; \bar{\xi}) \times \Sigma_t$, the welfare relevant measure of consumption inequality is now $\Sigma_t = (1 - \eta^h)\Sigma_t^{nh} + \eta^h\Sigma_t^h$, where Σ_t^{nh} denotes consumption inequality among unconstrained households and evolves according to (21), while Σ_t^h denotes consumption inequality among HtM households, and equals $\Sigma_t^h = \frac{1}{2}\gamma^2\tilde{\mu}_t^2 w_t^2 \sigma_t^2$. Since there is no wealth inequality among HtM households, unlike Σ_t^{nh} , Σ_t^h depends only on *current* consumption risk. However, since $\tilde{\mu}_t > \mu_t$, consumption inequality moves more for this group in response to changes in income risk. While the trade-offs facing the planner are qualitatively the same as in our baseline economy, quantitatively, monetary policy has even larger effects on Σ_t in the presence of HtM households:

Lemma 2. *The effect of a one-time increase in output engineered by monetary policy reduces inequality*

³⁹With $\Omega = 0$, introducing unequal distribution of dividends or initial wealth inequality due to a non-utilitarian planner would still introduce differences between HANK and RANK.

⁴⁰This is for the same reasons as in Bilbie (2008); Werning (2015); Acharya and Dogra (2020).

Σ_t by a larger amount, the larger the fraction of HtM households η^h :

$$\frac{\partial^2 \widehat{\Sigma}_t}{\partial \widehat{y}_t \partial \eta^h} = -\gamma \Omega \left\{ \left[1 - \widetilde{\beta}(1 - \Lambda) \right] \left(\Sigma^h (1 - \widetilde{\beta})^{-2} - \Sigma^{nh} \right) + \widetilde{\beta} \Lambda \Sigma^{nh} \right\} < 0 \quad \text{for} \quad \Omega > 0$$

Proof. See Appendix H.3. □

Since the main differences in optimal policy in HANK relative to RANK arise because monetary policy can affect inequality, a higher sensitivity of inequality to monetary policy magnifies these differences.

Productivity shocks Figure 6 shows the dynamics under optimal policy following a negative productivity shock in RANK (red dashed curves), HANK with no HtMs (blue solid curves) and HANK with 30% HtMs (magenta dot dashed curves).⁴¹ In our baseline ($\eta^h = 0$), monetary policy already prevents output from falling as much as \widehat{y}_t^n on impact, permitting some inflation. With $\eta^h > 0$, policy cushions the fall in output even more (see panel a), resulting in even higher inflation responses initially (see panel b). Quantitatively, the impact response of the output gap is about twice as large with HtM households, and that of inflation about two and a half times as large. Intuitively, a fall in output is more costly with $\eta^h > 0$ because it increases consumption inequality more for HtMs that cannot self-insure using the bond market. This can be seen by comparing the magenta dot-dashed curves in panel c, which plots consumption inequality among unconstrained households, with panel d which plots inequality among the HtMs. At its peak, the percentage increase in Σ_t^h is around ten times the increase in Σ_t^{nh} . Thus, the benefit of mitigating the fall in output, in terms of the effect on Σ_t , is much higher in the economy with HtMs. To see this, compare the magenta dot-dashed curve in panel e, which plots inequality under optimal policy in the economy with 30% HtMs, to the black dotted curve, which plots inequality if monetary policy uses the target criterion which would be optimal in an economy with no HtMs. The difference between these curves—the reduction in overall inequality due to a higher path of output—is much larger than the reduction in inequality among the unconstrained households, shown by the difference between the curves in panel c. Since inequality is more sensitive to the level of output in the presence of HtMs, the planner tolerates larger deviations from productive efficiency and price stability to mitigate the rise in inequality following an adverse shock.

Markup shocks Similarly, when studying markup shocks in our HANK economy with HtMs, the difference between optimal policy in HANK and RANK is qualitatively the same as in our baseline, but quantitatively amplified. To mitigate the increase in inequality, particularly among HtMs, monetary policy stabilizes output more (magenta dot-dashed curve relative to blue solid curve in panel a, Figure 7) at the cost of higher inflation (magenta dot-dashed curve relative to blue solid curve in panel b). Quantitatively, in the presence of HtMs, optimal policy shaves off around half the initial fall in output in RANK, while optimal policy only shaves off about a quarter in our baseline (absent HtMs). Similarly, the increase in inflation is about 50% larger with HtMs.

Overall, introducing MPC heterogeneity does not qualitatively change the trade-offs analyzed in our baseline. In fact, it accentuates the differences relative to RANK: with higher MPCs; i.e., higher pass-through from income to consumption risk, consumption inequality is even more sensitive to monetary

⁴¹ $\eta^h = 0.3$ is in line with Kaplan et al. (2014) who find that approximately 30% of US households are hand-to-mouth. Given our calibration, this implies an average MPC of around 17% (around 40% for HtMs and 7% for unconstrained households), which is in line with the range of MPCs reported in the empirical literature.




Figure 6: **Optimal policy in response to productivity shocks** In panels a and b, blue solid curves depicts dynamics in HANK with $\Omega > 0$ and no HtM agents; red dashed curves depict dynamics in RANK; and magenta dot-dashed lines depict the optimal response of an economy with 30% HtM households following a negative productivity shock. In panels c, d and e, the magenta dot-dashed lines present the evolution of Σ_t^{nh} , Σ_t^h and Σ_t , respectively, under optimal policy in the economy with 30% HtMs, while the black dotted line depicts the evolution of these variables in the economy with 30% HtMs if monetary policy implements the target criterion which would be optimal in an economy with no HtMs. All panels plot log-deviations from steady state except panel a which plots $(\hat{y}_t - \hat{y}_t^n)/y$.




Figure 7: **Optimal policy in response to markup shocks** In panels a and b, blue solid curves depicts dynamics in HANK with $\Omega > 0$ and no HtM agents; red dashed curves depict dynamics in RANK; and dot-dashed magenta lines depict the optimal response of an economy with 30% HtM households following a positive markup shock. In panels c,d and e, the dot-dashed magenta line presents the evolution of Σ_t^{nh} , Σ_t^h and Σ_t resp. under optimal policy in the economy with 30% HtMs, while the black dotted line depicts the evolution of these variables in the economy with 30% HtMs if monetary policy implements the target criterion which would be optimal in an economy with no HtMs. All panels plot log-deviations from steady state.

policy. Consequently, policy deviates even further from RANK to stabilize inequality. This suggests that the trade-offs we study analytically would be even more important in quantitative HANK economies with a substantial fraction of high MPC households.

6.2 Optimal monetary policy response to demand shocks

In Section 4, we focused on productivity and markup shocks, both of which affect the natural level of output y_t^n . The RANK literature also studies the optimal response to other shocks that do not affect y_t^n ; e.g. changes in households' discount factor. Following the literature, we term these *demand* shocks. Since

these shocks do not induce a trade-off between productive efficiency and price stability, the RANK planner simply implements $\widehat{y}_t = \widehat{y}_t^n = \pi_t = 0$ in response to these shocks by setting the interest rate equal to the *natural rate of interest* r_t^* ; i.e., the interest rate consistent with $y_t = y_t^n$ at all dates.

As shown in Section 4, the HANK planner generally does not implement $y_t = y_t^n$, even in response to productivity shocks which do not induce a trade-off between productive efficiency and price stability. This is because responding one-for-one to fluctuations in the natural level of output would adversely affect inequality. Similarly, in response to demand shocks, setting $y_t = y_t^n$ is in general not optimal because these shocks would affect inequality should monetary policy fully insulate output from them. Consequently, optimal policy lets output vary in order to offset these undesirable changes in inequality.

We study two demand shocks: (i) changes in households' discount factor and (ii) shocks to the variance of idiosyncratic shocks faced by households. We now assume that household preferences are given by

$$\mathbb{E}_s \sum_{t=s}^{\infty} (\beta\vartheta)^{t-s} \left(\prod_{k=s}^{t-1} \zeta_k \right) u \left(c_t^s(i), \ell_t^s(i); \xi_t^s(i) \right),$$

where ζ_t is a shock to the individual's discount factor between dates t and $t+1$. Appendix A shows that Proposition 1 remains true except that the aggregate Euler equation (11) becomes

$$C_t = -\frac{1}{\gamma} \ln \beta \zeta_t R_t + C_{t+1} - \frac{\gamma \mu_{t+1}^2 w_{t+1}^2 \sigma_{t+1}^2}{2}$$

The preference shock is internalized by the utilitarian planner who puts weight $\beta^s \left(\prod_{k=0}^{s-1} \zeta_k \right)$ on the lifetime utility of a household born at date $s > 0$.

We also introduce a shock to the variance of idiosyncratic risk faced by households (ξ) by assuming that this variance satisfies $\sigma_t^2 w_t^2 = \sigma^2 w^2 \exp \{2[\phi(y_t - y) + \varsigma_t]\}$. Higher ς_t increases the cross-sectional variance of cash-on-hand at date t . To the extent that the shock is persistent ($\varrho_\varsigma > 0$), this can also be thought of as a *risk shock*: higher ς_{t+1} increases the uncertainty households face at date t about the realization of the shock to disutility (and hence to cash-on-hand) at date $t+1$. When plotting IRFs, following Bayer et al. (2020), we set the persistence and standard deviation of risk shocks and discount-factor shocks to $\varrho_\varsigma = 0.68^4$, $\varrho_\zeta = 0.83^4$, $\sigma_\varsigma = 1.4$ and $\sigma_\zeta = 0.01$.

Both discount-factor shocks and risk shocks affect the evolution of consumption inequality. This can be seen through the linearized $\widehat{\Sigma}_t$ recursion (25) which now becomes⁴²

$$\frac{\widehat{\Sigma}_t}{\Sigma} = -\gamma \left(1 - \widetilde{\beta}\right) \Omega \widehat{y}_t - \frac{\widetilde{\beta} \Lambda}{1 - \widetilde{\beta} \varrho_\zeta (1 - \Lambda)} \widehat{\zeta}_t + \frac{(1 - \widetilde{\beta} \varrho_\varsigma) \Lambda}{1 - \widetilde{\beta} (1 - \Lambda) \varrho_\varsigma} \widehat{\varsigma}_t + \left(\frac{\widetilde{\beta}}{\beta} \right) \frac{\widehat{\Sigma}_{t-1}}{\Sigma} \quad (38)$$

An increase in $\widehat{\zeta}_t$ directly affects income risk and thus persistently affects consumption inequality. More subtly, a fall in households' discount factor $\widehat{\zeta}_t < 0$ increases the natural rate of interest, which in our economy is given by $r_t^* = -\frac{1 - \widetilde{\beta} \varrho_\zeta}{1 - \widetilde{\beta} (1 - \Lambda) \varrho_\zeta} \widehat{\zeta}_t - \frac{(1 - \widetilde{\beta} \varrho_\varsigma) \Lambda}{1 - \widetilde{\beta} (1 - \Lambda) \varrho_\varsigma} \widehat{\varsigma}_{t+1}$. Thus, if monetary policy keeps output unchanged in response to a fall in ζ_t , this entails a rise in interest rates, which increases the pass-through μ_t . For a given level of income risk, a higher pass-through increases consumption risk and hence the level of consumption inequality. A persistent increase in ς_t also reduces r_t^* as households attempt to increase their precautionary

⁴²See Appendix E.1 for a derivation. We have implicitly set $\widehat{z}_t = 0$ throughout this section.

savings in response to the increase in risk. This decline in interest rates reduces μ_t somewhat, offsetting some of the direct effect of a higher ς_t on consumption risk. However, a higher ς_t still increases Σ_t on net.

Since demand shocks affect inequality, the planner generally deviates from keeping output equal to its natural level and implementing zero inflation (even though this remains *feasible*) in order to mitigate the impact on inequality. This is formalized in the following proposition.⁴³

Proposition 9. *In response to demand shocks, the planner sets nominal interest rates so that the following target criterion holds at all dates $t \geq 0$:*

$$(\widehat{y}_t - y_t^*) + \frac{1 + \gamma\rho}{\Upsilon(\Omega)} \frac{\lambda}{\lambda - 1} \widehat{p}_t = 0 \quad (39)$$

where $y_t^* = -\chi(\Omega)\widehat{\zeta}_t + \Xi(\Omega)\widehat{\varsigma}_t$ is the desired level of output (in deviations from steady state). $\chi(\Omega)$ and $\Xi(\Omega)$ are defined in Appendix E.3 and satisfy $\chi(0) = \Xi(0) = 0$. $\Upsilon(\Omega)$ is the same as in Proposition 3. When risk is countercyclical ($\Theta > 1 \Rightarrow \Omega > \Omega^c$), $\chi(\Omega) > 0$ and $\Xi(\Omega) < 0$.

As described earlier, the target criterion (39) indicates that the planner seeks to minimize fluctuations in the price level while also keeping output close to its desired level y_t^* . When risk is acyclical or countercyclical, demand shocks that tend to increase consumption inequality—higher ς_t or lower ζ_t —increase y_t^* . That is, the planner targets a higher level of output because this tends to reduce consumption inequality when $\Omega \geq \Omega^c > 0$, mitigating the increase in inequality due to the shock. Since demand shocks keep y_t^n unchanged, adjusting output in response to these shocks entails some inflation; as discussed earlier, the HANK planner puts a smaller relative weight on price stability $\Upsilon(\Omega) > 1$ relative to the RANK planner.

Risk shocks We start by describing the dynamics under optimal policy in response to a risk shock $\widehat{\varsigma}_0 > 0$.

Proposition 10. *Under optimal policy with acyclical or countercyclical income risk, following an increase in risk ($\widehat{\varsigma}_0 > 0$), \widehat{y}_0 and π_0 both increase. In addition, there exists $T > 0$ such that for all $t \in (T, \infty)$, $\pi_t < 0$ and $\widehat{y}_t < 0$. Following a decline in risk ($\widehat{\varsigma}_0 < 0$) all these signs are reversed.*

Figure 8 plots the optimal response to an increase in risk in RANK and HANK (with $\Omega \geq \Omega^c$). In RANK, since households can trade Arrow securities, an increase in the cross-sectional dispersion of income does not result in any increase in consumption inequality. Since risk shocks do not affect y_t^n , the RANK planner keeps output fixed at $\widehat{y}_t = \widehat{y}_t^n = 0$, implying zero inflation $\pi_t = 0$ (red dashed lines).

In contrast, in HANK with $\Omega \geq \Omega^c$, monetary policy cuts nominal interest rates on impact (panel e) to raise output above its natural level $\widehat{y}_0 > \widehat{y}_0^n = 0$ in response to a positive risk shock (panel a). In the acyclical or countercyclical case ($\Omega \geq \Omega^c > 0$), higher output tends to reduce consumption inequality, partially offsetting the effect of the risk shock (see equation (25)). Lower interest rates and higher output (which implies higher wages) also make it easier for households to self-insure, lowering the pass-through from income to consumption risk; i.e., $\widehat{\mu}_0 < 0$ (panel f). Monetary policy trades off the benefit from mitigating the increase in inequality against the cost of higher inflation (panel b) and productive inefficiency ($\widehat{y}_t \neq \widehat{y}_t^n$). To mitigate this inflation, the planner commits to mildly lower output and inflation in the future. If instead,

⁴³The target criterion in Proposition 9 is a generalization of the target criterion in (27) to include demand shocks but abstracting from productivity shocks ($\widehat{z}_t = 0$). Appendix E.3 derives a general target criterion (equation (E.31)), which is valid in the presence of all four shocks that we study.

monetary policy implements $\hat{y}_t = \hat{y}_t^n = 0$ and $\pi_t = 0$ (which is optimal under RANK), this would result in higher inequality (black dotted curve in panel c).

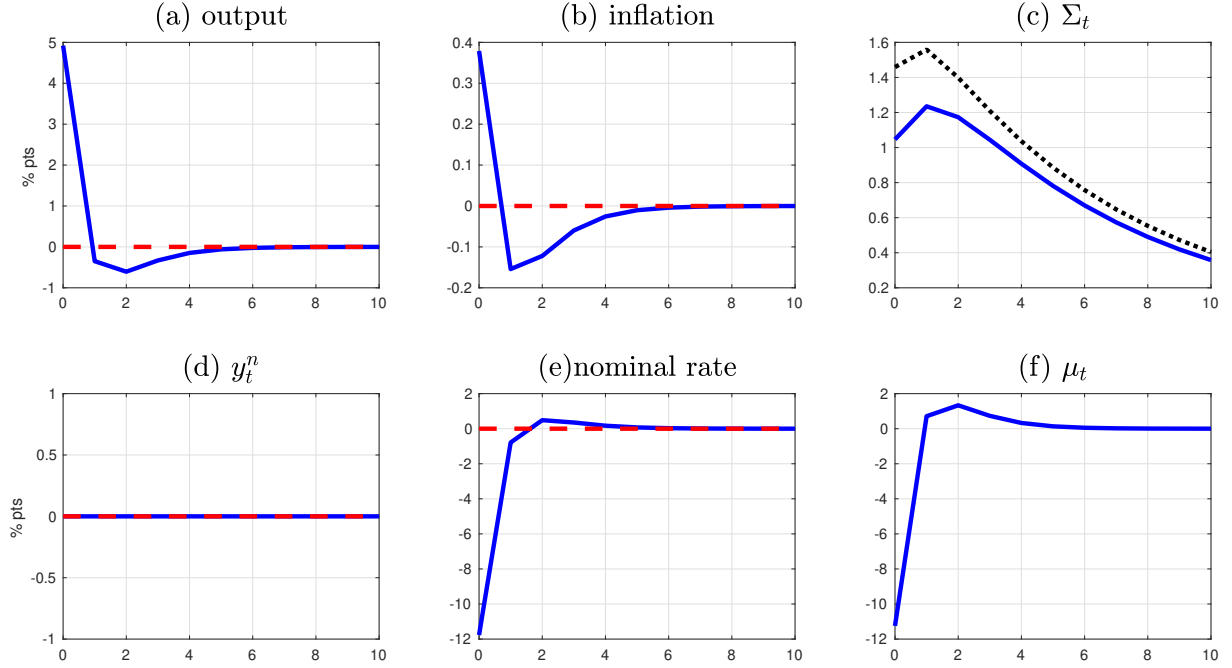


Figure 8: **Optimal policy in response to risk shocks** in HANK with $\Omega > 0$ (blue solid curves) and RANK (red dashed curves). Black dotted lines denote outcomes in HANK under non-optimal policy which sets $\hat{y}_t = \hat{y}_t^n = 0$, $\pi_t = 0 \forall t \geq 0$. All panels plot log-deviations from steady state.

Discount-factor shock A decrease in households' discount factor ($\hat{\zeta}_t < 0$) increases r_t^* , the interest rate consistent with $\hat{y}_t = \hat{y}_t^n = 0$ and $\pi_t = 0$. Consequently, the RANK planner raises interest rates one-for-one with r_t^* , keeping inflation and output unchanged. However, in HANK, this rise in interest rates would increase the pass-through μ_t and hence consumption inequality. Thus, as with a positive risk shock, monetary policy deviates from the flexible-price allocation ($\hat{y}_t = \hat{y}_t^n = \pi_t = 0$) to mitigate this rise in inequality.

Proposition 11. *Under optimal policy with acyclical or countercyclical income risk, following an decrease in households' discount factor ($\hat{\zeta}_0 < 0$), \hat{y}_0 and π_0 both increase. In addition, $\exists T > 0$ such that for all $t \in (T, \infty)$, $\pi_t < 0$ and $\hat{y}_t < 0$. Following a rise in households' discount factors, all these signs are reversed.*

Figure 9 plots the optimal dynamics following a negative discount-factor shock. As in RANK, the HANK planner raises rates (panel e), increasing the pass-through μ_t (panel f). This in turn tends to increase consumption inequality (panel c). However, the HANK planner does not increase rates one-for-one with r_t^* (panel d) as this would result in a larger increase in inequality (black dotted line in panel c). This lower path of interest rates increases output on impact (panel a), reducing the level of risk faced by households (when risk is countercyclical) and further curtailing the increase in inequality. To mitigate the rise in date 0 inflation, the planner commits to lower output and inflation in the future (panel b). However, these differences relative to RANK are fairly small given our calibration.




Figure 9: **Optimal policy in response to discount-factor shock** in HANK with $\Omega > 0$ (blue solid curves) and RANK (red dashed curves). Black dotted lines denote outcomes in HANK under non-optimal policy which sets $\widehat{y}_t - \widehat{y}_t^n = \pi_t = 0 \forall t \geq 0$. All panels plot log-deviations from steady state.

Absence of self-insurance channel in zero-liquidity HANK models The optimal response to discount-factor shocks highlights an important difference between our economy with $\Omega \geq \Omega^c > 0$ and zero-liquidity HANK economies (in which households cannot borrow and government debt is in zero net supply). In zero-liquidity models, interest rates do not affect households' ability to self-insure via the bond market, since they always consume their income in equilibrium. Thus, as in RANK, interest rates perform a single task in these economies: implementing the planner's desired path of output growth, which in turn affects inflation via the Phillips curve. Consequently, the planner can first choose output and inflation to maximize welfare subject to the Phillips curve, ignoring the IS curve. After this, the planner can use the IS equation to back out the interest rates implementing the desired path of output and inflation. Since discount-factor shocks only affect the IS curve, which can be dropped as a constraint, the planner in a zero-liquidity or RANK economy leaves output and inflation unchanged following such a shock, raising interest rates one-for-one with r_t^* .

In our HANK economy, the IS curve cannot be dropped as a constraint since the interest rate performs two tasks: (i) it affects output via the IS curve (15) and (ii) it affects the pass-through from income to consumption risk μ_t through (12). Formally, Appendix D.2 shows that the multiplier on the IS equation is non-zero in our HANK model but zero in RANK; it would also be 0 in a zero-liquidity HANK model. Our planner, therefore, faces a trade-off that is absent in both RANK and zero-liquidity economies: when choosing what path of output to target, they must also consider how the interest rates which implement the desired path of output affect consumption inequality. Thus, in response to a negative discount-factor shock, the HANK planner raises interest rates less than one-for-one with r_t^* , tolerating higher output and inflation to curtail the rise in inequality. While this difference relative to zero-liquidity HANK models is easiest to see with discount factor shocks, the same difference is also present in response to other shocks. For example,

one reason the planner does not let output fall as much as y_t^n following a negative productivity shock is that this would require a steeper increase in interest rates, impairing households' ability to self-insure using the bond market.

7 Conclusion

We use an analytically tractable HANK model to study how monetary policy affects inequality, and how this affects optimal monetary policy. Optimal policy differs between HANK and RANK because monetary policy may be able to reduce consumption inequality in HANK; our analytical framework sharply distinguishes between two broad ways in which monetary policy could do this. First, expansionary monetary policy can reduce idiosyncratic consumption risk when income risk is acyclical or countercyclical. Second, monetary policy can reduce fluctuations in between-group inequality arising from unequal exposures to aggregate shocks and policy. Both idiosyncratic risk and unequal exposures lead optimal monetary policy to put some weight on stabilizing output around a higher steady state level, and correspondingly less weight on productive efficiency and price stability, in response to productivity and markup shocks. In response to demand shocks, the same concern with inequality leads optimal policy in HANK to deviate from price stability (which is optimal in RANK); i.e., *divine coincidence* does not hold in HANK. A practical implication of our analysis is that monetary policymakers who are concerned with inequality do not necessarily need to explicitly incorporate some measure of inequality in their reaction function. Introducing the level of output in the target criterion—and accordingly reducing the relative weights on the output gap and prices—adequately captures the planner's concern for consumption inequality.

Our tractable framework can easily be extended to incorporate more shocks or other features of HANK economies. Thus, it can be used as a toolbox to study how these shocks or features affect optimal policy, without having to solve for optimal policy in a computationally demanding HANK model. It also complements quantitative HANK papers by uncovering which model features ultimately drive their results.

References

- Acharya, Sushant and Keshav Dogra**, “Understanding HANK: Insights From a PRANK,” *Econometrica*, May 2020, 88 (3), 1113–1158.
- Angeletos, George-Marios and Laurent Emmanuel Calvet**, “Idiosyncratic Production Risk, Growth and the Business Cycle,” *Journal of Monetary Economics*, 2006, 53, 1095 – 1115.
- Auclert, Adrien**, “Monetary Policy and the Redistribution Channel,” *American Economic Review*, 2019, 109 (6), 2333 – 2367.
- , **Matt Rognlie, and Ludwig Straub**, “The Intertemporal Keynesian Cross,” Working Paper 25020 2018.
- Bayer, Christian, Benjamin Born, and Ralph Luetticke**, “Shocks, Frictions, and Inequality in US Business Cycles,” CEPR Discussion Papers 14364 January 2020.
- Bhandari, Anmol, David Evans, Mikhail Golosov, and Thomas J Sargent**, “Inequality, Business Cycles and Monetary-Fiscal Policy,” *Econometrica*, Forthcoming 2021.
- Bilbiie, Florin and Xavier Ragot**, “Optimal Monetary Policy and Liquidity with Heterogeneous Households,” *Review of Economic Dynamics*, July 2021, 41, 71–95.
- Bilbiie, Florin O.**, “Limited Asset Markets Participation, Monetary Policy and (Inverted) Aggregate Demand Logic,” *Journal of Economic Theory*, May 2008, 140 (1), 162–196.
- , “Monetary Policy and Heterogeneity: An Analytical Framework,” Working Paper, November 2019.
- , “The New Keynesian Cross,” *Journal of Monetary Economics*, 2020, 114, 90–108.
- Blanchard, Olivier Jean and Jordi Galí**, “Real Wage Rigidities and the New Keynesian Model,” *Journal of Money, Credit and Banking*, 2007, 39, 35 – 65.
- Broer, Tobias, Niels-Jakob Harbo Hansen, Per Krusell, and Erik Öberg**, “The New Keynesian Transmission Mechanism: A Heterogenous-Agent Perspective,” *Review of Economic Studies*, 2019, 87 (1), 77–101.
- Caballero, Ricardo J.**, “Consumption Puzzles and Precautionary Savings,” *Journal of Monetary Economics*, 1990, 25 (1), 113 – 136.
- Calvet, Laurent Emmanuel**, “Incomplete Markets and Volatility,” *Journal of Economic Theory*, 2001, 98, 295 – 338.
- Calvo, Guillermo A and Maurice Obstfeld**, “Optimal Time-Consistent Fiscal Policy with Finite Lifetimes,” *Econometrica*, 1988, 56 (2), 411–432.
- Challe, Edouard**, “Uninsured Unemployment Risk and Optimal Monetary Policy in a Zero-Liquidity Economy,” *American Economic Journal: Macroeconomics*, April 2020, 12 (2), 241–83.

- Coibion, Olivier, Yuriy Gorodnichenko, Lorenz Kueng, and John Silvia**, “Innocent Bystanders? Monetary policy and inequality,” *Journal of Monetary Economics*, 2017, 88, 70–89.
- Correia, Isabel, Juan Pablo Nicolini, and Pedro Teles**, “Optimal Fiscal and Monetary Policy: Equivalence Results,” *Journal of Political Economy*, February 2008, 116 (1), 141–170.
- Debortoli, Davide and Jordi Galí**, “Monetary Policy with Heterogeneous Agents: Insights from TANK Models,” Working Paper 2018.
- Farhi, Emmanuel and Iván Werning**, “Monetary policy, bounded rationality, and incomplete markets,” *American Economic Review*, 2019, 109 (11), 3887–3928.
- Galí, Jordi**, *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications, Second Edition*, Princeton University Press, 2015.
- Galí, Jordi, David Lopez-Salido, and Javier Valles**, “Understanding the Effects of Government Spending on Consumption,” *Journal of the European Economic Association*, 2007, 5 (1), 227–270.
- Güvenen, Fatih, Serdar Ozkan, and Jae Song**, “The Nature of Countercyclical Income Risk,” *Journal of Political Economy*, 2014, 122 (3), 621–660.
- Hansen, Niels-Jakob Harbo, Alessandro Lin, and Rui Mano**, “Should Inequality Factor into Central Banks’ Decisions?,” IMF Working Papers September 2020.
- Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante**, “Monetary Policy According to HANK,” *American Economic Review*, March 2018, 108 (3), 697–743.
- , **Giovanni L. Violante, and Justin Weidner**, “The Wealthy Hand-to-Mouth,” *Brookings Papers on Economic Activity*, 2014, 45 (1 (Spring)), 77–153.
- Kekre, Rohan**, “Unemployment Insurance in Macroeconomic Stabilization,” Working Paper 2019.
- Le Grand, François, Alais Martin-Baillon, and Xavier Ragot**, “Should monetary policy care about redistribution? Optimal fiscal and monetary policy with heterogeneous agents,” Working Paper 2021.
- McKay, Alisdair and Ricardo Reis**, “The Role of Automatic Stabilizers in the U.S. Business Cycle,” *Econometrica*, January 2016, 84, 141–194.
- and – , “Optimal Automatic Stabilizers,” *The Review of Economic Studies*, 03 2021.
- , **Emi Nakamura, and Jon Steinsson**, “The Power of Forward Guidance Revisited,” *American Economic Review*, October 2016, 106 (10), 3133.
- Mumtaz, Haroon and Angeliki Theophilopoulou**, “The impact of monetary policy on inequality in the UK. An empirical analysis,” *European Economic Review*, 2017, 98 (C), 410–423.
- Nakajima, Makoto and Vladimir Smirnyagin**, “Cyclical Labor Income Risk,” Working Paper 2019.
- Nisticò, Salvatore**, “Optimal Monetary Policy And Financial Stability In A Non-Ricardian Economy,” *Journal of the European Economic Association*, October 2016, 14 (5), 1225–1252.

Nuño, Galo and Carlos Thomas, “Optimal Monetary Policy with Heterogeneous Agents,” Working Papers 1624, Banco de Espana November 2021.

Ravn, Morten O. and Vincent Sterk, “Macroeconomic Fluctuations with HANK and SAM: an Analytical Approach,” *Journal of the European Economic Association*, 06 2020, 19 (2), 1162–1202.

Rotemberg, Julio J., “Monopolistic Price Adjustment and Aggregate Output,” *Review of Economic Studies*, 1982, 49 (4), 517–531.

Storesletten, Kjetil, Chris Telmer, and Amir Yaron, “Cyclical Dynamics in Idiosyncratic Labor Market Risk,” *Journal of Political Economy*, 2004, 112 (3), 695–717.

Wang, Neng, “Caballero Meets Bewley: The Permanent-Income Hypothesis in General Equilibrium,” *American Economic Review*, June 2003, 93 (3), 927–936.

Werning, Ivan, “Incomplete Markets and Aggregate Demand,” Working Paper 21448, National Bureau of Economic Research August 2015.

Woodford, M., *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton University Press, 2003.

Appendix

A Proof of Proposition 1

The date s problem of an individual i born at date s can be written as

$$\max_{\{c_t^s(i), \ell_t^s(i), a_{t+1}^s(i)\}} -\mathbb{E}_s \sum_{t=s}^{\infty} (\beta\vartheta)^{t-s} \left(\prod_{k=s}^{t-1} \zeta_k \right) \left\{ \frac{1}{\gamma} e^{-\gamma c_t^s(i)} + \rho e^{\frac{1}{\rho}[\ell_t^s(i) - \xi_t^s(i)]} \right\}$$

s.t.

$$c_t^s(i) + q_t a_{t+1}^s(i) = w_t \ell_t^s(i) + (1 - \tau_t^a) a_t^s(i) + D_t - T_t, \quad (\text{A.1})$$

where $a_s^s(i) = \mathcal{T}_s$, $w_t = (1 - \tau^w) \tilde{w}_t$, $\tau_t^a = 0$ for $t > 0$ and ζ_t is the discount-factor shock introduced in Section 6. The optimal labor supply decision of household i is given by

$$\ell_t^s(i) = \rho \ln w_t - \rho \gamma c_t^s(i) + \xi_t^s(i) \quad (\text{A.2})$$

and the Euler equation for all dates $t > 0$ is given by

$$e^{-\gamma c_t^s(i)} = \beta \zeta_t R_t \mathbb{E}_t e^{-\gamma c_{t+1}^s(i)}, \quad (\text{A.3})$$

where we have used the fact that $q_t = \frac{\vartheta}{R_t}$. Next, guess that the consumption decision rule takes the form

$$c_t^s(i) = \mathcal{C}_t + \mu_t x_t^s(i) \quad (\text{A.4})$$

where $x_t^s(i) = (1 - \tau_t^a)a_t^s(i) + w_t(\xi_t^s(i) - \bar{\xi})$ is de-meanded cash on hand. Notice that $x_{t+1}^s(i)$ is normally distributed and so given the guess (A.4), $c_{t+1}^s(i)$ is also normally distributed with mean

$$\mathbb{E}_t c_{t+1}^s(i) = \mathcal{C}_{t+1} + \mu_{t+1} \frac{R_t}{\vartheta} [x_t^s(i) + w_t(\rho \ln w_t + \bar{\xi}) + D_t - T_t - (1 + \rho\gamma w_t) c_t^s(i)]$$

and variance

$$\mathbb{V}_t(c_{t+1}^s(i)) = \mu_{t+1}^2 w_{t+1}^2 \sigma_{t+1}^2$$

Taking logs of (A.3) and using the two expressions above

$$\begin{aligned} c_t^s(i) &= -\frac{1}{\gamma} \ln \beta R_t - \frac{1}{\gamma} \ln \mathbb{E}_t e^{-\gamma c_{t+1}^s(i)} \\ &= -\frac{1}{\gamma} \ln \beta R_t + \mathbb{E}_t c_{t+1}^s(i) - \frac{\gamma}{2} \mathbb{V}_t(c_{t+1}^s(i)) \\ &= -\frac{1}{\gamma} \ln \beta R_t + \mathcal{C}_{t+1} + \mu_{t+1} \frac{R_t}{\vartheta} [x_t^s(i) + w_t(\rho \ln w_t + \bar{\xi}) + D_t - T_t - (1 + \rho\gamma w_t) c_t^s(i)] - \frac{\gamma \mu_{t+1}^2 w_{t+1}^2 \sigma_{t+1}^2}{2} \end{aligned}$$

Combining the $c_t^s(i)$ terms and using (A.4), the above can be rewritten as

$$\begin{aligned} \frac{\mu_{t+1}}{\mu_t} \frac{R_t}{\vartheta} \{ \mathcal{C}_t + \mu_t x_t^s(i) \} &= -\frac{1}{\gamma} \ln \beta R_t + \mathcal{C}_{t+1} + \mu_{t+1} \frac{R_t}{\vartheta} [x_t^s(i) + w_t(\rho \ln w_t + \bar{\xi}) + D_t - T_t] \\ &\quad - \frac{\gamma \mu_{t+1}^2 w_{t+1}^2 \sigma_{t+1}^2}{2} \end{aligned}$$

Matching coefficients, we have for all $t \geq 0$:

$$\mathcal{C}_t = -\frac{\vartheta \mu_t}{\mu_{t+1} R_t} \frac{1}{\gamma} \ln \beta R_t + \frac{\vartheta \mu_t}{\mu_{t+1} R_t} \mathcal{C}_{t+1} + \mu_t [w_t(\rho \ln w_t + \bar{\xi}) + D_t - T_t] - \frac{\vartheta}{R_t} \frac{\mu_t}{\mu_{t+1}} \frac{\gamma \mu_{t+1}^2 w_{t+1}^2 \sigma_{t+1}^2}{2} \quad (\text{A.5})$$

$$\mu_t^{-1} = (1 + \rho\gamma w_t) + \frac{\vartheta}{R_t} \mu_{t+1}^{-1} \quad (\text{A.6})$$

Notice that (A.6) is the same as (12) in the main text. Next, aggregate hours worked are given by $\ell_t = \rho \ln w_t - \gamma \rho \mathcal{C}_t + \bar{\xi}$ and, hence, aggregate income is $y_t = w_t \ell_t + D_t - T_t = w_t \rho \ln w_t - \gamma \rho w_t \mathcal{C}_t + w_t \bar{\xi} + D_t - T_t$. Using this in (A.5) and the fact that $\mathcal{C}_t = y_t$ yields equation (11) in the main text.

B Derivation of Σ recursion

B.1 Evolution of within-cohort cash on hand

Given the consumption function and the definition of x , the evolution of cash on hand can be written as

$$\begin{aligned} x_{t+1}^s(i) &= a_{t+1}^s(i) + w_{t+1}(\xi_{t+1}^s(i) - \bar{\xi}) \\ &= \frac{R_t}{\vartheta} [x_t^s(i) + w_t(\rho \ln w_t + \bar{\xi}) + T_t + D_t - (1 + \rho\gamma w_t) y_t - (1 + \rho\gamma w_t) \mu_t x_t^s(i)] + w_{t+1}(\xi_{t+1}^s(i) - \bar{\xi}) \\ &= \frac{R_t}{\vartheta} [1 - (1 + \rho\gamma w_t) \mu_t] x_t^s(i) + w_{t+1}(\xi_{t+1}^s(i) - \bar{\xi}) \end{aligned}$$

where we have used the fact that $\tau_t^a = 0$ for all dates $t > 0$. In the last line, we have used the definition of aggregate income $y_t = w_t(\rho \ln w_t - \gamma \rho y_t + \bar{\xi}) + T_t + D_t$. Multiplying both sides by μ_{t+1} :

$$\mu_{t+1} x_{t+1}^s(i) = \mu_{t+1} \frac{R_t}{\vartheta} [1 - (1 + \rho \gamma w_t) \mu_t] x_t^s(i) + \mu_{t+1} w_{t+1} (\xi_{t+1}^s(i) - \bar{\xi})$$

and using (12), we have $\mu_{t+1} x_{t+1}^s(i) = \mu_t x_t^s(i) + \mu_{t+1} w_{t+1} (\xi_{t+1}^s(i) - \bar{\xi})$. That is, $\mu_t x_t^s(i)$ follows a random walk within cohort. This implies that in steady state with $\mu_t = \mu$, $x_t^s(i) \sim N(0, (t+1-s)w^2\sigma^2)$ and $a_t^s(i) \sim N(0, (t-s)w^2\sigma^2)$.

B.2 Objective function of planner

Substituting labor supply (10) into the objective function, we can write the date 0 expected utility of individual i from the cohort born at date s going forward as

$$W_0^s(i) = -\frac{1}{\gamma} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta_{t|0} \vartheta^t (1 + \gamma \rho w_t) e^{-\gamma c_t^s(i)} = -\frac{1}{\gamma} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta_{t|0} \vartheta^t (1 + \gamma \rho w_t) e^{-\gamma y_t - \gamma \mu_t x_t^s(i)},$$

where we have used the consumption function (9) and the fact that in equilibrium $C_t = y_t$. We assume that the planner puts a weight of $\varphi^s(i)$ on individual i born at date $s \leq 0$ and $\beta_{s|0} = \beta^s \prod_{k=0}^{s-1} \zeta_k$ on the lifetime welfare of individuals who will be born at date $s > 0$. Then the social welfare is

$$\mathbb{W}_0 = \underbrace{(1 - \vartheta) \sum_{s=-\infty}^0 \vartheta^{-s} \int \varphi^s(i) W_0^s(i) di}_{\text{welfare of those alive at date 0}} + \underbrace{(1 - \vartheta) \sum_{s=1}^{\infty} \beta_{s|0} \int W_s^s(i) di}_{\text{welfare of the unborn at date 0}}$$

Using the definition of $W_0^s(i)$ and $W_s^s(i)$, notice that \mathbb{W}_0 can be written as

$$\mathbb{W}_0 = -\frac{1}{\gamma} \sum_{t=0}^{\infty} \beta^t \underbrace{(1 + \gamma \rho w_t) e^{-\gamma y_t}}_{\text{utility of rep. agent}} \Sigma_t,$$

where Σ_t is defined as

$$\begin{aligned} \Sigma_t &= (1 - \vartheta) \sum_{s=-\infty}^0 \vartheta^{t-s} \int \varphi^s(i) e^{-\gamma(c_t^s(i) - c_t)} di + (1 - \vartheta) \sum_{s=1}^t \int \vartheta^{t-s} e^{-\gamma(c_t^s(i) - c_t)} di \\ &= (1 - \vartheta) \sum_{s=-\infty}^0 \vartheta^{t-s} \int \varphi^s(i) e^{-\gamma \mu_t x_t^s(i)} di + (1 - \vartheta) \sum_{s=1}^t \int \vartheta^{t-s} e^{-\gamma \mu_t x_t^s(i)} di \end{aligned} \quad (\text{B.1})$$

Thus, we can write \mathbb{W}_0 as

$$\mathbb{W}_0 = \sum_{t=0}^{\infty} \beta_{t|0} \mathbb{U}_t, \quad \text{where} \quad \mathbb{U}_t = -\frac{1}{\gamma} (1 + \gamma \rho w_t) e^{-\gamma y_t} \Sigma_t$$

Next, we write (B.1) as

$$\begin{aligned}\Sigma_t &= (1 - \vartheta) \sum_{s=-\infty}^0 \vartheta^{t-s} \int \wp^s(i) e^{-\gamma \mu_t x_t^s(i)} di + (1 - \vartheta) \sum_{s=1}^{t-1} \int \vartheta^{t-s} e^{-\gamma \mu_t x_t^s(i)} di + (1 - \vartheta) \int e^{-\gamma \mu_t x_t^t(i)} di \\ &= \vartheta \left\{ (1 - \vartheta) \sum_{s=-\infty}^0 \vartheta^{t-1-s} \int \wp^s(i) e^{-\gamma \mu_t x_t^s(i)} di + (1 - \vartheta) \sum_{s=1}^{t-1} \int \vartheta^{t-1-s} e^{-\gamma \mu_t x_t^s(i)} di \right\} + (1 - \vartheta) \int e^{-\gamma \mu_t x_t^t(i)} di\end{aligned}$$

Using $\mu_t x_t^s(i) = \mu_{t-1} x_{t-1}^s(i) + \mu_t w_t (\xi_t^s(i) - \bar{\xi})$ from Appendix B.1

$$\begin{aligned}\Sigma_t &= \vartheta \left\{ (1 - \vartheta) \sum_{s=-\infty}^0 \vartheta^{t-1-s} \int \wp^s(i) e^{-\gamma \{\mu_{t-1} x_{t-1}^s(i) + \mu_t w_t (\xi_t^s(i) - \bar{\xi})\}} di \right. \\ &\quad \left. + (1 - \vartheta) \sum_{s=1}^{t-1} \int \vartheta^{t-1-s} e^{-\gamma \{\mu_{t-1} x_{t-1}^s(i) + \mu_t w_t (\xi_t^s(i) - \bar{\xi})\}} di \right\} + (1 - \vartheta) \int e^{-\gamma \mu_t x_t^t(i)} di \\ &= \vartheta e^{\frac{1}{2} \gamma^2 \mu_t^2 w_t^2 \sigma_t^2} \left\{ (1 - \vartheta) \sum_{s=-\infty}^0 \vartheta^{t-1-s} \int \wp^s(i) e^{-\gamma \mu_{t-1} x_{t-1}^s(i)} di + (1 - \vartheta) \sum_{s=1}^{t-1} \int \vartheta^{t-1-s} e^{-\gamma \mu_{t-1} x_{t-1}^s(i)} di \right\} \\ &\quad + (1 - \vartheta) \int e^{-\gamma \mu_t x_t^t(i)} di \\ &= e^{\frac{1}{2} \gamma^2 \mu_t^2 w_t^2 \sigma_t^2} [1 - \vartheta + \vartheta \Sigma_{t-1}]\end{aligned}$$

Taking logs, this is the same as (21) in the main text. For date 0, we have

$$\begin{aligned}\Sigma_0 &= (1 - \vartheta) \sum_{s=-\infty}^0 \vartheta^{-s} \int \wp^s(i) e^{-\gamma \mu_0 x_0^s(i)} di \\ &= (1 - \vartheta) \sum_{s=-\infty}^0 \vartheta^{-s} \int \wp^s(i) e^{-\gamma \mu_0 (1 - \tau_0^a) a_0^s(i)} e^{-\gamma \mu_0 w_0 (\xi_0^s(i) - \bar{\xi})} di \\ &= (1 - \vartheta) e^{\frac{1}{2} \gamma^2 \mu_0^2 w_0^2 \sigma_0^2} \sum_{s=-\infty}^0 \vartheta^{-s} \int \wp^s(i) e^{-\gamma \mu_0 (1 - \tau_0^a) a_0^s(i)} di\end{aligned}$$

where we use the fact that $x_0^s(i) = (1 - \tau_0^a) a_0^s(i) + w_0 (\xi_0^s(i) - \bar{\xi})$. Next, we restrict $\wp^s(i) = e^{\gamma \alpha a_0^s(i)}$ where $\alpha \geq 0$ measures the planner's tolerance for pre-existing wealth inequality at date 0. Then we can write Σ_0 as

$$\Sigma_0 = (1 - \vartheta) e^{\frac{1}{2} \gamma^2 \mu_0^2 w_0^2 \sigma_0^2} \sum_{s=-\infty}^0 \vartheta^{-s} \int e^{-\gamma [\alpha - \mu_0 (1 - \tau_0^a)] a_0^s(i)} di$$

Since $a_0^s(i) \sim N(0, -s w^2 \sigma^2)$ for $s \leq 0$, this can be rewritten as

$$\Sigma_0 = (1 - \vartheta) e^{\frac{1}{2} \gamma^2 \mu_0^2 w_0^2 \sigma_0^2} \sum_{s=-\infty}^0 \left(\vartheta e^{\frac{\gamma^2 \mu_0^2 w_0^2 \sigma_0^2}{2} \left[\frac{\alpha - \mu_0 (1 - \tau_0^a)}{\mu} \right]^2} \right)^{-s} = \frac{(1 - \vartheta) e^{\frac{1}{2} \gamma^2 \mu_0^2 w_0^2 \sigma_0^2}}{1 - \vartheta e^{\frac{\gamma^2 \mu_0^2 w_0^2 \sigma_0^2}{2} \left[\frac{\alpha - \mu_0 (1 - \tau_0^a)}{\mu} \right]^2}}$$

This is the same as (30) and with $\alpha = 0$, this is the same as (22).

B.2.1 The utilitarian planner

The utilitarian planner is one who assigns $\wp^s(i) = 1$ for all households alive at date 0. In this case the expression for Σ_0 can be simplified to:

$$\Sigma_t = (1 - \vartheta) \sum_{s=-\infty}^t \vartheta^{t-s} e^{\frac{1}{2}\gamma^2\sigma_c^2(s,t)}$$

To see this, impose $\wp^s(i) = 1$ in (B.1), which can then be written as

$$\Sigma_t = (1 - \vartheta) \sum_{s=-\infty}^t \vartheta^{t-s} \int e^{-\gamma\mu_t x_i^s(i)} di$$

Given the consumption function (9) and the normality of shocks, the consumption of newly born individuals at any date s is normally distributed with mean y_s and variance $\sigma_c^2(s, s) = \mu_s^2 w_s^2 \sigma_s^2$ since they all have zero wealth. Given the linearity of the budget constraint, it follows that newly born agents' savings decisions $a_{s+1}^s(i)$ are also normally distributed with mean 0 and variance $\sigma_a^2(s+1, s) = \left(\frac{R_s}{\vartheta}\right)^2 [1 - (1 + \gamma\rho w_s) \mu_s]^2 w_s^2 \sigma_s^2$. By induction, it follows that for any cohort born at date s , the cross-sectional distribution of consumption at any date $t > s$ is normal with mean y_t and variance

$$\sigma_c^2(t, s) = \mu_t^2 \sigma_a^2(t, s) + \mu_t^2 w_t^2 \sigma_t^2, \quad (\text{B.2})$$

while the distribution of asset holdings is normal with mean 0 and variance

$$\sigma_a^2(t, s) = \frac{R_{t-1}^2}{\vartheta^2} [1 - (1 + \gamma\rho w_{t-1}) \mu_{t-1}]^2 [\sigma_a^2(t-1, s) + w_{t-1}^2 \sigma_{t-1}^2] \quad (\text{B.3})$$

C Some auxiliary results

In the proofs that follow, we shall make liberal use of the following assumptions and results.

Assumption 1. *Throughout the paper, we shall assume that:*

1. $\vartheta \geq \frac{1}{2}$
2. $\beta\vartheta > e^{-\frac{1}{2}} = 0.61$
3. $\sigma < \min\{\bar{\sigma}_1, \bar{\sigma}_2\}$ where $\bar{\sigma}_1 = \sqrt{\frac{2\rho^2 \ln \vartheta^{-1}}{\left(\frac{\gamma\rho}{1+\gamma\rho}\right)^2 \left(\frac{2(1-\frac{\phi}{\gamma}) \ln \vartheta^{-1}}{1+2 \ln \vartheta} + (1-\beta)\right)^2}}$ and $\bar{\sigma}_2 = \frac{\rho}{\sqrt{(1-\beta\vartheta)(1+\gamma\rho+1-\beta\vartheta)}}$.

Lemma 3. *Given that $\beta\vartheta > e^{-\frac{1}{2}}$, we have $\Lambda < 1$ and $\tilde{\beta} < 1$.*

Proof. Recall that in steady state, $\Lambda = \gamma^2 \mu^2 w^2 \sigma^2 > 0$, i.e.,

$$\Lambda = \frac{\sigma^2}{\rho^2} \left(\frac{\gamma\rho w}{1 + \gamma\rho w} \right)^2 \left(1 - \beta\theta e^{\frac{\Lambda}{2}} \right)^2$$

Rearranging:

$$f(\Lambda) \equiv \frac{\Lambda}{\left(1 - \beta\theta e^{\frac{\Lambda}{2}}\right)^2} = \frac{\sigma^2}{\rho^2} \left(\frac{\gamma\rho w}{1 + \gamma\rho w}\right)^2 \quad (\text{C.1})$$

Now, $f(\Lambda)$ is increasing for $\Lambda < \Lambda^* \equiv -2 \ln \beta\vartheta < 1$ given our assumption, and goes to ∞ as $\Lambda \rightarrow \Lambda^*$. For any values of σ and ρ , we can find some $0 < \bar{\Lambda} < \Lambda^*$ satisfying $f(\bar{\Lambda}) = \frac{\sigma^2}{\rho^2}$. Thus, any solution to (C.1) must satisfy $\Lambda \leq \bar{\Lambda} < \Lambda^* < 1$. By construction, for any $\Lambda < \Lambda^*$, $\tilde{\beta} = \beta\theta e^{\frac{\Lambda}{2}} < 1$. \square

Lemma 4. For $\sigma < [0, \bar{\sigma}_1)$, we have $\vartheta e^{\frac{\Lambda}{2}} < 1$.

Proof. First we show that $\vartheta e^{\frac{\Lambda}{2}} = 1$ implies that $\sigma = \bar{\sigma}$. Starting from the expression for wages in steady state, using $\vartheta e^{\frac{\Lambda}{2}} = 1$ we have

$$\frac{w - 1}{1 + \gamma\rho w} = \frac{\Theta - 1 + \Lambda}{(1 - \Lambda)(1 - \tilde{\beta})} = \frac{2\left(1 - \frac{\phi}{\gamma}\right) \ln \vartheta^{-1}}{(1 + 2 \ln \vartheta)(1 - \beta)}$$

Add 1 to both sides and multiply by $\frac{\gamma\rho}{1 + \gamma\rho}$ to get

$$\frac{\gamma\rho w}{1 + \gamma\rho w} = \left[\frac{2 \ln \vartheta^{-1} \left(1 - \frac{\phi}{\gamma}\right)}{(1 + 2 \ln \vartheta)(1 - \tilde{\beta})} + 1 \right] \frac{\gamma\rho}{1 + \gamma\rho}$$

Next, using the expression above in the definition of Λ , we have

$$\sigma^2 = \frac{2 \ln \vartheta^{-1}}{\left(\frac{\gamma\rho}{1 + \gamma\rho}\right)^2 \left(\frac{-2 \ln \vartheta \left(1 - \frac{\phi}{\gamma}\right)}{(1 + 2 \ln \vartheta)} + (1 - \beta)\right)^2},$$

which is the same as $\bar{\sigma}_1$ defined in Assumption 1. Second, note that when $\sigma^2 = 0$, we have $\Lambda = 0$ and $\vartheta e^{\frac{\Lambda}{2}} = \vartheta < 1$. By continuity it follows that for $\sigma \in [0, \bar{\sigma}_1)$, we have $\vartheta e^{\frac{\Lambda}{2}} < 1$. \square

Corollary 2. The following is true:

$$1 - \beta^{-1} \tilde{\beta}(1 - \Lambda) > 0$$

Proof.

$$1 - \beta^{-1} \tilde{\beta}(1 - \Lambda) = 1 - \vartheta e^{\frac{\Lambda}{2}}(1 - \Lambda) > 0$$

\square

D First-order condition of the planning problem

D.1 Optimally set fiscal instruments

The planner chooses τ_0^a and τ^w optimally absent shocks. This problem can be written as

$$\max_{\{w_t, y_t, \mu_t, \Sigma_t, \Pi_t\}_{t=0}^{\infty}, \tau_0^a, \tau^w} \sum_{t=0}^{\infty} \beta^t \left\{ -\frac{1}{\gamma} (1 + \gamma \rho w_t) e^{-\gamma y_t \Sigma_t} \right\}$$

s.t.

$$\gamma y_t = \gamma y_{t+1} - \ln \beta \vartheta + \ln \mu_{t+1} + \ln [\mu_t^{-1} - (1 + \gamma \rho w_t)] - \frac{\gamma^2 \mu_{t+1}^2 w^2 \sigma^2}{2} e^{2\phi(y_{t+1}-y)} \quad (\text{D.1})$$

$$(\Pi_t - 1) \Pi_t = \frac{\lambda}{(\lambda - 1) \Psi} \left[1 - \frac{(1 - \tau^w)}{w_t} \right] + \beta \left(\frac{w_{t+1}}{w_t} \right) (\Pi_{t+1} - 1) \Pi_{t+1} \quad (\text{D.2})$$

$$\ln \Sigma_t = \frac{\gamma^2 \mu_t^2 w^2 \sigma^2 e^{2\phi(y_t-y)}}{2} + \ln [1 - \vartheta + \vartheta \Sigma_{t-1}] + \mathbb{I}(t=0) \ln \left[\frac{1 - \vartheta e^{\frac{\Lambda}{2}}}{1 - \vartheta e^{\frac{\Lambda}{2} \left(\frac{\alpha - (1 - \tau_0^a) \mu_0}{\mu} \right)^2}} \right] \quad (\text{D.3})$$

$$y_t = \frac{\rho \ln w_t + \bar{\xi}}{1 + \gamma \rho} - \frac{\Psi}{2} (\Pi_t - 1)^2 \quad (\text{D.4})$$

The optimal decisions satisfy the following where $M_{1,t}$ denotes the multiplier on the date t aggregate Euler equation, $M_{2,t}$ is the multiplier on the date t Phillips curve, $M_{3,t}$ is the multiplier on the date t Σ recursion and $M_{4,t}$ is the multiplier on the relationship between y_t, w_t and Π_t

FOC wrt w_t

$$\begin{aligned} & \mathbb{U}_t \frac{\gamma \rho w_t}{1 + \gamma \rho w_t} + M_{2,t-1} \left(\frac{w_t}{w_{t-1}} \right) (\Pi_t - 1) \Pi_t - M_{1,t} \frac{\gamma \rho w_t}{\mu_t^{-1} - (1 + \gamma \rho w_t)} \\ & + M_{2,t} \left\{ \frac{\lambda_t}{\Psi (\lambda - 1)} \frac{(1 - \tau^w)}{w_t} - \beta \left(\frac{w_{t+1}}{w_t} \right) (\Pi_{t+1} - 1) \Pi_{t+1} \right\} - M_{4,t} \frac{\rho}{1 + \gamma \rho} = 0 \end{aligned} \quad (\text{D.5})$$

FOC y_t

$$-\gamma \mathbb{U}_t - \gamma M_{1,t} + \beta^{-1} M_{1,t-1} \left\{ \gamma - \phi \gamma^2 \mu_t^2 w^2 \sigma^2 e^{2[\phi(y_t-y)]} \right\} + M_{3,t} \phi \gamma^2 \mu_t^2 w^2 \sigma^2 e^{2\phi(y_t-y)} + M_{4,t} = 0 \quad (\text{D.6})$$

FOC μ_t

$$\begin{aligned} & -M_{1,t} \frac{\mu_t^{-1}}{\mu_t^{-1} - 1 - \gamma \rho w_t} + \beta^{-1} M_{1,t-1} \left[1 - \gamma^2 \sigma^2 w^2 \mu_t^2 e^{2\phi(y_t-y)} \right] + M_{3,t} \gamma^2 \sigma^2 w^2 \mu_t^2 e^{2\phi(y_t-y)} \\ & + \mathbb{I}(t=0) M_{3,0} \left\{ \frac{\vartheta e^{\frac{\Lambda}{2} \left(\frac{\alpha - (1 - \tau_0^a) \mu_0}{\mu} \right)^2}}{1 - \vartheta e^{\frac{\Lambda}{2} \left(\frac{\alpha - (1 - \tau_0^a) \mu_0}{\mu} \right)^2}} \Lambda \left(\frac{\alpha - (1 - \tau_0^a) \mu_0}{\mu} \right) \left(\frac{1 - \tau_0^a}{\mu} \right) \right\} = 0 \end{aligned} \quad (\text{D.7})$$

FOC Σ_t

$$\mathbb{U}_t - M_{3,t} + \beta M_{3,t+1} \frac{\vartheta \Sigma_t}{1 - \vartheta + \vartheta \Sigma_t} = 0 \quad (\text{D.8})$$

FOC Π_t

$$M_{2,t-1} \left(\frac{w_t}{w_{t-1}} \right) (2\Pi_t - 1) - M_{2,t} (2\Pi_t - 1) + \Psi M_{4,t} (\Pi_t - 1) = 0 \quad (\text{D.9})$$

FOC τ_0^a

$$M_{3,0} \frac{\vartheta e^{\frac{\Lambda}{2} \left(\frac{\alpha - (1 - \tau_0^a) \mu_0}{\mu} \right)^2}}{1 - \vartheta e^{\frac{\Lambda}{2} \left(\frac{\alpha - (1 - \tau_0^a) \mu_0}{\mu} \right)^2}} \Lambda \left(\frac{\alpha - (1 - \tau_0^a) \mu_0}{\mu} \right) \frac{\mu_0}{\mu} = 0 \quad (\text{D.10})$$

FOC τ^w

$$\sum_{t=0}^{\infty} \beta^t \frac{M_{2,t}}{w_t} = 0 \quad (\text{D.11})$$

We guess and verify that the optimal solution features $y_t = y, w_t = w, \mu_t = \mu$ and $\Pi_t = 1$ such that $\frac{w-1}{1+\gamma\rho w} = \Omega$. Plugging in the guesses into the FOCs, (D.9) implies $M_{2,t-1} = M_{2,t}$. Given this, (D.11) implies that $M_{2,t} = 0$ for all $t \geq 0$. Using $\mu_t = \mu$ in (D.10), we have

$$1 - \tau_0^a = \frac{\alpha}{\mu}$$

as long as $M_{3,0} \neq 0$. Next, we show that $M_{3,0} \neq 0$. To see this, notice that (D.8) can be rewritten as

$$1 - \frac{M_{3,t}}{\mathbb{U}_t} + \beta \vartheta \frac{M_{3,t+1}}{\mathbb{U}_{t+1}} \frac{\Sigma_{t+1}}{1 - \vartheta + \vartheta \Sigma_t} = 0 \quad \Rightarrow \quad \frac{M_{3,t}}{\mathbb{U}_t} = 1 + \tilde{\beta} \frac{M_{3,t+1}}{\mathbb{U}_{t+1}}, \quad (\text{D.12})$$

where we have used the fact that $\mathbb{U}_{t+1}/\mathbb{U}_t = \Sigma_{t+1}/\Sigma_t$ and $\frac{\Sigma_{t+1}}{1 - \vartheta + \vartheta \Sigma_t} = e^{\frac{\Lambda}{2}}$ since $y_t = y, w_t = w$ and $\mu_t = \mu$. Iterating forward, we get $M_{3,t}/\mathbb{U}_t = (1 - \tilde{\beta})^{-1} \neq 0$.

Using this, (D.5), (D.6) and (D.7) become

$$\frac{(1 + \gamma\rho)w}{1 + \gamma\rho w} + (1 - \tilde{\beta}^{-1}) \frac{M_{1,t}}{\mathbb{U}_t} \frac{(1 + \gamma\rho)w}{1 + \gamma\rho w} - \frac{1}{\gamma} \frac{M_{4,t}}{\mathbb{U}_t} = 0 \quad (\text{D.13})$$

$$-1 - \frac{M_{1,t}}{\mathbb{U}_t} + \beta^{-1} \frac{M_{1,t-1}}{\mathbb{U}_t} \left(1 - \frac{\phi\Lambda}{\gamma} \right) + \frac{1}{1 - \tilde{\beta}} \frac{\phi\Lambda}{\gamma} + \frac{1}{\gamma} \frac{M_{4,t}}{\mathbb{U}_t} = 0 \quad (\text{D.14})$$

and

$$-\tilde{\beta}^{-1} \frac{M_{1,t}}{\mathbb{U}_t} + \beta^{-1} (1 - \Lambda) \frac{M_{1,t-1}}{\mathbb{U}_t} + \Lambda \frac{1}{1 - \tilde{\beta}} = 0, \quad (\text{D.15})$$

where we have used $\mu^{-1} = \frac{1+\gamma\rho w}{1-\tilde{\beta}}$. Next, combining (D.13) and (D.14), we get

$$\frac{w-1}{1+\gamma\rho w} + \left[(1-\tilde{\beta}^{-1}) \frac{w-1}{1+\gamma\rho w} - \tilde{\beta}^{-1} \right] \frac{M_{1,t}}{\mathbb{U}_t} + \beta^{-1} \Theta \frac{M_{1,t-1}}{\mathbb{U}_t} + (1-\Theta) \frac{1}{1-\tilde{\beta}} = 0 \quad (\text{D.16})$$

Combining (D.15) with (D.16), we get

$$\left[\frac{w-1}{1+\gamma\rho w} - \frac{\Theta-1+\Lambda}{(1-\tilde{\beta})(1-\Lambda)} \right] \left[1 + \beta^{-1} (1-\tilde{\beta}) \frac{M_{1,t-1}}{\mathbb{U}_t} \right] = 0 \quad (\text{D.17})$$

In particular, this must be true at date 0 when $M_{1,-1} = 0$. This requires

$$\frac{w-1}{1+\gamma\rho w} = \frac{\Theta-1+\Lambda}{(1-\tilde{\beta})(1-\Lambda)},$$

which is the same as the definition of Ω in (23) in the main text. Given that w satisfies this restriction, (D.17) is also true at all subsequent dates. Since $\Pi = 1$, this implies from the Phillips curve that

$$1 - \tau^w = w^{-1} = \frac{1 + \Omega}{1 - \gamma\rho\Omega}$$

It follows that all FOCs and constraints are satisfied by our guesses and given the optimal values of τ_0^a and $\tau^w, y_t, \Pi_t, \mu_t, w_t$ remain at their steady state level absent shocks.

D.2 Steady state of the optimal plan

Imposing steady state on (D.3), one gets

$$\Sigma = \frac{(1-\vartheta) e^{\frac{\Lambda}{2}}}{1-\vartheta e^{\frac{\Lambda}{2}}}$$

We already know from (D.12) in steady state that $m_3 = \frac{1}{1-\tilde{\beta}}$ and that $m_2 = 0$ from (D.11) where $m_i = M_i/\mathbb{U}$ for $i = \{1, 2, 3, 4\}$. Next, imposing steady state in (D.15) yields

$$m_1 = \frac{\tilde{\beta}}{1-\tilde{\beta}} \left[\frac{\Lambda}{1-\beta^{-1}\tilde{\beta}(1-\Lambda)} \right] \quad (\text{D.18})$$

Notice that since $\Lambda = 0$ in RANK, we have $m_1 = 0$. Finally, using this in (D.13) and imposing steady state yields

$$m_4 = \gamma \frac{(1-\beta^{-1}\tilde{\beta})(1-\Lambda)}{1-\beta^{-1}\tilde{\beta}(1-\Lambda)} (1+\Omega) \quad (\text{D.19})$$

where $\Omega = \frac{\Theta-1+\Lambda}{(1-\Lambda)(1-\tilde{\beta})}$.

D.3 Optimal monetary policy given optimally set fiscal policy

The planning problem can be written as

$$\max_{\{w_t, y_t, \mu_t, \Sigma_t, \Pi_t\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t \left(\prod_{k=0}^{t-1} \zeta_k \right) \left\{ -\frac{1}{\gamma} (1 + \gamma \rho w_t) e^{-\gamma y_t \Sigma_t} \right\}$$

s.t.

$$\gamma y_t = \gamma y_{t+1} - \ln \beta \vartheta - \ln \zeta_t + \ln \mu_{t+1} + \ln [\mu_t^{-1} - (1 + \gamma \rho w_t)] - \frac{\gamma^2 \mu_{t+1}^2 w^2 \sigma^2}{2} e^{2\phi(y_{t+1}-y)+2\varsigma_{t+1}} \quad (\text{D.20})$$

$$(\Pi_t - 1) \Pi_t = \frac{\lambda_t}{(\lambda_t - 1) \Psi} \left[1 - \frac{1 + \Omega}{1 - \gamma \rho \Omega} \frac{z_t}{\lambda_t^{-1} w_t} \right] + \beta \left(\frac{z_t w_{t+1}}{z_{t+1} w_t} \right) (\Pi_{t+1} - 1) \Pi_{t+1} \quad (\text{D.21})$$

$$\ln \Sigma_t = \frac{\gamma^2 \mu_t^2 w^2 \sigma^2}{2} e^{2\phi(y_t-y)+\varsigma_t} + \ln [1 - \vartheta + \vartheta \Sigma_{t-1}] + \mathbb{I}(t=0) \ln \left[\frac{1 - \vartheta e^{\frac{\Lambda}{2}}}{1 - \vartheta e^{\left(\frac{\alpha}{\mu}\right)^2 \frac{\Lambda}{2} \left(\frac{\mu - \mu_0}{\mu}\right)^2}} \right] \quad (\text{D.22})$$

$$y_t = z_t \frac{\rho \ln w_t + \bar{\xi}}{1 + \gamma \rho z_t} - \frac{\Psi}{2} (\Pi_t - 1)^2 \quad (\text{D.23})$$

and $\Sigma_{-1} = 1$. The problem can be written as a Lagrangian:

$$\begin{aligned} \mathcal{L} = & \sum_{t=0}^{\infty} \beta^t \left(\prod_{k=0}^{t-1} \zeta_k \right) \left\{ -\frac{1}{\gamma} (1 + \gamma \rho w_t) e^{-\gamma y_t \Sigma_t} \right\} \\ & + \sum_{t=0}^{\infty} \beta^t \left(\prod_{k=0}^{t-1} \zeta_k \right) M_{1,t} \left\{ \gamma y_{t+1} - \ln \beta \vartheta - \ln \zeta_t + \ln \mu_{t+1} + \ln [\mu_t^{-1} - (1 + \gamma \rho w_t)] \right. \\ & \left. - \frac{\gamma^2 \mu_{t+1}^2 w^2 \sigma^2}{2} e^{2\phi(y_{t+1}-y)+2\varsigma_{t+1}} - \gamma y_t \right\} \\ & + \sum_{t=0}^{\infty} \beta^t \left(\prod_{k=0}^{t-1} \zeta_k \right) M_{2,t} \left\{ \frac{\lambda_t}{(\lambda_t - 1) \Psi} \left[1 - \frac{(1 - \tau^w) z_t}{(1 - \tau) \lambda_t w_t} \right] + \beta \left(\frac{z_t w_{t+1}}{z_{t+1} w_t} \right) (\Pi_{t+1} - 1) \Pi_{t+1} - (\Pi_t - 1) \Pi_t \right\} \\ & + M_{3,0} \left\{ \frac{\gamma^2 \mu_0^2 w^2 \sigma^2}{2} e^{2\phi(y_0-y)+2\varsigma_0} + \ln [1 - \vartheta + \vartheta \Sigma_{-1}] + \ln \left[\frac{1 - \vartheta e^{\frac{\Lambda}{2}}}{1 - \vartheta e^{\left(\frac{\alpha}{\mu}\right)^2 \frac{\Lambda}{2} \left(\frac{\mu - \mu_0}{\mu}\right)^2}} \right] - \ln \Sigma_0 \right\} \\ & + \sum_{t=1}^{\infty} \beta^t \left(\prod_{k=0}^{t-1} \zeta_k \right) M_{3,t} \left\{ \frac{\gamma^2 \mu_t^2 w^2 \sigma^2}{2} e^{2\phi(y_t-y)+2\varsigma_t} + \ln [1 - \vartheta + \vartheta \Sigma_{t-1}] - \ln \Sigma_t \right\} \\ & + \sum_{t=0}^{\infty} \beta^t \left(\prod_{k=0}^{t-1} \zeta_k \right) M_{4,t} \left\{ y_t - z_t \frac{\rho \ln w_t + \bar{\xi}}{1 + \rho \gamma z_t} + \frac{\Psi}{2} (\Pi_t - 1)^2 \right\} \end{aligned}$$

The optimal decisions satisfy:

FOC wrt w_t

$$\begin{aligned} & \mathbb{U}_t \frac{\gamma \rho w_t}{1 + \gamma \rho w_t} + \zeta_{t-1}^{-1} M_{2,t-1} \left(\frac{z_{t-1} w_t}{z_t w_{t-1}} \right) (\Pi_t - 1) \Pi_t - M_{1,t} \frac{\gamma \rho w_t}{\mu_t^{-1} - (1 + \gamma \rho w_t)} \\ & + M_{2,t} \left\{ \frac{\lambda_t}{\Psi (\lambda_t - 1) (1 - \tau) \lambda_t w_t} - \beta \left(\frac{z_t w_{t+1}}{z_{t+1} w_t} \right) (\Pi_{t+1} - 1) \Pi_{t+1} \right\} - M_{4,t} \frac{\rho z_t}{1 + \gamma \rho z_t} = 0 \quad (\text{D.24}) \end{aligned}$$

FOC y_t

$$-\gamma \mathbb{U}_t - \gamma M_{1,t} + \beta^{-1} \zeta_{t-1}^{-1} M_{1,t-1} \left\{ \gamma - \phi \gamma^2 \mu_t^2 w^2 \sigma^2 e^{2[\phi(y_t-y)+\varsigma_t]} \right\} + M_{3,t} \phi \gamma^2 \mu_t^2 w^2 \sigma^2 e^{2\phi(y_t-y)+2\varsigma_t} + M_{4,t} = 0 \quad (\text{D.25})$$

FOC μ_t

$$-M_{1,t} \frac{\mu_t^{-1}}{\mu_t^{-1} - 1 - \gamma \rho w_t} + \beta^{-1} \zeta_{t-1}^{-1} M_{1,t-1} \left[1 - \gamma^2 \sigma^2 w^2 \mu_t^2 e^{2\phi(y_t-y)+2\varsigma_t} \right] + M_{3,t} \gamma^2 \sigma^2 w^2 \mu_t^2 e^{2\phi(y_t-y)+2\varsigma_t} + \mathbb{I}(t=0) M_{3,0} \left\{ \frac{\vartheta e^{\left(\frac{\alpha}{\mu}\right)^2 \frac{\Lambda}{2} \left(\frac{\mu-\mu_0}{\mu}\right)^2}}{1 - \vartheta e^{\left(\frac{\alpha}{\mu}\right)^2 \frac{\Lambda}{2} \left(\frac{\mu-\mu_0}{\mu}\right)^2}} \Lambda \left(\frac{\mu-\mu_0}{\mu^2}\right) \left(\frac{\alpha}{\mu}\right)^2 \right\} = 0 \quad (\text{D.26})$$

FOC Σ_t

$$\mathbb{U}_t - M_{3,t} + \beta \zeta_t M_{3,t+1} \frac{\vartheta \Sigma_t}{1 - \vartheta + \vartheta \Sigma_t} = 0 \quad (\text{D.27})$$

FOC Π_t

$$\zeta_{t-1}^{-1} M_{2,t-1} \left(\frac{z_{t-1} w_t}{z_t w_{t-1}} \right) (2\Pi_t - 1) - M_{2,t} (2\Pi_t - 1) + \Psi M_{4,t} (\Pi_t - 1) = 0 \quad (\text{D.28})$$

D.4 State contingent τ_0^a

Unlike in the main paper, if we allowed the planner to set τ_0^a in a state contingent fashion (varying with shocks), the optimality condition with respect to τ_0^a given by equation (D.10) holds for any μ_0 , not just absent shocks. This implies that the tx is optimally set to:

$$1 - \tau_0^{a*} = \frac{\alpha}{\mu_0}$$

Consequently, (D.3) becomes

$$\ln \Sigma_t = \frac{\gamma^2 \mu_t^2 w^2 \sigma^2 e^{2\{\phi(y_t-y)\}}}{2} + \ln [1 - \vartheta + \vartheta \Sigma_{t-1}]$$

for any α at all dates $t \geq 0$. Since α does not appear explicitly in any of the other constraints or the objective function, it follows that the optimal path of all the variables is the same as that chosen by the utilitarian planner.

E Linearized equations

E.1 Linearized dynamics equations

In the baseline model with all four shocks, the linearized equations describing aggregate dynamics are

$$\hat{y}_t = \Theta \hat{y}_{t+1} - \frac{1}{\gamma} \left(\hat{i}_t - \pi_{t+1} + \hat{\zeta}_t \right) - \frac{\Lambda}{\gamma} \hat{\varsigma}_{t+1} - \frac{\Lambda}{\gamma} \hat{\mu}_{t+1} \quad (\text{E.1})$$

$$\hat{\mu}_t = -(1 - \tilde{\beta}) \frac{\gamma \rho w}{1 + \gamma \rho w} \frac{\hat{w}_t}{w} + \tilde{\beta} \left(\hat{\mu}_{t+1} + \hat{i}_t - \pi_{t+1} \right) \quad (\text{E.2})$$

$$\hat{y}_t = \frac{\rho}{1 + \gamma \rho} \frac{\hat{w}_t}{w} + \frac{y}{1 + \gamma \rho} \hat{z}_t \quad (\text{E.3})$$

$$\pi_t = \beta \pi_{t+1} + \kappa \left(\hat{y}_t - \frac{y + \rho}{1 + \gamma \rho} \hat{z}_t \right) + \frac{\kappa \rho}{1 + \gamma \rho} \hat{\lambda}_t \quad (\text{E.4})$$

Using (E.2) and (E.3) to substitute out i_t and \hat{w}_t/w , the IS equation (E.1) can be rewritten as

$$\begin{aligned} \gamma \left[1 + (1 - \tilde{\beta}) \frac{w - 1}{1 + \gamma \rho w} \right] \hat{y}_t + \hat{\mu}_t &= \tilde{\beta} (1 - \Lambda) \left[\gamma \frac{\Theta}{1 - \Lambda} \hat{y}_{t+1} + \hat{\mu}_{t+1} \right] - \tilde{\beta} \hat{\zeta}_t \\ &+ \gamma (1 - \tilde{\beta}) \left[1 + \frac{w - 1}{1 + \gamma \rho w} \right] \frac{y}{1 + \gamma \rho} \hat{z}_t - \tilde{\beta} \Lambda \hat{\varsigma}_{t+1} = 0 \end{aligned}$$

Using the fact that $\Omega = \frac{w-1}{1+\gamma\rho w}$ and $1 + (1 - \tilde{\beta}) \Omega = \frac{\Theta}{1-\Lambda}$, the IS can be written as

$$\begin{aligned} \gamma \left[1 + (1 - \tilde{\beta}) \Omega \right] \hat{y}_t + \hat{\mu}_t &= \tilde{\beta} (1 - \Lambda) \left\{ \gamma \left[1 + (1 - \tilde{\beta}) \Omega \right] \hat{y}_{t+1} + \hat{\mu}_{t+1} \right\} - \tilde{\beta} \hat{\zeta}_t \\ &+ (1 + \Omega) \frac{(1 - \tilde{\beta}) \gamma y}{1 + \gamma \rho} \hat{z}_t - \tilde{\beta} \Lambda \hat{\varsigma}_{t+1} = 0 \end{aligned}$$

Solving this equation forward yields

$$\gamma \left[1 + (1 - \tilde{\beta}) \Omega \right] \hat{y}_t + \hat{\mu}_t = \Gamma_t, \quad (\text{E.5})$$

where

$$\begin{aligned} \Gamma_t &= \sum_{s=0}^{\infty} \tilde{\beta}^s (1 - \Lambda)^s \left\{ \gamma (1 + \Omega) \frac{(1 - \tilde{\beta}) y}{1 + \gamma \rho} \hat{z}_{t+s} - \tilde{\beta} \hat{\zeta}_{t+s} - \tilde{\beta} \Lambda \hat{\varsigma}_{t+s+1} \right\} \\ &= \frac{\gamma y}{1 + \gamma \rho} \frac{(1 + \Omega) (1 - \tilde{\beta})}{1 - \tilde{\beta} (1 - \Lambda)} \hat{z}_t - \frac{\tilde{\beta}}{1 - \tilde{\beta} (1 - \Lambda)} \frac{\hat{\zeta}_t}{\varrho_{\zeta}} - \frac{\tilde{\beta} \varrho_{\varsigma} \Lambda}{1 - \tilde{\beta} (1 - \Lambda)} \frac{\varsigma_t}{\varrho_{\varsigma}}, \end{aligned} \quad (\text{E.6})$$

where we have used the fact that $\hat{z}_{t+s} = \varrho_z^s \hat{z}_t$, $\hat{\zeta}_{t+s} = \varrho_{\zeta}^s \hat{\zeta}_t$ and $\hat{\varsigma}_{t+k} = \varrho_{\varsigma}^k \hat{\varsigma}_t$ in the second equality. Next, the linearized Σ_t recursion is

$$\frac{\hat{\Sigma}_t}{\Sigma} = -\gamma (\Theta - 1) \hat{y}_t + \Lambda (\hat{\mu}_t + \varsigma_t) + \frac{\tilde{\beta}}{\beta} \frac{\hat{\Sigma}_{t-1}}{\Sigma}$$

Using equation (E.5), we can substitute out $\hat{\mu}_t$ from this expression:

$$\frac{\hat{\Sigma}_t}{\Sigma} = -\gamma(\Theta - 1)\hat{y}_t + \Lambda \left(\Gamma_t - \gamma \left[1 + (1 - \tilde{\beta})\Omega \right] \hat{y}_t + s_t \right) + \frac{\tilde{\beta}}{\beta} \frac{\hat{\Sigma}_{t-1}}{\Sigma}$$

where

$$\Gamma_t = \frac{\gamma y}{1 + \gamma \rho} \frac{(1 + \Omega)(1 - \tilde{\beta})}{1 - \tilde{\beta}(1 - \Lambda)} \hat{z}_t - \frac{\tilde{\beta}}{1 - \tilde{\beta}(1 - \Lambda)} \hat{\zeta}_t - \frac{\tilde{\beta} \varrho_\varsigma \Lambda}{1 - \tilde{\beta}(1 - \Lambda) \varrho_\varsigma} s_t$$

is defined in equation (E.6). Then we can write the linearized Σ_t recursion as

$$\frac{\hat{\Sigma}_t}{\Sigma} = -\gamma(1 - \tilde{\beta})\Omega \hat{y}_t + \Lambda \bar{\Gamma}_t + \frac{\tilde{\beta}}{\beta} \frac{\hat{\Sigma}_{t-1}}{\Sigma},$$

where $\bar{\Gamma}_t = \Gamma_z \hat{z}_t + \Gamma_\zeta \hat{\zeta}_t + \Gamma_\varsigma \hat{s}_t$, where

$$\begin{aligned} \Gamma_z &= \frac{\gamma y}{1 + \gamma \rho} \frac{(1 + \Omega)(1 - \tilde{\beta})}{1 - \tilde{\beta}(1 - \Lambda) \varrho_z} \\ \Gamma_\zeta &= -\frac{\tilde{\beta}}{1 - \tilde{\beta}(1 - \Lambda) \varrho_\zeta} \\ \Gamma_\varsigma &= \frac{1 - \tilde{\beta} \varrho_\varsigma}{1 - \tilde{\beta}(1 - \Lambda) \varrho_\varsigma} \end{aligned}$$

E.2 Linearized first-order conditions

Linearizing the first-order conditions (D.24)-(D.28) and constraints (D.1)-(D.4) around the steady state described in Appendix D.2 yields the following:

FOC w

$$\begin{aligned} & -\gamma(1 + \Omega)\hat{y}_t + (1 + \Omega) \frac{\hat{\Sigma}_t}{\Sigma} - \left(\frac{1 - \tilde{\beta}}{\tilde{\beta}} \right) (1 + \Omega) \hat{m}_{1,t} - m_1 \left(\frac{1 - \tilde{\beta}}{\tilde{\beta}} \right)^2 \frac{\gamma \rho}{1 + \gamma \rho} (1 + \Omega)^2 \frac{\hat{w}_t}{w} \\ & - \left(\frac{1 - \tilde{\beta}}{\tilde{\beta}^2} \right) (1 + \Omega) m_1 \hat{\mu}_t + \frac{\kappa}{\gamma} \hat{m}_{2,t} - \frac{\hat{m}_{4,t}}{\gamma} + \frac{m_4}{\gamma} \frac{\hat{w}_t}{w} - \frac{m_4}{\gamma} \frac{1}{(1 + \gamma \rho)} \hat{z}_t = 0 \end{aligned} \quad (\text{E.7})$$

FOC y

$$\begin{aligned} & -\frac{\gamma \rho (1 + \Omega)}{1 + \gamma \rho} \frac{\hat{w}_t}{w} + \gamma \left[1 + 2 \frac{(1 - \Theta)^2}{\Lambda} \left(m_3 - \frac{m_1}{\beta} \right) \right] \hat{y}_t - \frac{\hat{\Sigma}_t}{\Sigma} - \hat{m}_{1,t} + \frac{\Theta}{\beta} \left(\hat{m}_{1,t-1} - m_1 \hat{\zeta}_{t-1} \right) \\ & + 2(1 - \Theta) \left(m_3 - \frac{m_1}{\beta} \right) \hat{\mu}_t + (1 - \Theta) \hat{m}_{3,t} + \frac{\hat{m}_{4,t}}{\gamma} + 2(1 - \Theta) \left(m_3 - \frac{m_1}{\beta} \right) \hat{s}_t \end{aligned} \quad (\text{E.8})$$

FOC π

$$-\hat{m}_{2,t} + \hat{m}_{2,t-1} + \Psi m_4 \pi_t = 0 \quad (\text{E.9})$$

FOC μ

$$\begin{aligned}
& - \left(\frac{1-\tilde{\beta}}{\tilde{\beta}^2} \right) \frac{\gamma\rho(1+\Omega)}{1+\gamma\rho} m_1 \frac{\hat{w}_t}{w} + \left[2\Lambda \left(m_3 - \frac{m_1}{\beta} \right) - \frac{1-\tilde{\beta}}{\tilde{\beta}^2} m_1 \right] \hat{\mu}_t + \Lambda \hat{m}_{3,t} \\
& + 2\gamma(1-\Theta) \left(m_3 - \frac{m_1}{\beta} \right) \hat{y}_t - \frac{1}{\tilde{\beta}} \left(\hat{m}_{1,t} - \frac{\tilde{\beta}}{\beta} (1-\Lambda) (\hat{m}_{1,t-1} - m_1 \hat{\zeta}_{t-1}) \right) + 2\Lambda \left(m_3 - \frac{m_1}{\beta} \right) \hat{\varsigma}_t \\
& + \mathbb{I}(t=0) \Lambda \left(\frac{\alpha}{\mu} \right)^2 m_3 \left(\frac{\vartheta}{1-\vartheta} \right) \frac{\hat{\mu}_0}{\mu} = 0
\end{aligned} \tag{E.10}$$

FOC Σ

$$\frac{\gamma\rho w}{1+\gamma\rho w} \frac{\hat{w}_t}{w} - \gamma \hat{y}_t - \hat{m}_{3,t} + \tilde{\beta} \hat{m}_{3,t+1} + \frac{1-\beta^{-1}\tilde{\beta}^2}{1-\tilde{\beta}} \frac{\hat{\Sigma}_t}{\Sigma} + \tilde{\beta} m_3 \hat{\zeta}_t = 0 \tag{E.11}$$

where $\hat{m}_i = \frac{\hat{M}_i}{\mathbb{U}^i}$ for $i \in \{1, 2, 3, 4\}$.

E.3 Deriving the target criterion

Combine the FOC for \hat{w}_t (E.7) and \hat{y}_t (E.8):

$$\begin{aligned}
& -\gamma(1+\Omega) \hat{y}_t + \Omega \frac{\hat{\Sigma}_t}{\Sigma} - \frac{\gamma\rho(1+\Omega)}{1+\gamma\rho} \left[m_1 \left(\frac{1-\tilde{\beta}}{\tilde{\beta}} \right)^2 (1+\Omega) + 1 \right] \frac{\hat{w}_t}{w} + \frac{m_4}{\gamma} \frac{\hat{w}_t}{w} + \frac{\kappa}{\gamma} \hat{m}_{2,t} - \frac{m_4}{\gamma} \frac{1}{(1+\gamma\rho)} \hat{z}_t \\
& + \frac{\Theta}{\beta} (\hat{m}_{1,t-1} - m_1 \hat{\zeta}_{t-1}) - \left[\left(\frac{1-\tilde{\beta}}{\tilde{\beta}} \right) (1+\Omega) + 1 \right] \hat{m}_{1,t} + \gamma \left[1 + 2 \frac{(1-\Theta)^2}{\Lambda} \left(m_3 - \frac{m_1}{\beta} \right) \right] \hat{y}_t \\
& - \left[\left(\frac{1-\tilde{\beta}}{\tilde{\beta}^2} \right) (1+\Omega) m_1 - 2(1-\Theta) \left(m_3 - \frac{m_1}{\beta} \right) \right] \hat{\mu}_t \\
& + 2(1-\Theta) \left(m_3 - \frac{m_1}{\beta} \right) \hat{\varsigma}_t + (1-\Theta) \hat{m}_{3,t} = 0
\end{aligned} \tag{E.12}$$

Combine with (E.10):

$$\begin{aligned}
& \gamma \left[-\Omega + 2 \frac{(1-\Theta)^2}{\Lambda} \left(m_3 - \frac{m_1}{\beta} \right) - 2(1-\Theta) \left[1 + \Omega (1-\tilde{\beta}) \right] \left(m_3 - \frac{m_1}{\beta} \right) \right] \hat{y}_t + \Omega \frac{\hat{\Sigma}_t}{\Sigma} \\
& + \left\{ m_1 \left(\frac{1-\tilde{\beta}}{\tilde{\beta}^2} \right) \frac{\gamma\rho(1+\Omega)}{1+\gamma\rho} \left[1 + \Omega (1-\tilde{\beta}) \right] - \frac{\gamma\rho(1+\Omega)}{1+\gamma\rho} \left[m_1 \left(\frac{1-\tilde{\beta}}{\tilde{\beta}} \right)^2 (1+\Omega) + 1 \right] + \frac{m_4}{\gamma} \right\} \frac{\hat{w}_t}{w} \\
& - \left\{ \left[1 + \Omega (1-\tilde{\beta}) \right] \left[2\Lambda \left(m_3 - \frac{m_1}{\beta} \right) - \frac{1-\tilde{\beta}}{\tilde{\beta}^2} m_1 \right] + \left[\left(\frac{1-\tilde{\beta}}{\tilde{\beta}^2} \right) (1+\Omega) m_1 - 2(1-\Theta) \left(m_3 - \frac{m_1}{\beta} \right) \right] \right\} \hat{\mu}_t \\
& - (1-\tilde{\beta}) \Omega \hat{m}_{3,t} + \frac{\kappa}{\gamma} \hat{m}_{2,t} - \frac{m_4}{\gamma} \frac{1}{(1+\gamma\rho)} \hat{z}_t - 2(1-\tilde{\beta}) \Omega \left(m_3 - \frac{m_1}{\beta} \right) \hat{\varsigma}_t \\
& - \mathbb{I}(t=0) \Lambda \left(\frac{\alpha}{\mu} \right)^2 \left[1 + \Omega (1-\tilde{\beta}) \right] m_3 \left(\frac{\vartheta}{1-\vartheta} \right) \frac{\hat{\mu}_0}{\mu} = 0
\end{aligned}$$

Next, use the GDP definition (19) to substitute out for $\frac{\hat{w}_t}{w}$:

$$\begin{aligned}
& \gamma \left[-\Omega + 2 \frac{(1-\Theta)^2}{\Lambda} \left(m_3 - \frac{m_1}{\beta} \right) - 2(1-\Theta) \left[1 + \Omega (1 - \tilde{\beta}) \right] \left(m_3 - \frac{m_1}{\beta} \right) \right] \hat{y}_t + \Omega \frac{\hat{\Sigma}_t}{\Sigma} \\
& + \left\{ m_1 \left(\frac{1-\tilde{\beta}}{\tilde{\beta}^2} \right) \frac{\gamma \rho (1+\Omega)}{1+\gamma \rho} \left[1 + \Omega (1 - \tilde{\beta}) \right] - \frac{\gamma \rho (1+\Omega)}{1+\gamma \rho} \left[m_1 \left(\frac{1-\tilde{\beta}}{\tilde{\beta}} \right)^2 (1+\Omega) + 1 \right] + \frac{m_4}{\gamma} \right\} \frac{1+\gamma \rho}{\rho} \hat{y}_t \\
& - \left\{ m_1 \left(\frac{1-\tilde{\beta}}{\tilde{\beta}^2} \right) \frac{\gamma \rho (1+\Omega)}{1+\gamma \rho} \left[1 + \Omega (1 - \tilde{\beta}) \right] - \frac{\gamma \rho (1+\Omega)}{1+\gamma \rho} \left[m_1 \left(\frac{1-\tilde{\beta}}{\tilde{\beta}} \right)^2 (1+\Omega) + 1 \right] + \frac{m_4}{\gamma} \right\} \frac{y}{\rho} \hat{z}_t \\
& - \left\{ \left[1 + \Omega (1 - \tilde{\beta}) \right] \left[2\Lambda \left(m_3 - \frac{m_1}{\beta} \right) - \frac{1-\tilde{\beta}}{\tilde{\beta}^2} m_1 \right] + \left[\left(\frac{1-\tilde{\beta}}{\tilde{\beta}^2} \right) (1+\Omega) m_1 - 2(1-\Theta) \left(m_3 - \frac{m_1}{\beta} \right) \right] \right\} \hat{\mu}_t \\
& - (1-\tilde{\beta}) \Omega \hat{m}_{3,t} + \frac{\kappa}{\gamma} \hat{m}_{2,t} - \frac{m_4}{\gamma} \frac{1}{(1+\gamma \rho)} \hat{z}_t - 2(1-\tilde{\beta}) \Omega \left(m_3 - \frac{m_1}{\beta} \right) \hat{\varsigma}_t \\
& - \mathbb{I}(t=0) \Lambda \left(\frac{\delta}{\mu} \right)^2 \left[1 + \Omega (1 - \tilde{\beta}) \right] m_3 \left(\frac{\vartheta}{1-\vartheta} \right) \frac{\hat{\mu}_0}{\mu} = 0 \tag{E.13}
\end{aligned}$$

Next, using (E.5) to substitute out for $\hat{\mu}_t$ and using the definitions of m_1, m_3 and m_4 , (E.13) becomes:

$$\begin{aligned}
& \gamma \Omega \left[-1 - 2 \left(\frac{1-\Theta}{\Lambda} \right) \left(\frac{1-\beta^{-1}\tilde{\beta}}{1-\beta^{-1}\tilde{\beta}(1-\Lambda)} \right) + \left[\frac{2(1-\beta^{-1}\tilde{\beta}) + \Lambda}{1-\beta^{-1}\tilde{\beta}(1-\Lambda)} \right] \frac{\Theta}{1-\Lambda} \right] \hat{y}_t + \Omega \frac{\hat{\Sigma}_t}{\Sigma} \\
& + \frac{(1+\Omega)(1-\Lambda)(1-\beta^{-1}\tilde{\beta})}{\rho [1-\beta^{-1}\tilde{\beta}(1-\Lambda)]} \left[\hat{y}_t - \left(\frac{y+\rho}{1+\gamma \rho} \right) \hat{z}_t \right] \\
& - \Omega \left[\frac{2(1-\beta^{-1}\tilde{\beta}) + \Lambda}{1-\beta^{-1}\tilde{\beta}(1-\Lambda)} \right] \Gamma_t - (1-\tilde{\beta}) \Omega \hat{m}_{3,t} + \frac{\kappa}{\gamma} \hat{m}_{2,t} - 2(1-\tilde{\beta}) \Omega \left(m_3 - \frac{m_1}{\beta} \right) \hat{\varsigma}_t \\
& + \mathbb{I}(t=0) \gamma \Lambda \left(\frac{\alpha}{\mu} \right)^2 \left[1 + \Omega (1 - \tilde{\beta}) \right]^2 m_3 \left(\frac{\vartheta}{1-\vartheta} \right) \frac{1}{\mu} \hat{y}_0 \\
& - \mathbb{I}(t=0) \Lambda \left(\frac{\alpha}{\mu} \right)^2 \left[1 + \Omega (1 - \tilde{\beta}) \right] m_3 \left(\frac{\vartheta}{1-\vartheta} \right) \frac{1}{\mu} \Gamma_0 = 0 \tag{E.14}
\end{aligned}$$

Guess that

$$\hat{m}_{3,t} = \frac{1}{1-\tilde{\beta}} \frac{\hat{\Sigma}_t}{\Sigma} + \gamma \Omega \hat{y}_t + a_z \hat{z}_t + a_\zeta \hat{\zeta}_t + a_\varsigma \hat{\varsigma}_t \tag{E.15}$$

and use this in (E.11) with $\frac{\hat{w}_t}{w}$ substituted out using the definition of GDP:

$$\begin{aligned}
& \gamma (1-\tilde{\beta}) \Omega \hat{y}_{t+1} - \left(\frac{1-\tilde{\beta}}{\tilde{\beta}} \right) \frac{\gamma y (1+\Omega)}{1+\gamma \rho} \hat{z}_t + \frac{\hat{\Sigma}_{t+1}}{\Sigma} - \beta^{-1} \tilde{\beta} \frac{\hat{\Sigma}_t}{\Sigma} + \left(\frac{1-\tilde{\beta}}{\tilde{\beta}} \right) (\tilde{\beta} \varrho_z - 1) a_z \hat{z}_t \\
& + \left(\frac{1-\tilde{\beta}}{\tilde{\beta}} \right) (\tilde{\beta} \varrho_\zeta - 1) a_\zeta \hat{\zeta}_t + \left(\frac{1-\tilde{\beta}}{\tilde{\beta}} \right) (\tilde{\beta} \varrho_\varsigma - 1) a_\varsigma \hat{\varsigma}_t + \left(\frac{1-\tilde{\beta}}{\tilde{\beta}} \right) \tilde{\beta} m_3 \hat{\zeta}_t = 0,
\end{aligned}$$

using the fact that $\widehat{z}_{t+1} = \varrho_z \widehat{z}_t$, $\widehat{\zeta}_{t+1} = \varrho_\zeta \widehat{\zeta}_t$ and $\widehat{\varsigma}_{t+1} = \varrho_\varsigma \widehat{\varsigma}_t$. Using the expression for Σ_{t+1} (25) in the equation above:

$$\begin{aligned} & \left[\left(\frac{1-\widetilde{\beta}}{\widetilde{\beta}} \right) (\widetilde{\beta}\varrho_z - 1) a_z - \left(\frac{1-\widetilde{\beta}}{\widetilde{\beta}} \right) \frac{\gamma y (1+\Omega)}{1+\gamma\rho} + \Lambda \frac{\gamma\rho}{1+\gamma\rho} \frac{y (1+\Omega) (1-\widetilde{\beta})}{\rho (1-\widetilde{\beta})(1-\Lambda)} \varrho_z \right] \widehat{z}_t \\ & \quad + \left[\left(\frac{1-\widetilde{\beta}}{\widetilde{\beta}} \right) (\widetilde{\beta}\varrho_\zeta - 1) a_\zeta + 1 - \frac{\widetilde{\beta}\varrho_\zeta\Lambda}{1-\widetilde{\beta}(1-\Lambda)\varrho_\zeta} \right] \widehat{\zeta}_t \\ & \quad + \left[\left(\frac{1-\widetilde{\beta}}{\widetilde{\beta}} \right) (\widetilde{\beta}\varrho_\varsigma - 1) a_\varsigma - \widetilde{\beta} \frac{\varrho_\varsigma^2 \Lambda^2}{1-\widetilde{\beta}(1-\Lambda)\varrho_\varsigma} + \varrho_\varsigma \Lambda \right] \widehat{\varsigma}_t = 0 \end{aligned}$$

which implies that a_z , a_ζ and a_ς must satisfy:

$$a_z = -\frac{\gamma y}{1+\gamma\rho} \left[\frac{1+\Omega}{1-\widetilde{\beta}\varrho_z(1-\Lambda)} \right] \quad (\text{E.16})$$

$$a_\zeta = \frac{\widetilde{\beta}}{1-\widetilde{\beta}} \left[\frac{1}{1-\widetilde{\beta}(1-\Lambda)\varrho_\zeta} \right] \quad (\text{E.17})$$

$$a_\varsigma = \frac{1}{1-\widetilde{\beta}} \left[\frac{\widetilde{\beta}\varrho_\varsigma\Lambda}{1-\widetilde{\beta}(1-\Lambda)\varrho_\varsigma} \right] \quad (\text{E.18})$$

Using the expression (E.15) for $\widehat{m}_{3,t}$ in (E.14):

$$\Upsilon_0(\Omega)\widetilde{x}_0 = -\frac{\rho\kappa}{m_4}\widehat{m}_{2,0} \quad \text{for } t=0 \quad (\text{E.19})$$

$$\Upsilon(\Omega)x_t = -\frac{\rho\kappa}{m_4}\widehat{m}_{2,t} \quad \text{for } t>0 \quad (\text{E.20})$$

where

$$\widetilde{x}_0 = \widehat{y}_0 - \delta_0(\Omega) \frac{y+\rho}{1+\gamma\rho} \widehat{z}_0 + \chi_0(\Omega)\widehat{\zeta}_0 - \Xi_0(\Omega)\widehat{\varsigma}_0 \quad (\text{E.21})$$

$$x_t = \widehat{y}_t - \delta(\Omega) \frac{y+\rho}{1+\gamma\rho} \widehat{z}_t + \chi(\Omega)\widehat{\zeta}_t - \Xi(\Omega)\widehat{\varsigma}_t \quad (\text{E.22})$$

$$\Upsilon(\Omega) = 1 + \frac{\gamma\rho\Omega}{1+\Omega} \left\{ (1-\widetilde{\beta}) \Omega \left(\frac{2}{\Lambda(1-\Lambda)} - 1 \right) - 1 \right\} \quad (\text{E.23})$$

$$\delta(\Omega) = \frac{1}{\Upsilon(\Omega)} \left[1 + \left(\frac{1+\Lambda}{1-\Lambda} \right) \left(\frac{\gamma(1-\widetilde{\beta})\Omega}{1-\widetilde{\beta}\varrho_z(1-\Lambda)} \right) \frac{\rho y}{y+\rho} \right] \quad (\text{E.24})$$

$$\chi(\Omega) = \frac{1}{\Upsilon(\Omega)} \frac{\Omega}{1+\Omega} \left(\frac{1+\Lambda}{1-\Lambda} \right) \left[\frac{\widetilde{\beta}\rho}{1-\widetilde{\beta}\varrho_\zeta(1-\Lambda)} \right] \quad (\text{E.25})$$

$$\Xi(\Omega) = \frac{1}{\Upsilon(\Omega)} \frac{\rho\Omega}{1+\Omega} \left[2 - \left(\frac{1+\Lambda}{1-\Lambda} \right) \frac{\widetilde{\beta}\Lambda\varrho_\varsigma}{1-\widetilde{\beta}\varrho_\varsigma(1-\Lambda)} \right] \quad (\text{E.26})$$

$$\Upsilon_0(\Omega) = \Upsilon(\Omega) + \gamma \left[1 + \Omega \left(1 - \tilde{\beta} \right) \right] \mathcal{G} \quad (\text{E.27})$$

$$\delta_0(\Omega) = \frac{\Upsilon(\Omega)}{\Upsilon_0(\Omega)} \delta(\Omega) - \frac{\mathcal{G}}{\Upsilon_0(\Omega)} \frac{y}{y + \rho} \frac{(1 + \Omega) \left(1 - \tilde{\beta} \right)}{1 - \tilde{\beta} (1 - \Lambda)} \varrho_z \quad (\text{E.28})$$

$$\Xi_0(\Omega) = \frac{\Upsilon(\Omega)}{\Upsilon_0(\Omega)} \Xi(\Omega) - \frac{\mathcal{G}}{\Upsilon_0(\Omega)} \frac{\tilde{\beta} \varrho_c \Lambda}{1 - \tilde{\beta} (1 - \Lambda)} \varrho_c \quad (\text{E.29})$$

$$\chi_0(\Omega) = \frac{\Upsilon(\Omega)}{\Upsilon_0(\Omega)} \chi(\Omega) + \frac{\mathcal{G}}{\Upsilon_0(\Omega)} \frac{\tilde{\beta}}{1 - \tilde{\beta} (1 - \Lambda)} \varrho_c, \quad (\text{E.30})$$

where $\mathcal{G} = \frac{\gamma \rho}{m_4} \Lambda \left(\frac{\alpha}{\mu} \right)^2 \left[1 + \Omega \left(1 - \tilde{\beta} \right) \right] m_3 \left(\frac{\vartheta}{1 - \vartheta} \right) \frac{1}{\mu}$. Note that in the baseline with the utilitarian planner ($\alpha = 0$), we have $\mathcal{G} = 0$ and $\Upsilon_0(\Omega) = \Upsilon(\Omega)$, $\delta_0(\Omega) = \delta(\Omega)$, $\Xi_0(\Omega) = \Xi(\Omega)$ and $\chi(\Omega) = \chi_0(\Omega)$. Since $\hat{m}_{2,-1} = 0$, (E.9) implies that $\hat{m}_{2,t} = \Psi m_4 \hat{p}_t$, where $\pi_t = \hat{p}_t - \hat{p}_{t-1}$. Combining this with (E.19)-(E.20) and using the definition of κ yields

$$\begin{aligned} \Upsilon_0(\Omega) \tilde{x}_0 + (1 + \gamma \rho) \frac{\lambda}{\lambda - 1} \hat{p}_0 &= 0 \quad \text{for } t = 0 \\ \Upsilon(\Omega) x_t + (1 + \gamma \rho) \frac{\lambda}{\lambda - 1} \hat{p}_t &= 0 \quad \text{for } t > 0 \end{aligned} \quad (\text{E.31})$$

This general target criterion can be specialized to yield the target criterion in Proposition 3 for the utilitarian planner (setting $\alpha = 0$ and $\hat{\zeta}_t = \hat{\varsigma}_t = 0$), Proposition 7 for the non-utilitarian planner (again, setting $\hat{\zeta}_t = \hat{\varsigma}_t = 0$) and Proposition 9 for demand shocks (setting $\alpha = 0$ and $\hat{z}_t = \hat{\lambda}_t = 0$).

Claim 1. $\Upsilon(\Omega) > 1$ with countercyclical risk

Proof.

$$\begin{aligned} \Upsilon(\Omega) &= 1 + \frac{\rho \gamma \Omega}{1 + \Omega} \left[\left(\frac{2}{\Lambda (1 - \Lambda)} - 1 \right) \left(1 - \tilde{\beta} \right) \Omega - 1 \right] \\ &> 1 + \frac{\rho \gamma \Omega}{1 + \Omega} \left[\left(\frac{2}{\Lambda (1 - \Lambda)} - 1 \right) \left(1 - \tilde{\beta} \right) \frac{\Lambda}{(1 - \tilde{\beta}) (1 - \Lambda)} - 1 \right], \end{aligned}$$

where we have used the fact that $\Omega = \frac{\Theta - 1 + \Lambda}{(1 - \tilde{\beta})(1 - \Lambda)}$ and for countercyclical risk ($\Theta > 1$), we have $\Omega > \frac{\Lambda}{(1 - \tilde{\beta})(1 - \Lambda)}$. Then, the above can be simplified to

$$\Upsilon(\Omega) > 1 + \frac{\rho \gamma \Omega}{1 + \Omega} \frac{1 + \Lambda}{(1 - \Lambda)^2} > 1$$

□

Claim 2. $0 < \delta(\Omega) < 1$ with countercyclical risk

Proof. Using the expression for $\Upsilon(\Omega)$ in $\delta(\Omega)$, we have:

$$\delta(\Omega) = \frac{1 + \Omega + (\Omega + \Omega^2) \frac{\gamma(1-\tilde{\beta})}{1-\tilde{\beta}\rho_z(1-\Lambda)} \frac{\rho y}{\rho+y} \left[\frac{1+\Lambda}{1-\Lambda} \right]}{1 + (1-\rho\gamma)\Omega + \rho\gamma\Omega^2 \left(\frac{2}{\Lambda(1-\Lambda)} - 1 \right) (1-\tilde{\beta})} \quad (\text{E.32})$$

We need to show that $\delta(\Omega) < 1$, i.e.,

$$1 + \Omega + (\Omega + \Omega^2) \frac{\gamma(1-\tilde{\beta})}{1-\tilde{\beta}\rho_z(1-\Lambda)} \frac{\rho y}{\rho+y} \left[\frac{1+\Lambda}{1-\Lambda} \right] < 1 + (1-\rho\gamma)\Omega + \rho\gamma\Omega^2 \left(\frac{2}{\Lambda(1-\Lambda)} - 1 \right) (1-\tilde{\beta})$$

This expression can be simplified to yield

$$1 + \frac{(1-\tilde{\beta})}{1-\tilde{\beta}\rho_z(1-\Lambda)} \frac{y}{\rho+y} \left(\frac{1+\Lambda}{1-\Lambda} \right) < \Omega(1-\tilde{\beta}) \left[\left(\frac{2}{\Lambda(1-\Lambda)} - 1 \right) - \frac{1}{1-\tilde{\beta}\rho_z(1-\Lambda)} \frac{y}{\rho+y} \left(\frac{1+\Lambda}{1-\Lambda} \right) \right] \quad (\text{E.33})$$

First, we show that the term in the square brackets on the RHS of (E.33) is positive, i.e.,

$$2 > \Lambda \left[1 - \Lambda + \frac{1+\Lambda}{1-\tilde{\beta}\rho_z(1-\Lambda)} \frac{y}{\rho+y} \right]$$

The worst case for this to be true is if y were very large and $\rho_z = 1$. In that case, for the expression above to be true, it must be that

$$\tilde{\beta} < \frac{2}{2 - (1-\Lambda)\Lambda},$$

which is true since $\tilde{\beta} < 1$ and $\frac{2}{2-(1-\Lambda)\Lambda} > 1$ since we know that $0 < \Lambda < 1$ from Appendix C. Thus, the term in the square brackets on the RHS of (E.33) is positive. Next, to show that (E.33) holds with countercyclical risk, it suffices to show that it holds for the lowest Ω consistent with non-procyclical risk, i.e. $\Omega = \frac{\Lambda}{(1-\tilde{\beta})(1-\Lambda)}$. Plug in $\Omega = \frac{\Lambda}{(1-\tilde{\beta})(1-\Lambda)}$ into (E.33), i.e.,

$$1 + \frac{(1-\tilde{\beta})(1+\Lambda)}{1-\tilde{\beta}\rho_z(1-\Lambda)} \frac{y}{\rho+y} < \left[\Lambda \left(\frac{2}{\Lambda(1-\Lambda)} - 1 \right) - \frac{1+\Lambda}{1-\tilde{\beta}\rho_z(1-\Lambda)} \frac{y}{\rho+y} \left(\frac{\Lambda}{1-\Lambda} \right) \right]$$

Again the worst case for this condition to be satisfied is if $\rho_z = 1$. Suppose that is the case. Then, the expression can be further simplified to

$$\frac{y}{\rho+y} < 1,$$

which is true since steady state output is positive. \square

Claim 3. $\chi(\Omega) > 0$ with countercyclical risk

Proof. It is clear from the expression for $\chi(\Omega)$ that for countercyclical risk $\Omega \geq \Omega^c > 0$, $\chi(\Omega) > 0$. \square

Claim 4. $\Xi(\Omega) > 0$ with countercyclical risk.

Proof. It suffices to show that $2 - \left(\frac{1+\Lambda}{1-\Lambda}\right) \frac{\tilde{\beta}\Lambda\varrho_\varsigma}{1-\tilde{\beta}(1-\Lambda)\varrho_\varsigma} > 0$. This is clearly true for $\varrho_z = 0$ and as long as $\varrho_\varsigma < \frac{2(1-\Lambda)}{\tilde{\beta}[2-3\Lambda(1-\Lambda)]} = \bar{\varrho}_\varsigma$, the claim is true. Thus a sufficient condition for $\Xi(\Omega) > 0$ is for $\varrho_\varsigma \in [0, \bar{\varrho}_\varsigma]$. \square

Claim 5. $\Upsilon(0) = \delta(0) = 1$ and $\chi(0) = \Xi(0) = 0$ when $\alpha = 0$.

Proof. True by inspection of equations (E.23), (E.24), (E.25) and (E.26). \square

Claim 6. $\Upsilon_0(\Omega) > \Upsilon(\Omega)$ and $\delta_0(\Omega) < \delta(\Omega)$ when $\alpha \neq 0$

Proof. The first claim that $\Upsilon_0 > \Upsilon$ is true by inspection of equations (E.27), since $\mathcal{G} > 0$ for $\Omega \geq \Omega^c > 0$. For the second, rewrite (E.28) as

$$\delta_0(\Omega) = \frac{\Upsilon(\Omega)}{\Upsilon_0(\Omega)}\delta(\Omega) - \frac{\mathcal{G}}{\Upsilon_0(\Omega)} \frac{y}{y+\rho} \frac{(1+\Omega)(1-\tilde{\beta})}{1-\tilde{\beta}(1-\Lambda)\varrho_z}$$

Since $0 < \Upsilon(\Omega)/\Upsilon_0(\Omega) < 1$ and $\mathcal{G} > 0$ for $\Omega \geq \Omega^c > 0$, it follows that $\delta_0(\Omega) < \delta(\Omega)$ when $\alpha \neq 0$. \square

Claim 7. $\Upsilon_0(\Omega)$ is increasing in α and $\delta_0(\Omega)$ is decreasing in α for $\alpha > 0$.

Proof. Substituting the definition of \mathcal{G} into (E.27):

$$\Upsilon_0(\Omega) = \Upsilon(\Omega) + \gamma \left[1 + \Omega(1-\tilde{\beta})\right]^2 \frac{\gamma\rho}{m_4} \Lambda \left(\frac{\alpha}{\mu}\right)^2 \left(\frac{\vartheta}{1-\vartheta}\right) \frac{m_3}{\mu}$$

This is clearly increasing in α for $\alpha > 0$. Dividing (E.27) by $\Upsilon_0(\Omega)$ and substituting the expression for $\Upsilon(\Omega)/\Upsilon_0(\Omega)$ in to (E.28) and rearranging:

$$\delta_0(\Omega) = \delta(\Omega) - \left[\gamma \left[1 + \Omega(1-\tilde{\beta})\right] \delta(\Omega) + \frac{y}{y+\rho} \frac{(1+\Omega)(1-\tilde{\beta})}{1-\tilde{\beta}(1-\Lambda)\varrho_z} \right] \frac{\mathcal{G}}{\Upsilon_0(\Omega)}$$

This expression is clearly decreasing in $\mathcal{G}/\Upsilon_0(\Omega)$ which in turn is increasing in \mathcal{G} and, hence, increasing in α for $\alpha > 0$. Hence $\delta_0(\Omega)$ is decreasing in α . \square

F Optimal Dynamics

As shown in Appendix E.3, the dynamics of x_t and π_t are given by the target criterion (27)

$$x_t - x_{t-1} + \varepsilon\pi_t = 0 \tag{F.1}$$

and the Phillips curve

$$\pi_t = \beta\pi_{t+1} + \kappa \left(x_t - [1 - \delta(\Omega)] \frac{y+\rho}{1+\gamma\rho} \hat{z}_t - \chi(\Omega)\hat{\zeta}_t + \Xi(\Omega)\hat{\varsigma}_t + \frac{\rho}{1+\gamma\rho} \hat{\lambda}_t \right),$$

where we have used the definition of x_t from equation (E.22) and defined $\varepsilon = \frac{\lambda}{\lambda-1} \frac{1+\gamma\rho}{\Upsilon(\Omega)}$. Substituting the target criterion into the Phillips curve, we get a second-order difference equation

$$x_{t+1} - \left[1 + \frac{\kappa\varepsilon + 1}{\beta}\right] x_t + \frac{1}{\beta} x_{t-1} = \frac{\varepsilon\kappa}{\beta} \left[-[1 - \delta(\Omega)] \frac{y + \rho}{1 + \gamma\rho} \widehat{z}_t - \chi(\Omega) \widehat{\zeta}_t + \Xi(\Omega) \widehat{\varsigma}_t + \frac{\rho}{1 + \gamma\rho} \widehat{\lambda}_t \right]$$

The solution to this system has the form:

$$x_t = \mathcal{A}_x x_{t-1} + \mathcal{A}_z \widehat{z}_t + \mathcal{A}_\zeta \widehat{\zeta}_t + \mathcal{A}_\lambda \widehat{\lambda}_t + \mathcal{A}_\varsigma \widehat{\varsigma}_t \quad (\text{F.2})$$

$$\pi_t = \mathcal{B}_x x_{t-1} + \mathcal{B}_z \widehat{z}_t + \mathcal{B}_\zeta \widehat{\zeta}_t + \mathcal{B}_\lambda \widehat{\lambda}_t + \mathcal{B}_\varsigma \widehat{\varsigma}_t \quad (\text{F.3})$$

Using the method of undetermined coefficients, it is straightforward to see that \mathcal{A}_x satisfies the characteristic polynomial

$$\mathcal{P}(\mathcal{A}_x) = \mathcal{A}_x^2 - \left[1 + \frac{\kappa\varepsilon + 1}{\beta}\right] \mathcal{A}_x + \frac{1}{\beta} = 0 \quad (\text{F.4})$$

We know that $\mathcal{P}(0) = \beta^{-1} > 0$ and $\mathcal{P}(1) = -\beta^{-1}\kappa\varepsilon < 0$. Thus, we have $\mathcal{A}_x \in (0, 1)$ and the coefficients can be written as

$$\mathcal{A}_x = \frac{1}{2} \left(1 + \frac{\kappa\varepsilon + 1}{\beta} - \sqrt{\left[1 + \frac{\kappa\varepsilon + 1}{\beta}\right]^2 - \frac{4}{\beta}} \right) \in (0, 1) \quad (\text{F.5})$$

$$\mathcal{A}_z = \frac{\kappa\beta^{-1}\varepsilon}{\kappa\beta^{-1}\varepsilon + (1 - \mathcal{A}_x) + \left(\frac{1}{\beta} - \varrho_z\right)} [1 - \delta(\Omega)] \frac{y + \rho}{1 + \gamma\rho} \quad (\text{F.6})$$

$$\mathcal{A}_\zeta = \frac{\kappa\beta^{-1}\varepsilon}{\kappa\beta^{-1}\varepsilon + (1 - \mathcal{A}_x) + \left(\frac{1}{\beta} - \varrho_\zeta\right)} \chi(\Omega) \quad (\text{F.7})$$

$$\mathcal{A}_\lambda = -\frac{\kappa\beta^{-1}\varepsilon}{\kappa\beta^{-1}\varepsilon + (1 - \mathcal{A}_x) + \left(\frac{1}{\beta} - \varrho_\lambda\right)} \frac{\rho}{1 + \gamma\rho} \quad (\text{F.8})$$

$$\mathcal{A}_\varsigma = -\frac{\kappa\beta^{-1}\varepsilon}{\kappa\beta^{-1}\varepsilon + (1 - \mathcal{A}_x) + \left(\frac{1}{\beta} - \varrho_\varsigma\right)} \Xi(\Omega) \quad (\text{F.9})$$

$$\mathcal{B}_x = \frac{1 - \mathcal{A}_x}{\varepsilon} \quad (\text{F.10})$$

$$\mathcal{B}_i = -\frac{1}{\varepsilon} \mathcal{A}_i \quad \text{for } i \in \{z, \zeta, \lambda, \varsigma\} \quad (\text{F.11})$$

Claim 8. *The following statements are true:*

1. $\frac{\kappa\beta^{-1}\varepsilon}{\kappa\beta^{-1}\varepsilon + (1 - \mathcal{A}_x) + \left(\frac{1}{\beta} - \varrho_i\right)} \in (0, 1)$ for $i \in \{z, \zeta, \lambda, \varsigma\}$

2. $\mathcal{B}_x > 0$

Proof. To see that $\frac{\kappa\beta^{-1}\varepsilon}{\kappa\beta^{-1}\varepsilon + (1 - \mathcal{A}_x) + \left(\frac{1}{\beta} - \varrho_i\right)} \in (0, 1)$ for $i \in \{z, \zeta, \lambda, \varsigma\}$, notice that since $\mathcal{A}_x \in (0, 1)$, we have $1 - \mathcal{A}_x \geq 0$. Furthermore, since $\beta < 1$, so $\beta^{-1} - \varrho_i > 0$ for $i \in \{z, \zeta, \lambda, \varsigma\}$ as long as $\varrho_i \in (0, 1)$, which is

a maintained assumption. Again, since $\mathcal{A}_x \in (0, 1)$, it is immediate that $\mathcal{B}_x > 0$. It follows that $\mathcal{A}_z > 0$, $\mathcal{A}_\zeta > 0$, and $\mathcal{A}_\lambda < 0$. □

F.1 Proof of Propositions 4, 5, 10 and 11

F.1.1 Impact effects following a productivity shock

Since $x_{-1} = 0$, it follows from equations (F.2) and (F.3) that the impact effect of a productivity shock is:

$$\frac{\partial x_0}{\partial \widehat{z}_0} = \mathcal{A}_z > 0 \quad \text{and} \quad \frac{\partial \pi_0}{\partial \widehat{z}_0} = \mathcal{B}_z < 0$$

Using $\widehat{y}_t = x_t + \delta(\Omega) \frac{y+\rho}{1+\gamma\rho} \widehat{z}_t$, we have:

$$\frac{\partial \widehat{y}_0}{\partial \widehat{z}_0} = \mathcal{A}_z + \delta(\Omega) \frac{y+\rho}{1+\gamma\rho} = \frac{y+\rho}{1+\gamma\rho} \left(\delta(\Omega) + [1-\delta(\Omega)] \frac{\kappa\beta^{-1}\varepsilon}{\kappa\beta^{-1}\varepsilon + (1-\mathcal{A}_x) + \left(\frac{1}{\beta} - \varrho_z\right)} \right) \in \left(0, \frac{y+\rho}{1+\gamma\rho}\right),$$

where we have used the fact that $\delta(\Omega) \in (0, 1)$ for $\Omega \geq \Omega^c$. In other words, \widehat{y}_0 falls less than $\widehat{y}_0^n = \frac{y+\rho}{1+\gamma\rho} \widehat{z}_0$ for $\widehat{z}_0 < 0$.

F.1.2 Impact effects following a markup shock

Since $x_{-1} = 0$ and $\widehat{y}_t = x_t$ (since all shocks other than markup shocks are 0 in this case), it follows immediately from equations (F.2) and (F.3) that:

$$\begin{aligned} \frac{\partial \widehat{y}_0}{\partial \widehat{\lambda}_0} &= \mathcal{A}_\lambda < 0 \\ \frac{\partial \pi_0}{\partial \widehat{\lambda}_0} &= \mathcal{B}_\lambda > 0 \end{aligned}$$

Following a markup shock, the dynamics of π_t and x_t are described by the same equations (F.1) and (F) except that $\varepsilon(\Omega) = \frac{\lambda}{\lambda-1} \frac{1+\gamma\rho}{\Upsilon(\Omega)}$ is smaller in HANK since $\Upsilon(\Omega) > 1$ while $\Upsilon = 1$ in RANK. Thus, to show that in HANK, output decreases less and inflation increases more following a positive markup shock, it suffices to show that $\frac{\partial \mathcal{A}_\lambda}{\partial \varepsilon} < 0$ (output falls less on impact when ε is lower) and $\frac{\partial \mathcal{B}_\lambda}{\partial \varepsilon} < 0$ (inflation increases more on impact when ε is lower). We have

$$\begin{aligned} \mathcal{A}_\lambda &= -\frac{\kappa\beta^{-1}\varepsilon}{\kappa\beta^{-1}\varepsilon + (1-\mathcal{A}_x) + \left(\frac{1}{\beta} - \varrho_\lambda\right)} \frac{\rho}{1+\gamma\rho} \\ &= -\frac{2\beta^{-1}\kappa}{\sqrt{\frac{[1+\beta^{-1}(\kappa\varepsilon+1)]^2 - 4\beta^{-1}}{\varepsilon^2}} + \left(\frac{1-\varrho_\lambda}{\varepsilon}\right) + \frac{\beta^{-1}-\varrho_\lambda}{\varepsilon} + \beta^{-1}\kappa} \frac{\rho}{1+\gamma\rho} \end{aligned}$$

where we have plugged in the expression for \mathcal{A}_x from (F.5) in the second line. Since

$$\frac{\partial}{\partial \varepsilon} \left[\frac{[1 + \beta^{-1}(\kappa\varepsilon + 1)]^2 - 4\beta^{-1}}{\varepsilon^2} \right] = -2 \left[\frac{[1 + \beta^{-1}(\kappa\varepsilon + 1)](1 + \beta^{-1}) + 4\beta^{-1}}{\varepsilon^3} \right] < 0$$

it is clear that the denominator of \mathcal{A}_λ is decreasing in ε . Since the numerator is negative, it follows that \mathcal{A}_λ is decreasing in ε . We also know that $\mathcal{B}_\lambda = -\frac{1}{\varepsilon}\mathcal{A}_\lambda$ which implies

$$\mathcal{B}_\lambda = \frac{2\beta^{-1}\kappa}{\sqrt{[1 + \beta^{-1}(\kappa\varepsilon + 1)]^2 - 4\beta^{-1} + (1 - \varrho_\lambda) + (\beta^{-1} - \varrho_\lambda) + \beta^{-1}\kappa\varepsilon}} \frac{\rho}{1 + \gamma\rho}$$

Clearly, the denominator is increasing in ε so \mathcal{B}_λ is decreasing in ε .

F.1.3 Impact effects following a discount-factor shock

Since $x_{-1} = 0$ and $y_t = x_t - \chi(\Omega)\widehat{\zeta}_t$, the response of \widehat{y}_0 to $\widehat{\zeta}_0$ is

$$\frac{d\widehat{y}_0}{d\widehat{\zeta}_0} = \mathcal{A}_\zeta - \chi(\Omega) = - \left[1 - \frac{\kappa\beta^{-1}\varepsilon}{\kappa\beta^{-1}\varepsilon + (1 - \mathcal{A}_x) + \left(\frac{1}{\beta} - \varrho_\zeta\right)} \right] \chi(\Omega) < 0$$

while the impact response of π_0 is given by $\mathcal{B}_\zeta = -\frac{1}{\varepsilon}\mathcal{A}_\zeta < 0$.

F.2 Impact effects following a risk shock

Since $x_{-1} = 0$ and $\widehat{y}_t = x_t + \Xi(\Omega)\widehat{\varsigma}_t$, the response of \widehat{y}_0 to $\widehat{\varsigma}_0$ is

$$\frac{d\widehat{y}_0}{d\widehat{\varsigma}_0} = \mathcal{A}_\varsigma + \Xi(\Omega) = \left[1 - \frac{\kappa\beta^{-1}\varepsilon}{\kappa\beta^{-1}\varepsilon + (1 - \mathcal{A}_x) + \left(\frac{1}{\beta} - \varrho_\varsigma\right)} \right] \Xi(\Omega) > 0,$$

for $\Omega \geq \Omega^c$ and ϱ_ς not too large. Similarly, the impact response of π_0 is given by $\mathcal{B}_\varsigma = -\frac{1}{\varepsilon}\mathcal{A}_\varsigma > 0$.

F.3 Response of $\widehat{y}_t - \widehat{y}_t^n$ and π_t for large t

The following lemma characterizes the behavior of $\widehat{y}_t - \widehat{y}_t^n$ and π_t following a generic shock \mathcal{S}_0 , where $\mathcal{S}_0 \in \{\widehat{z}_0, \widehat{\lambda}_0, \widehat{\zeta}_0, \widehat{\varsigma}_0\}$ for large t . In doing so, the lemma provides a proof of the claims made in Propositions 4, 5, 10 and 11 about long-run behavior of $\widehat{y}_t - \widehat{y}_t^n$ and π_t .

Lemma 5. *After any date 0 shock \mathcal{S}_0 , where $\mathcal{S}_0 \in \{\widehat{z}_0, \widehat{\lambda}_0, \widehat{\zeta}_0, \widehat{\varsigma}_0\}$, for large enough t ,*

$$\text{sign} \left(\frac{\partial \pi_t}{\partial \mathcal{S}_0} \right) = \text{sign} \left(\frac{\partial (\widehat{y}_t - \widehat{y}_t^n)}{\partial \mathcal{S}_0} \right) = -1 \times \text{sign} \left(\frac{\partial \pi_0}{\partial \mathcal{S}_0} \right)$$

Proof. We know that $\frac{\partial \pi_0}{\partial \mathcal{S}_0} = \mathcal{B}_\mathcal{S} = -\frac{1}{\varepsilon}\mathcal{A}_\mathcal{S}$. Thus, we need to show that for large enough t , π_t and $\widehat{y}_t - \widehat{y}_t^n$ have the same sign as $\mathcal{A}_\mathcal{S} \times \mathcal{S}_0$. The dynamics of x_t and π_t in response to a shock \mathcal{S}_0 are given by the system of two equations

$$x_t = \mathcal{A}_x x_{t-1} + \mathcal{A}_\mathcal{S} \mathcal{S}_t \quad \text{and} \quad \mathcal{S}_t = \varrho_\mathcal{S} \mathcal{S}_{t-1},$$

with \mathcal{S}_0 given. The solution of this system is given by

$$x_t = \mathcal{A}_\mathcal{S} \frac{\varrho_\mathcal{S}^{t+1} - \mathcal{A}_x^{t+1}}{\varrho_\mathcal{S} - \mathcal{A}_x} \mathcal{S}_0$$

as long as $\varrho_S \neq \mathcal{A}_x$. Using this in (27), the dynamics of inflation can then be written as

$$\pi_t = -\frac{\mathcal{A}_S}{\varepsilon} \left(\frac{\varrho_S^{t+1} - \mathcal{A}_x^{t+1}}{\varrho_S - \mathcal{A}_x} - \frac{\varrho_S^t - \mathcal{A}_x^t}{\varrho_S - \mathcal{A}_x} \right) \mathcal{S}_0,$$

where $\varepsilon > 0$, $\mathcal{A}_S > 0$ and $0 < \mathcal{A}_x < 1$ are defined in Appendix F. For large enough $t > 0$, the dynamics of x_t and π_t are governed by the dominant eigenvalue $\max\{\mathcal{A}_x, \varrho_S\}$. If $\varrho_S < \mathcal{A}_S$, dividing the expression for π_t above by \mathcal{A}_x^t and taking the limit $t \rightarrow \infty$, we have

$$\lim_{t \rightarrow \infty} \mathcal{A}_x^{-t} \pi_t = \frac{1}{\varepsilon} \left(\frac{\mathcal{A}_S}{\mathcal{A}_x - \varrho_S} \right) \underbrace{(1 - \mathcal{A}_x)}_{>0} \mathcal{S}_0,$$

which has the same sign as $\mathcal{A}_S \times \mathcal{S}_0$. Similarly, dividing the Phillips curve by \mathcal{A}_x^t and taking limits as $t \rightarrow \infty$ yields

$$\lim_{t \rightarrow \infty} (1 - \beta \mathcal{A}_x) \frac{\pi_t}{\mathcal{A}_x^t} = \kappa \lim_{t \rightarrow \infty} \left(\frac{\hat{y}_t - \hat{y}_t^n}{\mathcal{A}_x^t} \right)$$

This implies that $\hat{y}_t - \hat{y}_t^n$ has the same sign as $\mathcal{A}_S \times \mathcal{S}_0$ for large t . Instead if $\varrho_S > \mathcal{A}_x$, dividing by ϱ_S^t and taking limits, we have:

$$\lim_{t \rightarrow \infty} \varrho_S^{-t} \pi_t = \frac{1}{\varepsilon} \underbrace{(1 - \varrho_S)}_{>0} \frac{\mathcal{A}_S}{\varrho_S - \mathcal{A}_x} \mathcal{S}_0 \quad \text{and} \quad \lim_{t \rightarrow \infty} (1 - \beta \varrho_S) \frac{\pi_t}{\varrho_S^t} = \kappa \lim_{t \rightarrow \infty} \left(\frac{\hat{y}_t - \hat{y}_t^n}{\varrho_S^t} \right)$$

This implies that again, π_t and $\hat{y}_t - \hat{y}_t^n$ have the same sign as $\mathcal{A}_S \times \mathcal{S}_0$ for large t . Finally, in the special case where both eigenvalues are identical, $\mathcal{A}_x = \varrho_S$, the solution for x_t is instead given by:

$$x_t = (t + 1) \mathcal{A}_S \varrho_S^t \mathcal{S}_0$$

and so the target criterion implies that the path of inflation can be written as

$$\pi_t = -\frac{\mathcal{A}_S}{\varepsilon} ((t + 1) \varrho_S^t - t \varrho_S^{t-1}) \mathcal{S}_0$$

Divide this by $(t + 1) \varrho_S^t$ and take limits

$$\lim_{t \rightarrow \infty} \frac{\pi_t}{(t + 1) \varrho_S^t} = \frac{\mathcal{A}_S}{\varepsilon} \left(\frac{1 - \varrho_S}{\varrho_S} \right) \mathcal{S}_0$$

Following the same steps as above with the Phillips curve and taking limits yields

$$\left(1 - \frac{\rho_z}{R} \right) \lim_{t \rightarrow \infty} \frac{\pi_t}{(t + 1) \varrho_S^t} = \kappa \lim_{t \rightarrow \infty} \left(\frac{\hat{y}_t - \hat{y}_t^n}{(t + 1) \varrho_S^t} \right)$$

Thus, even in this case, the sign of $\hat{y}_t - \hat{y}_t^n$ and π_t is the same as that of $\mathcal{A}_S \times \mathcal{S}_0$ for large t . \square

F.4 Interest rate rules

We have already seen that under optimal policy, the dynamics of x_t and π_t can be written as functions of x_{t-1} and shocks shown in equations (F.2) and (F.3). Substituting (E.5) into the linearized IS equation

(17):

$$\widehat{y}_t = \left[1 + (1 - \widetilde{\beta}) \Omega\right] \widehat{y}_{t+1} - \frac{1}{\gamma} (i_t - \pi_{t+1}) - \frac{\Lambda}{\gamma} \Gamma_z \varrho_z \widehat{z}_t$$

We can use this equation along with equations (F.2) and (F.3) to express i_t in terms of x_{t-1} and the shocks

$$\begin{aligned} i_t &= \gamma \left[1 + (1 - \widetilde{\beta}) \Omega\right] \widehat{y}_{t+1} - \gamma \widehat{y}_t + \pi_{t+1} - \Lambda \Gamma_z \varrho_z \widehat{z}_t \\ &= \gamma \left[1 + (1 - \widetilde{\beta}) \Omega\right] x_{t+1} - \gamma x_t + \pi_{t+1} + \left\{ \gamma (1 - \widetilde{\beta}) \Omega \delta(\Omega) \frac{y + \rho}{1 + \gamma \rho} \varrho_z - \Lambda \Gamma_z \varrho_z - \gamma \delta(\Omega) \frac{y + \rho}{1 + \gamma \rho} (1 - \varrho_z) \right\} \widehat{z}_t \\ &= \left(\gamma \left[1 + (1 - \widetilde{\beta}) \Omega\right] \mathcal{A}_x - \gamma + \mathcal{B}_x \right) x_t \\ &\quad + \left\{ \gamma (1 - \widetilde{\beta}) \Omega \delta(\Omega) \frac{y + \rho}{1 + \gamma \rho} \varrho_z - \Lambda \Gamma_z \varrho_z - \gamma \delta(\Omega) \frac{y + \rho}{1 + \gamma \rho} (1 - \varrho_z) + \gamma \left[1 + (1 - \widetilde{\beta}) \Omega\right] \mathcal{A}_z \varrho_z + \mathcal{B}_z \varrho_z \right\} \widehat{z}_t \\ &\quad + \left(\gamma \left[1 + (1 - \widetilde{\beta}) \Omega\right] \mathcal{A}_\lambda \varrho_\lambda \widehat{\lambda}_t + \mathcal{B}_\lambda \varrho_\lambda \right) \widehat{\lambda}_t \\ &= \Phi_x x_{t-1} + \Phi_z \widehat{z}_t + \Phi_\lambda \widehat{\lambda}_t, \end{aligned}$$

where

$$\begin{aligned} \Phi_x &= \left\{ \gamma \left[1 + (1 - \widetilde{\beta}) \Omega\right] \mathcal{A}_x - \gamma + \mathcal{B}_x \right\} \mathcal{A}_x \\ \Phi_z &= \left\{ \gamma \left[1 + (1 - \widetilde{\beta}) \Omega\right] \mathcal{A}_x - \gamma + \mathcal{B}_x \right\} \mathcal{A}_z + \gamma (1 - \widetilde{\beta}) \Omega \delta(\Omega) \frac{y + \rho}{1 + \gamma \rho} \varrho_z - \Lambda \Gamma_z \varrho_z \\ &\quad - \gamma \delta(\Omega) \frac{y + \rho}{1 + \gamma \rho} (1 - \varrho_z) + \gamma \left[1 + (1 - \widetilde{\beta}) \Omega\right] \mathcal{A}_z \varrho_z + \mathcal{B}_z \varrho_z \\ \Phi_\lambda &= \gamma \left[1 + (1 - \widetilde{\beta}) \Omega\right] \mathcal{A}_\lambda \varrho_\lambda \widehat{\lambda}_t + \mathcal{B}_\lambda \varrho_\lambda + \left(\gamma \left[1 + (1 - \widetilde{\beta}) \Omega\right] \mathcal{A}_x - \gamma + \mathcal{B}_x \right) \mathcal{A}_\lambda \end{aligned}$$

Next, we show that (29) implements the optimal allocations uniquely. First, note that first-differencing the target criterion (27) yields

$$\phi \pi_t + \phi_x \Delta x_t = 0,$$

where $x_t = \widehat{y}_t - \delta(\Omega) \frac{y + \rho}{1 + \gamma \rho} \widehat{z}_t$ and ϕ and ϕ_x are as defined in the main text in Section 4.3. Since by definition, we have $i_t = i_t^*$ under optimal policy, it follows that the rule (29) is satisfied at the optimal allocation. To see that this rule implements to optimal allocations uniquely, it suffices to look at the determinacy property of the system comprised by the IS curve, the Phillips curve and the interest rate rule absent shocks. This system can be written as

$$\begin{aligned} (\gamma + \phi_x) x_t &= \gamma \left[1 + (1 - \widetilde{\beta}) \Omega\right] x_{t+1} - \Phi_x x_{t-1} - \phi \pi_t + \phi_x x_{t-1} + \pi_{t+1} \\ \pi_t &= \beta \pi_{t+1} + \kappa x_t \end{aligned}$$

In matrix form, this can be written as

$$\begin{bmatrix} x_{t+1} \\ \pi_{t+1} \\ Lx_{t+1} \end{bmatrix} = \begin{bmatrix} \frac{\beta \gamma + \beta \phi_x + \kappa}{\beta \gamma [1 + (1 - \widetilde{\beta}) \Omega]} & -\frac{1 - \beta \phi}{\beta \gamma [1 + (1 - \widetilde{\beta}) \Omega]} & \frac{\Phi_x - \phi_x}{\gamma [1 + (1 - \widetilde{\beta}) \Omega]} \\ -\frac{\kappa}{\beta} & \frac{1}{\beta} & 0 \\ 1 & 0 & 0 \end{bmatrix} \begin{bmatrix} x_t \\ \pi_t \\ Lx_t \end{bmatrix},$$

where $Lx_t \equiv x_{t-1}$. The characteristic polynomial of this system is given by

$$\begin{aligned} \mathcal{P}(\aleph) = & -\left(\frac{1}{\beta} - \aleph\right) \left(\frac{\Phi_x - \phi_x}{\gamma \left[1 + (1 - \tilde{\beta})\Omega\right]} \right) \\ & - \aleph \left\{ \left(\frac{\beta\gamma + \beta\phi_x + \kappa}{\beta\gamma \left[1 + (1 - \tilde{\beta})\Omega\right]} - \aleph \right) \left(\frac{1}{\beta} - \aleph \right) - \frac{1 - \beta\phi}{\beta\gamma \left[1 + (1 - \tilde{\beta})\Omega\right]} \frac{\kappa}{\beta} \right\} \end{aligned}$$

Notice that

$$\mathcal{P}(0) = \frac{\phi \frac{\Upsilon(\Omega)}{1+\gamma\rho} \left(\frac{\lambda-1}{\lambda}\right) - \Phi_x}{\beta\gamma \left[1 + (1 - \tilde{\beta})\Omega\right]} \quad \mathcal{P}(1) = \frac{\kappa(1 - \beta\phi)}{\beta^3\gamma \left[1 + (1 - \tilde{\beta})\Omega\right]}$$

Clearly, for large enough ϕ , we have $\mathcal{P}(0) > 0$ and $\mathcal{P}(1) < 0$, implying that there is at least one root inside the unit circle. Also, note that

$$\frac{\partial \mathcal{P}(\aleph)}{\partial \phi} = \frac{1}{\beta\gamma \left[1 + (1 - \tilde{\beta})\Omega\right]} \left[\frac{\Upsilon(\Omega)}{1 + \gamma\rho} \left(\frac{\lambda-1}{\lambda}\right) (1 - \beta\aleph)(1 - \aleph) - \kappa\aleph \right],$$

which is positive for a finite $\bar{\aleph} > \beta^{-1} > 1$. It follows that for sufficiently large ϕ , $\mathcal{P}(\bar{\aleph}) > 0$. Finally,

$$\lim_{\aleph \rightarrow \infty} \mathcal{P}(\aleph) = -\infty$$

implying that for sufficiently large ϕ , there are two roots above 1. Thus, the system has one stable and two unstable eigenvalues as we have two jump variables (π_t and x_t) and one predetermined variable Lx_t .

G Unequal distribution of profits

The date s problem of an individual i who is a stockholder (d) or a non-stockholder (nd) born at date s can be written as

$$\max_{\{c_t^s(i), \ell_t^s(i), a_{t+1}^s(i)\}} -\mathbb{E}_s \sum_{t=s}^{\infty} (\beta\vartheta)^{t-s} \left(\prod_{k=s}^{t-1} \zeta_k \right) \left\{ \frac{1}{\gamma} e^{-\gamma c_t^s(i)} + \rho e^{\frac{1}{\rho}[\ell_t^s(i) - \xi_t^s(i)]} \right\}$$

s.t.

$$c_t^s(i) + q_t a_{t+1}^s(i) = w_t \ell_t^s(i) + (1 - \tau_t^a) a_t^s(i) + \mathbb{T}_t(i), \quad (\text{G.1})$$

where $a_s^s(i) = 0$ and $w_t = (1 - \tau^w) \tilde{w}_t$ and $\tau_t^a = 0$ for $t > 0$. For a stockholder i , $\mathbb{T}_t(i) = \frac{D_t}{\eta^d} - T_t - J$ where J is the lump-sum tax on stockholders and D_t/η^d is the dividend received by each of the η^d stockholders. For a non-stockholder $\mathbb{T}_t(i) = -T_t + \frac{\eta^d}{1-\eta^d} J$. The individual decision problem then is the same as in Appendix A replacing $D_t - T_t$ with $\mathbb{T}_t(i)$. Thus, following the steps in Appendix A, it is easy to see that the consumption function for stockholders can be written as

$$c_t^s(i; d) = \mathcal{C}_t^d + \mu_t x_t^s(i; d)$$

and for non-stockholders

$$c_t^s(i; nd) = C_t^{nd} + \mu_t x_t^s(i; nd),$$

where the definition of $x = a + w(\xi - \bar{\xi})$ is the same as in the baseline model.

$$C_t^d = -\frac{\vartheta \mu_t}{\mu_{t+1} R_t \gamma} \ln \beta R_t + \frac{\vartheta \mu_t}{\mu_{t+1} R_t} C_{t+1}^d + \mu_t \left[w_t (\rho \ln w_t + \bar{\xi}) + \frac{D_t}{\eta^d} - T_t - J \right] - \frac{\vartheta}{R_t} \frac{\mu_t}{\mu_{t+1}} \frac{\gamma \mu_{t+1}^2 w_{t+1}^2 \sigma_{t+1}^2}{2} \quad (\text{G.2})$$

$$C_t^{nd} = -\frac{\vartheta \mu_t}{\mu_{t+1} R_t \gamma} \ln \beta R_t + \frac{\vartheta \mu_t}{\mu_{t+1} R_t} C_{t+1}^{nd} + \mu_t \left[w_t (\rho \ln w_t + \bar{\xi}) - T_t + \frac{\eta^d}{1 - \eta^d} J \right] - \frac{\vartheta}{R_t} \frac{\mu_t}{\mu_{t+1}} \frac{\gamma \mu_{t+1}^2 w_{t+1}^2 \sigma_{t+1}^2}{2} \quad (\text{G.3})$$

$$\mu_t^{-1} = (1 + \rho \gamma w_t) + \frac{\vartheta}{R_t} \mu_{t+1}^{-1} \quad (\text{G.4})$$

Since $x_t^s(i)$ has mean zero at any date and both types of households have the same μ_t , the goods market clearing condition can be written as

$$\eta^d C_t^d + (1 - \eta^d) C_t^{nd} = y_t$$

Multiplying (G.2) by η^d and (G.3) by $1 - \eta^d$ and adding the two along with market clearing and rearranging yields the aggregate Euler equation which is the same as in the baseline model

$$y_t = -\frac{1}{\gamma} \ln \beta R_t + y_{t+1} - \frac{\gamma}{2} \mu_{t+1}^2 w_{t+1}^2 \sigma_{t+1}^2 \quad (\text{G.5})$$

Combining (G.2) and (G.5):

$$\left(C_t^d - y_t \right) = \frac{\vartheta}{R_t} \frac{\mu_t}{\mu_{t+1}} \left(C_{t+1}^d - y_{t+1} \right) + \mu_t \left(\frac{1 - \eta^d}{\eta^d} d_t - J \right) \quad (\text{G.6})$$

Iterating forward:

$$\mathcal{V}_t \equiv \frac{\eta^d}{1 - \eta^d} \left(\frac{C_t^d - y_t}{\mu_t} \right) = \sum_{s=0}^{\infty} \frac{\vartheta^s}{\prod_{k=0}^{s-1} R_{t+k}} \left[D_{t+s} - \frac{\eta^d}{1 - \eta^d} J \right]$$

In other words, we have $C_t^d = y_t + \frac{1 - \eta^d}{\eta^d} \mu_t \mathcal{V}_t$ as in the main text. Market clearing, then implies that $C_t^{nd} = y_t - \mu_t \mathcal{V}_t$. As claimed in the main text, $J = \frac{1 - \eta^d}{\eta^d} D$ implies that $\mathcal{V} = 0$ in steady state and average consumption of stockholders and non-stockholders is the same, $C^d = C^{nd}$. Thus, as in the main text, we can rewrite the definition of \mathcal{V}_t as

$$\mathcal{V}_t = (D_t - D) + \frac{\vartheta}{R_t} \mathcal{V}_{t+1} \quad (\text{G.7})$$

Since dividends can be written as

$$D_t = \left(1 - \frac{1}{\lambda(1 - \tau^w)} \frac{w_t}{z_t} \right) y_t - \frac{1}{\lambda(1 - \tau^w)} \frac{w_t}{z_t} \frac{\Psi}{2} (\Pi - 1)^2$$

we can write \widehat{D}_t as

$$\widehat{D}_t = \frac{y(y+\rho)}{\rho\lambda} \widehat{z}_t + \underbrace{\frac{1}{\lambda} \left[\lambda - 1 - \left(\frac{1+\gamma\rho}{\rho} \right) y \right]}_{\frac{\partial D}{\partial y}} \widehat{y}_t \quad (\text{G.8})$$

Using the above equation, it is straightforward to derive the expression for \widehat{V}_t in the main text.

G.1 Derivation of the Σ recursion

Even in this case, the objective function of the planner can be written as

$$\mathbb{W}_0 = \sum_{t=0}^{\infty} \beta^t u(c_t, n_t; \bar{\xi}) \Sigma_t,$$

where, as before, Σ_t is defined by

$$\Sigma_t = (1 - \vartheta) \sum_{s=-\infty}^t \int \vartheta^{t-s} e^{-\gamma(c_t^s(i) - c_t)} di$$

Since we have stockholders and non-stockholders, this can be further expanded

$$\begin{aligned} \Sigma_t &= (1 - \vartheta) \left\{ \sum_{s=-\infty}^{t-1} \int \vartheta^{t-s} e^{-\gamma(c_t^s(i) - y_t)} di + \int e^{-\gamma(c_t^t(i) - y_t)} di \right\} \\ &= (1 - \vartheta) \left\{ \sum_{s=-\infty}^{t-1} \int \vartheta^{t-s} e^{-\gamma(c_t^s(i) - y_t)} di + \eta^d \int e^{-\gamma(c_t^t(i;d) - y_t)} di + (1 - \eta^d) \int e^{-\gamma(c_t^t(i;nd) - y_t)} di \right\} \end{aligned}$$

Since $x_t^t(i) = w_t (\xi_t^t(i) - \bar{\xi})$, we have:

$$\begin{aligned} \Sigma_t &= (1 - \vartheta) \sum_{s=-\infty}^{t-1} \int \vartheta^{t-s} e^{-\gamma(c_t^s(i) - y_t)} di \\ &\quad + (1 - \vartheta) \left\{ \eta^d \int e^{-\gamma(c_t^d - y_t + \mu_t w_t (\xi_t^t(i) - \bar{\xi}))} di + (1 - \eta^d) \int e^{-\gamma(c_t^{nd} - y_t + \mu_t w_t (\xi_t^t(i) - \bar{\xi}))} di \right\} \end{aligned}$$

For dates $t > 0$, we can additionally write Σ_t as

$$\begin{aligned}
\Sigma_t &= (1 - \vartheta) \sum_{s=-\infty}^{t-1} \int \vartheta^{t-s} e^{-\gamma(c_{t-1}^s(i) - y_{t-1})} e^{-\gamma(c_t^s(i) - c_{t-1}^s(i) - y_t + y_{t-1})} di \\
&\quad + (1 - \vartheta) \left\{ \eta^d \int e^{-\gamma(C_t^d - y_t + \mu_t w_t (\xi_t^d(i) - \bar{\xi}))} di + (1 - \eta^d) \int e^{-\gamma(C_t^{nd} - y_t + \mu_t w_t (\xi_t^d(i) - \bar{\xi}))} di \right\} \\
&= (1 - \vartheta) \sum_{s=-\infty}^{t-1} \int \vartheta^{t-s} e^{-\gamma(c_{t-1}^s(i) - y_{t-1})} e^{-\gamma \mu_t w_t (\xi_t^s(i) - \bar{\xi})} di \\
&\quad + (1 - \vartheta) \left\{ \eta^d \int e^{-\gamma(C_t^d - y_t + \mu_t w_t (\xi_t^d(i) - \bar{\xi}))} di + (1 - \eta^d) \int e^{-\gamma(C_t^{nd} - y_t + \mu_t w_t (\xi_t^d(i) - \bar{\xi}))} di \right\} \\
&= \vartheta (1 - \vartheta) \sum_{s=-\infty}^{t-1} e^{\frac{\gamma^2 \mu_t^2 w_t^2 \sigma_t^2}{2}} \int \vartheta^{t-1-s} e^{-\gamma(c_{t-1}^s(i) - y_{t-1})} di \\
&\quad + (1 - \vartheta) e^{\frac{\gamma^2 \mu_t^2 w_t^2 \sigma_t^2}{2}} \left[\eta^d e^{-\gamma(C_t^d - y_t)} + (1 - \eta^d) e^{-\gamma(C_t^{nd} - y_t)} \right] \\
&= [\vartheta \Sigma_{t-1} + (1 - \vartheta) \mathbb{B}_t] e^{\frac{\gamma^2 \mu_t^2 w_t^2 \sigma_t^2}{2}}
\end{aligned}$$

or

$$\ln \Sigma_t = \frac{\gamma^2 \mu_t^2 w_t^2 \sigma_t^2}{2} + \ln [\vartheta \Sigma_{t-1} + (1 - \vartheta) \mathbb{B}_t] \quad \text{for } t > 0,$$

where $\mathbb{B}_t = \eta^d e^{-\gamma(C_t^d - y_t)} + (1 - \eta^d) e^{-\gamma(C_t^{nd} - y_t)}$. Given the properties, of C_t^d and C_t^{nd} , we have

$$\mathbb{B}_t = \mathbb{B}(\mu_t \mathcal{V}_t) \equiv \eta^d e^{-\gamma \left(\frac{1 - \eta^d}{\eta^d} \right) \mu_t \mathcal{V}_t} + (1 - \eta^d) e^{\gamma \mu_t \mathcal{V}_t}$$

At date 0, since the utilitarian planner sets $\tau_0^a = 1$, there is no pre-existing wealth inequality and $x_0^s(i) = w_0 (\xi_0^s(i) - \bar{\xi})$ for stockholders and non-stockholders born at some date $s \leq 0$. Thus, we have

$$\begin{aligned}
\Sigma_0 &= (1 - \vartheta) \sum_{s=-\infty}^0 \int \vartheta^{-s} e^{-\gamma(c_0^s(i) - y_0)} di \\
&= (1 - \vartheta) \sum_{s=-\infty}^0 \vartheta^{-s} e^{\frac{1}{2} \gamma^2 \mu_0^2 w_0^2 \sigma_0^2} \left\{ \eta^d e^{-\gamma(C_0^d - y_0)} + (1 - \eta^d) e^{-\gamma(C_0^{nd} - y_0)} \right\} \\
&= e^{\frac{1}{2} \gamma^2 \mu_0^2 w_0^2 \sigma_0^2} \mathbb{B}_0
\end{aligned}$$

or

$$\ln \Sigma_0 = \frac{1}{2} \gamma^2 \mu_0^2 w_0^2 \sigma_0^2 + \ln \mathbb{B}(\mu_0 \mathcal{V}_0)$$

Note that $\mathbb{B}(0) = \mathbb{B}'(0) = 0$ and $\mathbb{B}''(0) = \gamma^2 \left(\frac{1 - \eta^d}{\eta^d} \right) > 0$.

G.2 The planning problem

The planner maximizes

$$\mathbb{W}_0 = \sum_{t=0}^{\infty} \beta^t \left\{ -\frac{1}{\gamma} (1 + \gamma \rho w_t) e^{-\gamma y_t} \Sigma_t \right\}$$

s.t.

$$\begin{aligned}
\gamma y_t &= \gamma y_{t+1} - \ln \beta \vartheta + \ln \mu_{t+1} + \ln [\mu_t^{-1} - (1 + \gamma \rho w_t)] - \frac{\gamma^2 \mu_{t+1}^2 w^2 \sigma^2}{2} e^{2\phi(y_{t+1}-y)} \\
(\Pi_t - 1) \Pi_t &= \frac{\lambda_t}{(\lambda_t - 1) \Psi} \left[1 - \frac{1 + \Omega}{1 - \gamma \rho \Omega} \frac{z_t}{\lambda^{-1} \lambda_t w_t} \right] + \beta \left(\frac{z_t w_{t+1}}{z_{t+1} w_t} \right) (\Pi_{t+1} - 1) \Pi_{t+1} \\
\ln \Sigma_0 &= \frac{\gamma^2 \mu_t^2 w^2 \sigma^2}{2} e^{2\phi(y_t-y)} + \ln \mathbb{B}(\mu_0 \mathcal{V}_0) \quad \text{for } t = 0 \\
\ln \Sigma_t &= \frac{\gamma^2 \mu_t^2 w^2 \sigma^2}{2} e^{2\phi(y_t-y)} + \ln [(1 - \vartheta) \mathbb{B}(\mu_t \mathcal{V}_t) + \vartheta \Sigma_{t-1}] \quad \text{for } t > 0 \\
y_t &= z_t \frac{\rho \ln w_t + \bar{\xi}}{1 + \gamma \rho z_t} - \frac{\Psi}{2} (\Pi_t - 1)^2 \\
\mathcal{V}_t &= \left[1 - \frac{1}{\lambda(1 - \tau^w)} \frac{w_t}{z_t} \right] y_t - \frac{1}{\lambda(1 - \tau^w)} \frac{w_t}{z_t} \frac{\Psi}{2} (\Pi - 1)^2 - \frac{\eta}{1 - \eta} J + \left[\frac{\mu_t^{-1} - 1 - \gamma \rho w_t}{\mu_{t+1}^{-1}} \right] \mathcal{V}_{t+1}
\end{aligned}$$

The first-order condition for \mathcal{V}_t for $t > 0$ is:

$$0 = M_{3,t} \frac{(1 - \vartheta) \mu_t \mathbb{B}'(\mu_t \mathcal{V}_t)}{(1 - \vartheta) \mathbb{B}(\mu_t \mathcal{V}_t) + \vartheta \Sigma_{t-1}} - M_{5,t} + \beta^{-1} \frac{\vartheta}{R_{t-1}} M_{5,t-1}$$

In steady state $\mathcal{V}_t = 0$, and thus we have $M_5 = 0$ in steady state since $\mathcal{B}'(0) = 0$, where $M_{5,t}$ denotes the multiplier on the \mathcal{V}_t recursion. Taking the rest of the first-order conditions and linearizing around the steady state in which the average consumption of stockholders and non-stockholders is equal, we have the following.

FOC w_t

$$\begin{aligned}
& -\gamma(1 + \Omega) \hat{y}_t + (1 + \Omega) \frac{\hat{\Sigma}_t}{\Sigma} - \left(\frac{1 - \tilde{\beta}}{\tilde{\beta}} \right) (1 + \Omega) \hat{m}_{1,t} - m_1 \left(\frac{1 - \tilde{\beta}}{\tilde{\beta}} \right)^2 \frac{\gamma \rho}{1 + \gamma \rho} (1 + \Omega)^2 \frac{\hat{w}_t}{w} \\
& - \left(\frac{1 - \tilde{\beta}}{\tilde{\beta}^2} \right) (1 + \Omega) m_1 \hat{\mu}_t + \frac{\kappa}{\gamma} \hat{m}_{2,t} - \frac{\hat{m}_{4,t}}{\gamma} + \frac{m_4}{\gamma} \frac{\hat{w}_t}{w} - \frac{m_4}{\gamma} \frac{1}{(1 + \gamma \rho)} \hat{z}_t - \frac{1 + \gamma \rho}{\gamma \rho} \hat{m}_{5,t} \frac{y}{\lambda} = 0 \quad (\text{G.9})
\end{aligned}$$

FOC y_t

$$\begin{aligned}
& -\frac{\gamma \rho (1 + \Omega)}{1 + \gamma \rho} \frac{\hat{w}_t}{w} + \gamma \left[1 + 2 \frac{(1 - \Theta)^2}{\Lambda} \left(m_3 - \frac{m_1}{\beta} \right) \right] \hat{y}_t - \frac{\hat{\Sigma}_t}{\Sigma} - \hat{m}_{1,t} + \frac{\Theta}{\beta} \hat{m}_{1,t-1} \\
& + 2(1 - \Theta) \left(m_3 - \frac{m_1}{\beta} \right) \hat{\mu}_t + (1 - \Theta) \hat{m}_{3,t} + \frac{\hat{m}_{4,t}}{\gamma} + \frac{\hat{m}_{5,t}}{\gamma} \left(\frac{\lambda - 1}{\lambda} \right) = 0 \quad (\text{G.10})
\end{aligned}$$

FOC Σ_t

$$\frac{\gamma \rho w}{1 + \gamma \rho w} \frac{\hat{w}_t}{w} - \gamma \hat{y}_t - \hat{m}_{3,t} + \tilde{\beta} \hat{m}_{3,t+1} + \frac{1 - \beta^{-1} \tilde{\beta}^2}{1 - \tilde{\beta}} \frac{\hat{\Sigma}_t}{\Sigma} = 0 \quad (\text{G.11})$$

FOC Π_t

$$-\hat{m}_{2,t} + \hat{m}_{2,t-1} + \Psi m_4 \pi_t = 0 \quad (\text{G.12})$$

FOC μ_t

$$\begin{aligned}
& - \left(\frac{1 - \tilde{\beta}}{\tilde{\beta}^2} \right) \frac{\gamma \rho (1 + \Omega)}{1 + \gamma \rho} m_1 \frac{\hat{w}_t}{w} + \left[2\Lambda \left(m_3 - \frac{m_1}{\beta} \right) - \frac{1 - \tilde{\beta}}{\tilde{\beta}^2} m_1 \right] \hat{\mu}_t + \Lambda \hat{m}_{3,t} \\
& + 2\gamma (1 - \Theta) \left(m_3 - \frac{m_1}{\beta} \right) \hat{y}_t - \frac{1}{\tilde{\beta}} \left(\hat{m}_{1,t} - \frac{\tilde{\beta}}{\beta} (1 - \Lambda) \hat{m}_{1,t-1} \right) = 0
\end{aligned} \tag{G.13}$$

FOC \mathcal{V}_t

$$\begin{aligned}
& \gamma^2 \frac{\mu^2}{1 - \tilde{\beta}} \left(\frac{1 - \eta^d}{\eta^d} \right) \hat{\mathcal{V}}_0 - \hat{m}_{5,0} = 0 \quad \text{for } t = 0 \\
& \gamma^2 \frac{\mu^2}{1 - \tilde{\beta}} \frac{1 - \vartheta}{1 - \vartheta + \vartheta \Sigma} \left(\frac{1 - \eta^d}{\eta^d} \right) \hat{\mathcal{V}}_t - \hat{m}_{5,t} + \beta^{-1} \tilde{\beta} \hat{m}_{5,t-1} = 0 \quad \text{for } t > 0,
\end{aligned} \tag{G.14}$$

where $\hat{m}_{5,t} = \widehat{M}_{5,t}/\mathbb{U}$.

Following the same steps as in Appendix E.3, we can arrive at the following expression which is the analog of equations (E.19)-(E.20) in that Appendix

$$\Upsilon(\Omega) \left[\hat{y}_t - \delta(\Omega) \frac{y + \rho}{1 + \gamma \rho} \hat{z}_t \right] = - (1 + \gamma \rho) \frac{\lambda}{\lambda - 1} \frac{\hat{m}_{2,t}}{\Psi m_4} - \frac{\rho}{m_4} \left(\frac{\partial D}{\partial y} \right) \hat{m}_{5,t}$$

Since $\hat{m}_{2,-1} = 0$, (G.12) implies that $\frac{\hat{m}_{2,t}}{\Psi m_4} = \hat{p}_t$. Using this in the expression above

$$\Upsilon(\Omega) x_t = - (1 + \gamma \rho) \frac{\lambda}{\lambda - 1} \hat{p}_t - \frac{\rho}{m_4} \left(\frac{\partial D}{\partial y} \right) \hat{m}_{5,t},$$

where $x_t = \hat{y}_t - \delta(\Omega) \frac{y + \rho}{1 + \gamma \rho} \hat{z}_t$. Next, for $t = 0$, combining this expression with equation (G.14), one gets the target criterion for date $t = 0$

$$\Upsilon(\Omega) x_0 + (1 + \gamma \rho) \frac{\lambda}{\lambda - 1} \hat{p}_0 + \mathbb{K}(\eta^d) \left(\frac{\partial D}{\partial y} \right) \hat{\mathcal{V}}_0 = 0,$$

where $\mathbb{K}(\eta^d) = \frac{\rho}{m_4} \frac{\gamma^2 \mu^2}{1 - \tilde{\beta}} \left(\frac{1 - \eta^d}{\eta^d} \right)$. Similarly for dates $t > 0$ we have

$$\Upsilon(\Omega) \left(x_t - \frac{\tilde{\beta}}{\beta} x_{t-1} \right) + (1 + \gamma \rho) \frac{\lambda}{\lambda - 1} \left(\hat{p}_t - \frac{\tilde{\beta}}{\beta} \hat{p}_{t-1} \right) + \mathbb{K}(\eta^d) \left(1 - \frac{\tilde{\beta}}{\beta} \right) \left(\frac{\partial D}{\partial y} \right) \hat{\mathcal{V}}_t = 0,$$

which is the same as in Proposition 8. Clearly, $\mathbb{K}(1) = 0$ and $\mathbb{K}'(\eta^d) = -\frac{\rho}{m_4} \frac{\gamma^2 \mu^2}{1 - \tilde{\beta}} \left(\frac{1}{\eta^d} \right)^2 < 0$.

Finally, it is easy to see that with no idiosyncratic risk ($\sigma = 0 \Rightarrow \Omega = 0$), the target criterion becomes

$$\begin{aligned}
& x_0 + (1 + \gamma \rho) \frac{\lambda}{\lambda - 1} \hat{p}_0 + \mathbb{K}(\eta^d) \left(\frac{\partial D}{\partial y} \right) \hat{\mathcal{V}}_0 = 0 \quad \text{for } t = 0 \\
& \left(x_t - \frac{\tilde{\beta}}{\beta} x_{t-1} \right) + (1 + \gamma \rho) \frac{\lambda}{\lambda - 1} \left(\hat{p}_t - \frac{\tilde{\beta}}{\beta} \hat{p}_{t-1} \right) + \mathbb{K}(\eta^d) \left(1 - \frac{\tilde{\beta}}{\beta} \right) \left(\frac{\partial D}{\partial y} \right) \hat{\mathcal{V}}_t = 0 \quad \text{for } t > 0
\end{aligned}$$

As is clear, even in this case, the target criterion is different from RANK and there is a motive to stabilize

\mathcal{V}_t , since $\mathbb{K} \neq 0$.

H Hand-to-mouth households

H.1 Decision problem of HtM households

A HtM agent's problem at any date t can be written as

$$\max_{c_t^s(i;h), \ell_t^s(i;h)} -\frac{1}{\gamma} e^{-\gamma c_t^s(i;h)} - \rho e^{\rho(\ell_t^s(i;h) - \xi_t^s(i))}$$

s.t.

$$c_t^s(i;h) = w_t \ell_t^s(i;h) + D_t - T_t$$

The optimal labor supply can be written as

$$\ell_t^s(i;h) = \rho \ln w_t - \gamma \rho c_t^s(i;h) + \xi_t^s(i;h), \quad (\text{H.1})$$

which is the same as that for the non-HtM households (10). Aggregating the individual labor supply across all HtM and non-HtM households, multiplying by w_t and adding $D_t - T_t$:

$$w_t \ell_t + D_t - T_t = w_t \ln w_t - \gamma \rho w_t y_t + w_t \bar{\xi} + D_t - T_t$$

The LHS of this expression is simply y_t , so we have

$$y_t = \frac{w_t (\ln w_t + \bar{\xi}) + D_t - T_t}{1 + \gamma \rho w_t}$$

Using this and the individual labor supply in the budget constraint for HtM households yields:

$$c_t^s(i;h) = y_t + \tilde{\mu}_t x_t^s(i;h)$$

where $x_t^s(i;h) = w_t (\xi_t^s(i) - \bar{\xi})$ and $\tilde{\mu}_t = (1 + \gamma \rho w_t)^{-1}$.

Since the average consumption of HtM households is y_t , market clearing implies that the average consumption of unconstrained households is also $C_t^{nh} = y_t$. Thus, it follows that the same aggregate Euler equation as in the baseline still holds with a fraction $\eta^h > 0$ of HtM households.

H.2 Deriving the Σ recursion

Even in this case, the objective function of the planner can be written as

$$\mathbb{W}_0 = \sum_{t=0}^{\infty} \beta^t u(c_t, n_t; \bar{\xi}) \Sigma_t$$

where, as before, Σ_t is defined by

$$\Sigma_t = (1 - \vartheta) \sum_{s=-\infty}^t \int \vartheta^{t-s} e^{-\gamma(c_t^s(i) - c_t)} di$$

Since we have HtM and non-HtM households, this can be further expanded

$$\Sigma_t = \underbrace{(1 - \eta^h)(1 - \vartheta) \sum_{s=-\infty}^t \int \vartheta^{t-s} e^{-\gamma(c_t^s(i;nh) - y_t)} di}_{\Sigma_t^{nh}} + \underbrace{\eta^h(1 - \vartheta) \sum_{s=-\infty}^t \int \vartheta^{t-s} e^{-\gamma(c_t^s(i;h) - y_t)} di}_{\Sigma_t^h}$$

Since $c_s^t(i; h) = y_t + \tilde{\mu}_t w_t (\xi_t^s(i) - \bar{\xi})$, we have Σ_t^h

$$\Sigma_t^h = (1 - \vartheta) \sum_{s=-\infty}^t \vartheta^{t-s} \int e^{-\gamma \tilde{\mu}_t w_t (\xi_t^s(i;h) - \bar{\xi})} = e^{\frac{1}{2} \gamma^2 \tilde{\mu}_t^2 w_t^2 \sigma_t^2}$$

Since the consumption function of unconstrained households is the same as in the baseline model, it follows that Σ_t^{nh} evolves as

$$\ln \Sigma_t^{nh} = \frac{\gamma^2 \mu_t^2 w_t^2 \sigma_t^2}{2} + \ln[1 - \vartheta + \vartheta \Sigma_{t-1}^{nh}]$$

H.3 Sensitivity of inequality w.r.t. monetary policy with HTMs

In the presence of HTMs, the welfare relevant measure of inequality at any date t (up to first order) is given by

$$\widehat{\Sigma}_t = (1 - \eta^h) \widehat{\Sigma}_t^{nh} + \eta^h \widehat{\Sigma}_t^h,$$

where (24) describes the evolution of $\widehat{\Sigma}_t$. Up to first order, the relationship between $\Sigma_t^h = \frac{1}{2} \left(\frac{\gamma w_t \sigma_t}{1 + \gamma \rho w_t} \right)^2$ and y_t can be expressed as

$$\widehat{\Sigma}_t^h = \frac{\Lambda}{(1 - \tilde{\beta})^2} \left\{ -\frac{\gamma \rho}{1 + \gamma \rho} \left(\frac{(1 + \gamma \rho) w}{1 + \gamma \rho w} \right) \frac{\widehat{w}_t}{w} + \phi \widehat{y}_t \right\} = -\frac{\gamma}{(1 - \tilde{\beta})^2} \left[(\Theta - 1 + \Lambda) + \Lambda \frac{w - 1}{1 + \gamma \rho w} \right] \widehat{y}_t$$

where we have used the equilibrium relationship between wages and output (19) (we have also set all shocks to zero without loss of generality). Thus, we have

$$\begin{aligned} \widehat{\Sigma}_t &= (1 - \eta^h) \widehat{\Sigma}_t^{nh} + \eta^h \widehat{\Sigma}_t^h \\ &= -\gamma \left[(1 - \eta^h) \Sigma_t^{nh} (\Theta - 1) + \eta^h \Sigma_t^h \frac{(\Theta - 1) + \Lambda \left(\frac{(1 + \gamma \rho) w}{1 + \gamma \rho w} \right)}{(1 - \tilde{\beta})^2} \right] \widehat{y}_t + (1 - \eta^h) \Lambda \widehat{\mu}_t + (1 - \eta^h) \frac{\vartheta}{\beta R} \widehat{\Sigma}_{t-1} \end{aligned}$$

We consider a one-time change in $\widehat{y}_t > 0$ engineered by monetary policy. Since equations (17)-(20), which describe the evolution of macroeconomic aggregates, are purely forward looking, monetary policy can implement this with a change in the nominal interest rate only at date t without affecting the trajectory of macroeconomic aggregates in the future. The change in nominal rates which implement this one-time

increase in date t output can be derived by setting all $t + 1$ variables (and all shocks) in (17)-(19) to zero:

$$\begin{aligned}\widehat{y}_t &= -\frac{1}{\gamma}i_t \\ \widehat{\mu}_t &= -\frac{(1-\tilde{\beta})\gamma\rho w}{1+\gamma\rho w_t}\frac{\widehat{w}_t}{w} + \tilde{\beta}i_t \\ \frac{1+\gamma\rho}{\rho}\widehat{y}_t &= \frac{\widehat{w}_t}{w},\end{aligned}$$

where the first, second and third equations are (17), (18) and (19), respectively. Combining the three equations and eliminating \widehat{w}_t/w yields

$$\widehat{\mu}_t = -\gamma \left[1 + (1 - \tilde{\beta}) \left(\frac{w - 1}{1 + \gamma\rho w} \right) \right] \widehat{y}_t$$

Using the above equation in the expression for $\widehat{\Sigma}_t$ yields

$$\widehat{\Sigma}_t = -\gamma \left[\left(1 - \eta^h \right) \Sigma^{nh} \left[(\Theta - 1) + \Lambda + (1 - \tilde{\beta}) \Lambda \Omega \right] + \eta^h \Sigma^h \left(\frac{(\Theta - 1) + \Lambda + \Lambda \Omega}{(1 - \tilde{\beta})^2} \right) \right] \widehat{y}_t + (1 - \eta^h) \Sigma^h \frac{\vartheta}{\beta R} \frac{\widehat{\Sigma}_{t-1}}{\Sigma}$$

Taking the derivative with respect to η^h :

$$\frac{\partial^2 \widehat{\Sigma}_t}{\partial \eta \partial \widehat{y}_t} = -\gamma \left\{ \left[(\Theta - 1) + \Lambda \frac{(1 + \gamma\rho)w}{1 + \gamma\rho w} \right] \left(\Sigma^h (1 - \tilde{\beta})^{-2} - \Sigma^{nh} \right) + \tilde{\beta} \Lambda \left(\frac{w - 1}{1 + \gamma\rho w} \right) \Sigma^{nh} \right\},$$

which is negative for countercyclical and acyclical risk ($\Theta \geq 1$). Thus, a higher fraction of HTMs (η^h) implies that Σ_t falls more in response to the same increase in output.

H.4 The planning problem

The utilitarian planner maximizes:

$$\mathbb{W}_0 = \sum_{t=0}^{\infty} \beta^t \left\{ -\frac{1}{\gamma} (1 + \gamma\rho w_t) e^{-\gamma y_t} \left[(1 - \eta^h) \Sigma_t^{nh} + \eta^h \Sigma_t^h \right] \right\}$$

s.t.

$$\begin{aligned}\gamma y_t &= \gamma y_{t+1} - \ln \beta \vartheta + \ln \mu_{t+1} + \ln [\mu_t^{-1} - (1 + \gamma\rho w_t)] - \frac{\gamma^2 \mu_{t+1}^2 w^2 \sigma^2}{2} e^{2\phi(y_{t+1}-y)} \\ (\Pi_t - 1) \Pi_t &= \frac{\lambda_t}{(\lambda_t - 1) \Psi} \left[1 - (1 - \tau^w) \frac{z_t}{\lambda^{-1} \lambda_t w_t} \right] + \beta \left(\frac{z_t w_{t+1}}{z_{t+1} w_t} \right) (\Pi_{t+1} - 1) \Pi_{t+1} \\ \ln \Sigma_t^{nh} &= \frac{\gamma^2 \mu_t^2 w^2 \sigma^2}{2} e^{2\phi(y_t-y)} + \ln \left[(1 - \vartheta) + \vartheta \Sigma_{t-1}^{nh} \right] \\ \Sigma_t^h &= \frac{\gamma^2 (1 + \gamma\rho w_t)^{-2} w^2 \sigma^2}{2} e^{2\phi(y_t-y)} \\ y_t &= z_t \frac{\rho \ln w_t + \bar{\xi}}{1 + \gamma\rho z_t} - \frac{\Psi}{2} (\Pi_t - 1)^2\end{aligned}$$

As in the baseline model, the planner chooses τ^w optimally absent shocks. To plot Figures 6 and 7, we first solve for τ^w numerically, then we linearize the first-order conditions and compute the optimal dynamics in response to shocks numerically.

I Irrelevance of government debt in the baseline model

Since we assume that newborn households receive a transfer from the government equal to average wealth B_t/ϑ , our baseline economy features a form of Ricaridan equivalence – the level and path of government debt does not affect real allocations.

To see this, notice that the household's problem is the same as in the baseline with $\mathcal{T}_t = \frac{B_t}{\vartheta}$. Following the same steps in Appendix A, we can arrive at

$$C_t = -\frac{\vartheta\mu_t}{\mu_{t+1}R_t\gamma} \ln \beta R_t + \frac{\vartheta\mu_t}{\mu_{t+1}R_t} C_{t+1} + \mu_t [w_t (\rho \ln w_t + \bar{\xi}) + D_t - T_t] - \frac{\vartheta}{R_t} \frac{\mu_t}{\mu_{t+1}} \frac{\gamma\mu_{t+1}^2 w_{t+1}^2 \sigma_{t+1}^2}{2}, \quad (\text{G.1})$$

which is the same as equation (A.5) in Appendix A. We also know that since $c_t^s(i) = C_t + \mu_t a_t^s(i)$, aggregate consumption is given by

$$c_t = C_t + \mu_t \frac{B_t}{\vartheta} \quad \because (1 - \vartheta) \sum_{s=-\infty}^t \int a_t^s(i) di = \frac{B_t}{\vartheta}$$

Next, aggregating households' labor supply (10):

$$\ell_t = \rho \ln w_t - \gamma \rho \left(C_t + \mu_t \frac{B_t}{\vartheta} \right) + \bar{\xi}$$

Aggregating household budget constraints:

$$\begin{aligned} \frac{1}{R_t} B_{t+1} &= w_t \ell_t + \frac{B_t}{\vartheta} + D_t - T_t - c_t \\ &= w_t [\rho \ln w_t - \gamma \rho c_t + \bar{\xi}] + \frac{B_t}{\vartheta} + D_t - T_t - c_t \\ &= [w_t (\rho \ln w_t + \bar{\xi}) + D_t - T_t] + \frac{B_t}{\vartheta} - (1 + \gamma \rho w_t) c_t \end{aligned}$$

OR

$$w_t (\rho \ln w_t + \bar{\xi}) + D_t - T_t = \frac{1}{R_t} B_{t+1} - \frac{B_t}{\vartheta} + (1 + \gamma \rho w_t) c_t$$

Using this expression in equation (G.1) along with $c_t = C_t + \mu_t \frac{B_t}{\vartheta}$:

$$\begin{aligned} c_t &= -\frac{\vartheta\mu_t}{\mu_{t+1}R_t\gamma} \ln \beta R_t + \frac{\vartheta\mu_t}{\mu_{t+1}R_t} c_{t+1} + \mu_t (1 + \gamma \rho w_t) c_t - \frac{\vartheta}{R_t} \frac{\mu_t}{\mu_{t+1}} \frac{\gamma\mu_{t+1}^2 w_{t+1}^2 \sigma_{t+1}^2}{2} \\ &\quad + \underbrace{\left[\frac{\mu_t}{R_t} B_{t+1} - \frac{\mu_t}{R_t} B_{t+1} - \mu_t \frac{B_t}{\vartheta} + \mu_t \frac{B_t}{\vartheta} \right]}_{=0} \end{aligned}$$

Thus, we can follow the same steps as in Appendix A to derive the same aggregate Euler equation as in the baseline. Since $c_t = y_t$ in equilibrium even with positive government debt, equation (14) which defines

GDP remains the same. Also, it is straightforward to see that the Phillips curve is unaffected by the level of government debt and so it remains to show that the Σ_t recursion is unaffected by non-zero debt. For this, notice that the consumption function can be written as

$$c_t^s(i) = C_t + \mu_t x_t^s(i) = y_t + \mu_t \frac{B_t}{\vartheta} + \mu_t \left(x_t^s(i) - \frac{B_t}{\vartheta} \right),$$

where $x_t^s(i) - \frac{B_t}{\vartheta}$ has mean zero. Following the same steps in Appendix B.2 and replacing $x_t^s(i)$ with $x_t^s(i) - \frac{B_t}{\vartheta}$, it is straightforward to derive the same Σ_t recursion as in the baseline.

Finally, it is worth mentioning that if newborn households did not receive a transfer from the government equal to average wealth, the level and path of debt would matter for allocations. In particular, with positive government debt, there would be across-cohort wealth and consumption inequality and monetary policy would have an additional incentive to address that.