

Optimal Monetary Policy According to HANK*

Sushant Acharya[†]

Edouard Challe[‡]

Keshav Dogra[§]

August 11, 2022

Abstract

We study optimal monetary policy in an analytically tractable Heterogeneous Agent New Keynesian model with rich cross-sectional heterogeneity. Optimal policy differs from a Representative Agent benchmark because monetary policy can affect consumption inequality, by stabilizing consumption risk arising from both idiosyncratic shocks and unequal exposures to aggregate shocks. The tradeoff between consumption inequality, productive efficiency and price stability is summarized in a simple linear-quadratic problem yielding interpretable target criteria. Stabilizing consumption inequality requires putting some weight on stabilizing the level of output, and correspondingly reducing the weights on the output gap and price level relative to the representative agent benchmark.

Keywords: New Keynesian Model, Incomplete Markets, Optimal Monetary Policy

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

*We are grateful to Adrien Auclert, Anmol Bhandari, Florin Bilbiie, Christopher Carroll, Russell Cooper, Davide Debortoli, Konstantin Egorov, 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, EIEF, NUS, Université Paris-Dauphine, Université Paris 8, Banque de France, Drexel University, University of Essex, University of Melbourne and CREST, as well as from conference participants at the Barcelona Summer Forum (Monetary Policy and Central Banking), the NBER SI (Micro Data and Macro Models), the Salento Macro Meetings, ASSA-AEA, SED, EEA, 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.

[†]Bank of Canada and CEPR **Email:** sacharya@bankofcanada.ca

[‡]European University Institute **Email:** edouard.challe@eui.eu

[§]Federal Reserve Bank of New York **Email:** keshav.dogra@ny.frb.org

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; Ravn and Sterk 2020; Bilbiie 2021), 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, henceforth BEGS; Le Grand et al. 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 inequality affects welfare and hence optimal monetary policy. Crucially, in our framework the welfare cost of consumption inequality is summarized by a scalar variable that evolves recursively. This makes the planner’s problem tractable, allowing us to solve explicitly 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 policy can differ in HANK and RANK because uninsurable consumption risk (trivially absent in RANK) reduces 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 uninsurable consumption risk. Our analytical framework distinguishes between two broad ways in which monetary policy can affect consumption risk. First, monetary policy may reduce consumption risk arising from *idiosyncratic* shocks faced by households. Second, it may reduce consumption risk 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 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 baseline, monetary policy can 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 passthrough from individual income shocks to consumption (the *self-insurance channel*). This is because low interest rates facilitate self-insurance both directly through the bond market (by making it easier to borrow to insulate consumption from income shocks), and indirectly through the labor market, due to their expansionary impact on current and future wages (against which households can borrow). The effect of monetary policy on consumption risk via both channels can be summarized by a sufficient statistic: the *cyclicity of consumption risk*, i.e., the effect on consumption risk of a change in output induced by monetary policy. Importantly, we show that when consumption risk is countercyclical, monetary policy can mitigate inefficient fluctuations in consumption risk by stabilizing the level of output.

Our analysis yields both methodological and substantive insights. Methodologically, we (i) derive the welfare-based quadratic objective of the HANK planner, (ii) use this to characterize optimal monetary policy as a solution to a linear-quadratic (LQ) problem, and (iii) express the optimal monetary policy rule in terms of a simple *target criterion* which summarizes the tradeoffs 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 planner’s quadratic loss function places weight on stabilizing the output gap and inflation because the relevant tradeoff in RANK is between departures from productive efficiency and price stability. In HANK, however, the planner also seeks to minimize fluctuations in consumption risk. Our main substantive result is that, in the empirically relevant case of acyclical or countercyclical income risk, this desire to stabilize consumption risk 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. Intuitively, this is because when income risk is acyclical or countercyclical, both the income-risk and self-insurance channels described above make consumption risk countercyclical, implying that stabilizing output mitigates fluctuations in consumption risk. In our calibrated model, the HANK loss function and target criterion feature roughly equal weight on the level of output and output gap and feature a 50% smaller weight on price stability than in RANK. 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 equally exposed to aggregate shocks, uninsurable idiosyncratic risk can substantially change optimal monetary policy.

Our methodological and substantive results carry through to the case where monetary policy affects consumption risk via unequal exposures to aggregate shocks, in addition to idiosyncratic risk. We study two different sources of unequal exposures. First, we allow for unequally distributed dividends by assuming that only a fraction of households 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 increase consumption inequality between stockholders and nonstockholders, since these groups lack access to complete markets to efficiently share aggregate risk. The planner’s desire to avoid such between-group inequality and compensate for missing markets is captured by the presence of the present discounted value of dividends in the quadratic loss function (in addition to output, output gap and the price level).

Second, 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 and incomplete markets against aggregate risk, a surprise interest rate hike redistributes consumption from poor debtors to rich savers (the unhedged interest rate exposure (URE) channel described in Auclert 2019), providing an additional reason to avoid large interest rate hikes in response to aggregate shocks. This motive is absent in our baseline since the utilitarian planner uses fiscal policy to eliminate pre-existing wealth inequality. While the effect of the URE channel is quantitatively small given our calibration, its implications for optimal monetary policy are similar to those of unequally distributed dividends: the non-utilitarian planner places an even higher weight on output stabilization and

implements an even smaller fall in output on impact following a decline in aggregate productivity. Overall, while compensating for missing markets against aggregate risk is conceptually different from facilitating insurance against idiosyncratic income risk, both these motives lead optimal monetary policy to put more weight on output stabilization relative to RANK.

In Appendices [H](#), [I](#) and [J](#), we illustrate the versatility of our framework by extending it in a number of dimensions. First, we study how the presence of *hand-to-mouth* (HtM) households, who have high marginal propensity 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 risk within this group higher and more sensitive to monetary policy than that within the group of unconstrained households – amplifying differences between optimal policy in HANK and RANK. Second, we relax the assumption of i.i.d. idiosyncratic income risk (maintained in our baseline for tractability). As with HtM households, introducing persistent risk does not qualitatively change our results, but quantitatively magnifies the sensitivity of consumption risk to policy and the differences between HANK and RANK. Third, we characterize the optimal monetary policy response to *demand shocks*, i.e., shocks which do not affect the level of output under flexible prices. While optimal policy in RANK features *divine coincidence* ([Blanchard and Galí, 2007](#)) in response to these shocks, the HANK planner deviates from implementing productive efficiency and price stability in order to reduce fluctuations in consumption risk, even though productive efficiency and price stability remains feasible.

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 need not incorporate an explicit measure of inequality either in their objective function or in their reaction function. Instead, these concerns can be addressed by stabilizing the *level* of output, in addition to the output gap and the price level. Stabilizing output can itself stabilize inequality, both by reducing idiosyncratic risk and by preventing aggregate shocks from adversely impacting more vulnerable groups.

Related Literature The papers closest to ours are [BEGS](#) 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 permits more flexibility in the structure of preferences and idiosyncratic shocks, allowing for a quantitative assessment of the importance of heterogeneity for optimal policy, while the second makes it easier to qualitatively isolate and understand the channels by which monetary policy optimally affects consumption inequality. More recently, [McKay and Wolf \(2022\)](#) use sequence-space methods to characterize optimal policy rules in HANK.

[Nuño and Thomas \(2022\)](#) study how URE and unexpected inflation (the 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 tradeoff 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

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.

5.2. More recently, Davila and Schaab (2022) finds that optimal monetary policy in HANK economies under discretion features an inflationary bias: the planner has an incentive to engineer surprise cuts in real interest rates to redistribute towards high marginal utility debtors. This bias is absent in our paper since our planner can use fiscal instruments to deliver the desired level of redistribution, leaving monetary policy free to focus on facilitating insurance against idiosyncratic and aggregate risk, rather than redistribution. An earlier version of our paper considered the case where the planner has a more restricted set of fiscal instruments. In this case, the Ramsey planner had a incentive to engineer a surprise rate cut at date 0 in order to redistribute to high marginal utility debtors, as in Davila and Schaab (2022).

Several authors study optimal monetary policy in New Keynesian economies with limited household heterogeneity (Bilbiie, 2008, 2021; 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). This precludes monetary policy from facilitating self-insurance via asset markets because in equilibrium households do not borrow or lend, spending all their income on consumption. More generally, our paper belongs to the literature studying 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, 2021), social insurance (McKay and Reis, 2016, 2021; Kekre, 2022), and fiscal policy (Auclert et al., 2018; Bilbiie, 2021).

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. Section 4 characterizes 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 describes how various extensions to our baseline model affect optimal monetary policy. 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 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; this ensures that the model features a stationary wealth distribution.³ 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

²See also 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.

³As we discuss in Section 2.1, if we had infinitely lived agents, our model would not feature a stationary wealth distribution.

cohort born at $s < t$ is $(1 - \vartheta)^{\vartheta^{t-s}}$. The date s problem of an individual i born at date s is:

$$\begin{aligned} & \max_{\{c_t^s(i), \ell_t^s(i), a_t^s(i)\}} \mathbb{E}_s \sum_{t=s}^{\infty} (\beta\vartheta)^{t-s} u\left(c_t^s(i), \ell_t^s(i); \xi_t^s(i)\right) \\ \text{s.t.} \quad & c_t^s(i) + q_t a_{t+1}^s(i) = (1 - \tau^w) \tilde{w}_t \ell_t^s(i) + (1 - \tau_t^a) a_t^s(i) + D_t^s(i) - T_t \quad (1) \\ & a_s^s(i) = \mathcal{T}_s \quad (2) \end{aligned}$$

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

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

Each agent i saves in riskless real actuarial bonds, issued by financial intermediaries, 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 baseline does not feature hard borrowing constraints.⁴ 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 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 5.1 considers the case with unequally distributed 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 $\xi_t^s(i) \sim N(\bar{\xi}, \sigma_t^2)$ to their disutility from labor. In our baseline, $\xi_t^s(i)$ is independent across time and individuals; Appendix I allows for persistence in ξ . A larger $\xi_t^s(i)$ reduces disutility and, given wages, increases household 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.⁵ Defining leisure as $l_t^s(i) = \xi_t^s(i) - \ell_t^s(i)$, 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

⁴Appendix H studies a model variant in which a fraction of households (the Hand to Mouth) cannot access asset markets.

⁵We thank Gianluca Violante for suggesting this interpretation.

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_t according to the CES aggregator $y_t = \left[\int_0^1 y_t(j)^{\frac{\varepsilon_t-1}{\varepsilon_t}} dj \right]^{\frac{\varepsilon_t}{\varepsilon_t-1}}$, where ε_t is the elasticity of substitution between varieties. We allow ε_t to vary over time in order to introduce ‘‘cost-push’’ shocks, i.e., shocks to intermediate goods producers’ desired markup $\varepsilon_t/(\varepsilon_t - 1)$. The final good producer’s demand for variety j is:

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

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 $\frac{\Psi}{2} \left(\frac{P_t(j)}{P_{t-1}(j)} - 1 \right)^2 y_t$ 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

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

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:

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

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 $\tau^* = \varepsilon^{-1}$ where ε denotes the steady state elasticity of substitution, 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:

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

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 (see Appendix K). 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:

$$y_t = c_t \equiv (1 - \vartheta) \sum_{s=-\infty}^t \vartheta^{t-s} \int_i c_t^s(i) di, \quad \int_0^1 n_t(j) dj = (1 - \vartheta) \sum_{s=-\infty}^t \vartheta^{t-s} \int_i \ell_t^s(i) di$$

and

$$0 = \frac{B_{t+1}}{\vartheta} = (1 - \vartheta) \sum_{s=-\infty}^t \vartheta^{t-s} \int_i a_{t+1}^s(i) di$$

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 $\varepsilon_0/(\varepsilon_0 - 1)$, which decay geometrically: $\ln z_t = \varrho_z^t \ln z_0$, $\ln\left(\frac{\varepsilon_t}{\varepsilon_t - 1}\right) - \ln\left(\frac{\varepsilon}{\varepsilon - 1}\right) = \varrho_\varepsilon^t \left[\ln\left(\frac{\varepsilon_0}{\varepsilon_0 - 1}\right) - \ln\left(\frac{\varepsilon}{\varepsilon - 1}\right) \right]$. 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 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.⁶

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

$$c_t^s(i) = C_t + \mu_t x_t^s(i) \tag{9}$$

$$\ell_t^s(i) = \rho \ln w_t - \gamma \rho c_t^s(i) + \xi_t^s(i) \tag{10}$$

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

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

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

Proof. See Appendix A. □

⁶Since households have perfect foresight of aggregate variables, only the post-tax real interest rate matters for their decisions. Thus, setting $\tau_t^a \neq 0$ at date $t > 0$ instead of $\tau_t^a = 0$ does not change the set of implementable allocations. Starting from an allocation with $\tau_t^a = 0$ where the pre-tax interest rate 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 prices and allocations unchanged.

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 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; (12) describes its evolution. Iterating this forwards yields

$$\mu_t = \left[\sum_{\tau=0}^{\infty} Q_{t+\tau|t} (1 + \gamma \rho w_{t+\tau}) \right]^{-1} \quad \text{where} \quad Q_{t+\tau|t} = \prod_{k=0}^{\tau-1} \frac{\vartheta}{R_{t+k}}$$

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

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 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 passthrough 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 $y_t = z_t n_t - \frac{\Psi}{2} (\Pi_t - 1)^2 y_t$. Aggregating labor supply (10) and using goods and labor market clearing, we have

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

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:

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

Time varying σ_t Following McKay and Reis (2021), we allow the variance of ξ to vary endogenously with aggregate output to generate 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 $\sigma_t^2 w_t^2 = \sigma^2 w^2 \exp \{2\varphi(y_t - y)\}$ where y denotes steady state output and $\varphi = \frac{\partial \ln \mathbb{V}_t(x)}{\partial y}$ 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 increasing in y_t (*procyclical* risk), when $\varphi > 0$; decreasing in y_t (*countercyclical* risk), when $\varphi < 0$; or independent of y_t (*acyclical* risk) when $\varphi = 0$.⁷ Importantly, what we mean by cyclicity of income risk, and what is measured by φ , is the effect of an increase in output on income risk holding all shocks constant, rather than the correlation between output and income risk. In general, correlation between output and income risk could also arise because aggregate shocks affect both output and idiosyncratic risk.

2.1 Steady state

We now characterize the zero-inflation steady state which, as we show in Section 3.3, is optimal. We normalize steady state productivity to $z = 1$. Since $\tau^* = \varepsilon^{-1}$, imposing $\Pi_t = \Pi_{t+1} = 1$ in (7) requires that $\tilde{w} = 1$ and $w = 1 - \tau w$; steady state output is $y = \frac{\rho \ln w + \bar{\xi}}{1 + \gamma \rho}$. 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. 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

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

⁷More 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 generate countercyclical risk without incorporating a search model. This also allows us to keep our analysis close to the standard NK model.

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 equilibrium conditions 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. Log-linearizing around the zero-inflation steady state and using (14) to substitute out for wages, we have:

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

$$\widehat{\mu}_t = -\gamma \mu w y \left[(1 + \gamma \rho) \widehat{y}_t - \widehat{z}_t \right] + \widetilde{\beta} (\widehat{\mu}_{t+1} + \widehat{i}_t - \pi_{t+1}) \quad (18)$$

$$\pi_t = \beta \pi_{t+1} + \kappa (\widehat{y}_t - \widehat{y}_t^n) \quad (19)$$

where $\widehat{i}_t = \ln(1 + i_t) - \ln R$, $\Theta = 1 - \frac{\Lambda \varphi}{\gamma}$, $\kappa = \frac{\varepsilon}{\Psi} \frac{1 + \gamma \rho}{\rho / y}$, $\widehat{y}_t^n = \frac{(1 + \rho / y) \widehat{z}_t - (\rho / y) \widehat{\varepsilon}_t}{1 + \gamma \rho}$ is the log deviation from steady state of the ‘‘natural’’ level of output i.e. which would prevail under flexible prices ($\Psi = 0$), and $\widehat{\varepsilon}_t = \ln\left(\frac{\varepsilon_t}{\varepsilon_t - 1}\right) - \ln\left(\frac{\varepsilon}{\varepsilon - 1}\right)$.⁸ 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 $\varphi < 0$ implies $\Theta > 1$, procyclical income risk $\varphi < 0$ 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 passthrough, $\widehat{\mu}_{t+1}$, also enters the IS curve as it affects desired precautionary savings. In contrast, idiosyncratic risk does not affect the linearized Phillips curve (19) which is the same as in RANK.

2.3 Calibration

While our results are primarily analytical, when plotting IRFs we parameterize the model as follows. We calibrate the model to an annual frequency and target $r = 4\%$. When choosing the parameters affecting idiosyncratic income risk and its cyclicalities, we calibrate the equilibrium of the HANK economy in which the labor income tax is absent ($\tau^w = 0$). We choose $\bar{\xi}$ to normalize steady state output y to 1 in this economy. We choose the standard deviation of $\xi_t^s(i)$, σ , so that the standard deviation of income in steady state equals $w\sigma(1 - \gamma\rho\mu w) = 0.5$. This is in line with Guvenen et al. (2014) who using administrative data find the standard deviation of 1 year log earnings growth rate to be slightly above 0.5. We set the parameter controlling the cyclicalities of income risk $\varphi = -5.76$ which is broadly consistent with Storesletten et al. (2004).⁹ We set the slope of the Phillips curve $\kappa = 0.1$, and the elasticity of substitution ε to 10, implying a 10% steady state markup. Throughout, we set γ and ρ so that the coefficient of relative

⁸Note that we define $\widehat{\varepsilon}_t$ as the log deviation of desired markups (not the elasticity of substitution) from steady state. That is, $\widehat{\varepsilon}_t > 0$ implies that desired markups are higher, and the elasticity of substitution is lower, than in steady state.

⁹Storesletten et al. (2004) 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 $y \frac{d\sigma_y}{d\gamma} = \frac{0.12 - 0.21}{0.03} = -3$. Using $\sigma_{y,t} = (1 - \gamma\rho\mu_t w_t) w \sigma e^{\varphi(y_t - y)}$, the equilibrium relationship

risk aversion, $-\frac{cu''(c)}{u'(c)} = \gamma c$ and the Frisch elasticity (ρ/y) of the median household equal 2 and 1/3 in steady state respectively, within the range of estimates from the micro literature. We set the persistence of productivity and markup shocks $\rho_z = 0.95^4$ and $\rho_\varepsilon = 0.9^4$ (Bayer et al., 2020). When plotting IRFs, 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 $\sigma_\varepsilon = 0.034$ following Bayer et al. (2020). We set $\vartheta = 0.85$, similar to Nisticò (2016) and Farhi and Werning (2019).

3 Setting up the planning problem

3.1 Social welfare function

In our baseline model, we consider a utilitarian planner who attaches equal weights 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 5.2, 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:¹⁰

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

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*

$$\mathbb{U}_t = u\left(c_t, n_t; \bar{\xi}\right) \times \Sigma_t \quad \text{where} \quad \Sigma_t = (1 - \vartheta) \sum_{s=-\infty}^t \vartheta^{t-s} e^{\frac{1}{2}\gamma^2 \sigma_c^2(s,t)} \quad (20)$$

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 (We will often simply refer to Σ_t as consumption inequality). Absent risk, there would be no consumption inequality and hence $\Sigma_t = 1$ at all dates. However, in the presence of

between μ_t , w_t and y_t , and because we are calibrating cyclical risk in the economy with $\tau^w = 0$, we have:

$$\varphi = \frac{d \ln \sigma_{y,t}}{d \ln y_t} + \frac{\gamma(1 - \tilde{\beta})}{1 + \tilde{\beta}\gamma\rho}$$

Given our calibration, $\varphi = -5.76$ implies $y \frac{d\sigma_y}{dy} = -3$.

¹⁰Note 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.

risk, $\Sigma_t > 1$, reducing welfare relative to RANK. Recall that $u(\cdot) < 0$ and so higher Σ_t reduces welfare. Appendix B.2 shows that Σ_t evolves according to

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

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

The evolution of consumption inequality is an increasing function of consumption risk $\mu_t^2 w_t^2 \sigma_t^2$, which is in turn increasing in both income risk $w_t^2 \sigma_t^2$ and passthrough μ_t^2 . 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).¹¹ Finally, as we describe shortly, surprise changes in μ_0 have an additional effect on consumption inequality which is not present at all other dates.

3.2 Optimal Policy Problem

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. 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 aggregate shocks; τ_0^a and τ^w are fixed at the level which would be optimal absent aggregate shocks.¹² 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.

¹¹Within-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 cash-on-hand distribution. However, since cohorts gradually shrink 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.

¹²Since 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.

3.3 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 suppose 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) (Auclert, 2019), an unanticipated change in interest rates affects consumption inequality. Suppose that at date 0, the planner temporarily cuts real interest rates. This benefits debtors, reducing their interest payments and allowing them to increase consumption; conversely, lower rates reduce savers' interest income and consumption. Thus, lower rates reduce the MPC out of wealth μ_0 , reducing consumption inequality Σ_0 . Using $\Sigma_{-1} = \Sigma = \frac{(1-\vartheta)e^{\frac{\Lambda}{2}}}{1-\vartheta e^{\frac{\Lambda}{2}}}$ and $\mu = \mathbb{E}_{-1}\mu_0$ in (22):

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

While the first two terms on the RHS above 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: 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 optimal monetary policy different at date 0 than at all subsequent dates.¹³ However, the planner also has another instrument which can be used to redistribute from lenders to borrowers, namely the wealth tax. While this instrument is less flexible than monetary policy since it cannot be set in a state contingent way, Appendix D.1 shows that the utilitarian planner optimally sets this tax at a level $\tau_0^a = 1$ which completely eliminates pre-existing wealth inequality, setting the second term in (22) to zero. This not only eliminates the incentive of monetary policy to deliver a surprise rate cut absent shocks, it also leaves households 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

¹³A 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.

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 $t > 0$.¹⁴ Since all inequality at dates $t \geq 0$ arises from uninsurable idiosyncratic risk, the planner's desire to keep inequality low at subsequent dates does not reflect any redistributive motive, but rather the desire to compensate for missing markets against idiosyncratic shocks.

Labor-income tax We also allow the planner to optimally set the constant labor income tax τ^w absent aggregate shocks. As we show in Appendix D.1, this implies that zero inflation is optimal in steady state, and the planner need not use monetary policy to affect inequality on average. This income tax cannot be adjusted in response to aggregate shocks, reflecting the idea that fiscal policy is slow to adjust. Thus, monetary policy still has a role in dealing with changes in inequality in response to aggregate shocks.

In the absence of consumption risk (i.e. in RANK) the optimal labor income tax is $\tau^w = 0$ and the associated steady state level of output is $\bar{\xi}/(1 + \gamma\rho)$ – which is equal to 1 by our normalization (see Sections 2.1 and 2.3). In the presence of consumption risk the planner in general chooses $\tau^w \neq 0$, implying that $w \neq 1$ (the post-tax wage differs from the marginal product of labor) and thus steady state output $y \neq 1$. The planner trades off this productive inefficiency against the benefits of reducing consumption inequality. Appendix D.1 shows that this tradeoff can be summarized by the following optimality condition:

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

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 output does not affect income risk. Thus, raising steady-state output above its productively-efficient level does not reduce income risk (second term on the RHS of (23) is zero). However, higher output and wages still facilitate self-insurance through the labor market and reduce the passthrough from income shocks into consumption, measured by the first term of the RHS, reducing consumption risk. Thus, even with acyclical risk $\Omega = \Omega^c \equiv \frac{\Lambda}{(1 - \tilde{\beta})(1 - \Lambda)} > 0$, the planner subsidizes labor ($\tau^w < 0$) to raise steady state output above its productively efficient level.

Next, consider the income risk channel. With countercyclical income risk ($\Theta > 1$), pushing output above its productively efficient level 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

¹⁴This result is special to the case of the utilitarian planner. In Section 5.2, 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 $t > 0$.

higher output is even larger than if $\Theta = 1$ – both LHS components in (23) are positive and Ω is larger ($\Omega > \Omega^e$). Consequently, the planner subsidizes labor income even more, pushing steady state output further above its productively efficient level.

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 its productively efficient level, 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 then has no first order effect on consumption inequality and the planner does not distort productive efficiency in steady state, setting $\tau^w = 0$ as in RANK. $\Omega = 0$ will be a useful benchmark in what follows.

Importantly, the planner always has an incentive to reduce consumption risk. However, in the steady state with optimal fiscal policy, this incentive is exactly balanced by a first-order cost of reducing productive efficiency further. Given that fiscal policy optimally trades off consumption risk and productive efficiency, monetary policy has no further incentive to increase output in order to reduce consumption risk in steady state. Thus, in response to shocks, monetary policy seeks to stabilize both consumption risk and productive efficiency around their constrained efficient steady state levels, as we will show in Section 4.




Figure 1: **Comparative Statics of Ω .** The curves plot the values of Ω for different values of σ, φ, γ , in each case holding all parameters other than that on the x-axis fixed at their levels in our baseline calibration with countercyclical risk ($\varphi < 0$).

Figure 1 plots comparative statics of Ω with respect to σ, γ and φ . As the previous discussion suggests, only φ affects the sign of Ω (panel c): countercyclical or mildly procyclical risk $\varphi \leq \gamma$ implies $\Omega \geq 0$ while more strongly procyclical risk $\varphi > \gamma$ implies $\Omega < 0$. Higher income risk (higher σ) or higher risk aversion γ increase the welfare cost of inequality, and thus the absolute value of Ω , but do not affect the sign.

3.4 Productive efficiency and the output gap

From equation (19), setting $\hat{y}_t = \hat{y}_t^n$ would implement zero inflation, but this would in general not be efficient. Just as in RANK, deviations in productive efficiency in our model are captured by the “welfare-

relevant” output gap $\hat{y}_t - \hat{y}_t^e$, where \hat{y}_t^e does not respond to inefficient cost-push shocks:¹⁵

$$\hat{y}_t^c = \frac{1 + \rho/y}{1 + \gamma\rho} \hat{z}_t \quad \text{so that} \quad \hat{y}_t^n = \hat{y}_t^c - \frac{\varepsilon}{\kappa\Psi} \hat{\varepsilon}_t$$

This implies that the Phillips curve (19) can equivalently be written as

$$\pi_t = \beta\pi_{t+1} + \kappa(\hat{y}_t - \hat{y}_t^e) + \frac{\varepsilon}{\Psi} \hat{\varepsilon}_t \quad (24)$$

In what follows, with some abuse of terminology, we refer to \hat{y}_t^e as the productively efficient level of output.¹⁶

3.5 How does monetary policy affect inequality?

The key force which will make optimal monetary policy different in HANK versus RANK is the presence of consumption inequality (i.e. $\Sigma_t > 1$) and its sensitivity to monetary policy. Recall from (21) that the dynamics of consumption inequality are driven by consumption risk, which in turn depends on both income risk and the passthrough from income to consumption risk. Thus, the effect of monetary policy on both income risk and passthrough crucially affects how optimal monetary policy in HANK differs from that in RANK. Linearizing (21) and using our assumptions about $w_t\sigma_t$, we have

$$\hat{\Sigma}_t = \underbrace{\Lambda\hat{\mu}_t - \gamma y(\Theta - 1)\hat{y}_t}_{\text{consumption risk}} + \beta^{-1}\tilde{\beta}\hat{\Sigma}_{t-1} \quad (25)$$

(25) reveals that there are two ways in which monetary policy can affect consumption risk. First, monetary policy can lower consumption risk through the self-insurance channel by lowering interest rates and reducing the passthrough $\hat{\mu}_t$. As long as income risk is not acyclical, monetary policy can also affect income risk by affecting the level of output (the income risk channel), captured by the term $-\gamma y(\Theta - 1)\hat{y}_t$. For example, with countercyclical income risk $\Theta > 1$, raising output \hat{y}_t reduces income and hence, consumption risk.

But the planner cannot vary \hat{y}_t and $\hat{\mu}_t$ independently since they only have one instrument – the interest rate. To understand the overall effect of monetary policy on consumption risk through both the self-insurance channels and income risk channels, suppose monetary policy implements a mean reverting cut in interest rates. Figure 2 plots the response to output, μ_t and Σ_t following a 100 bps cut at date 0, after which the real rate is given by $\hat{r}_t = (0.5)^t \hat{r}_0$ for $t > 0$ (panel a). The lower rates reduce passthrough μ_t as shown in panel c. Recall that passthrough is lower when real interest rates are lower or when real wages are higher. Lower rates make it easier for households to self-insure using asset markets, reducing the passthrough of income shocks to consumption: the red dashed-lines in panel b show the response of passthrough $\hat{\mu}_t$ due to the low real rates but keeping wages unchanged. Lower rates also increase output (panel c) and hence wages, making it easier to self-insure using the labor market, lowering $\hat{\mu}_t$ further: the blue line in panel c shows the total effect on passthrough via both real interest rates and wages. This

¹⁵To understand why the output gap captures deviations from productive efficiency, it is useful to relate it to the *labor wedge*, defined as the ratio between household’s marginal rate of substitution between consumption and leisure and the marginal productivity of labor, which is given by w_t/z_t . Up to first-order, the log-deviation of the labor wedge from its steady state value can be expressed as $\frac{1+\gamma\rho}{\rho/y}(\hat{y}_t - \hat{y}_t^e)$, i.e., it is proportional to the output gap.

¹⁶To be clear, a zero output gap $\hat{y}_t - \hat{y}_t^e$ does not imply that output is at its productively efficient level. This is because in steady state, the HANK planner may optimally deviate from productive efficiency by setting $\tau^w \neq 0$ to reduce consumption inequality. A zero output gap implies that the labor wedge takes the same value as in this constrained efficient steady state.

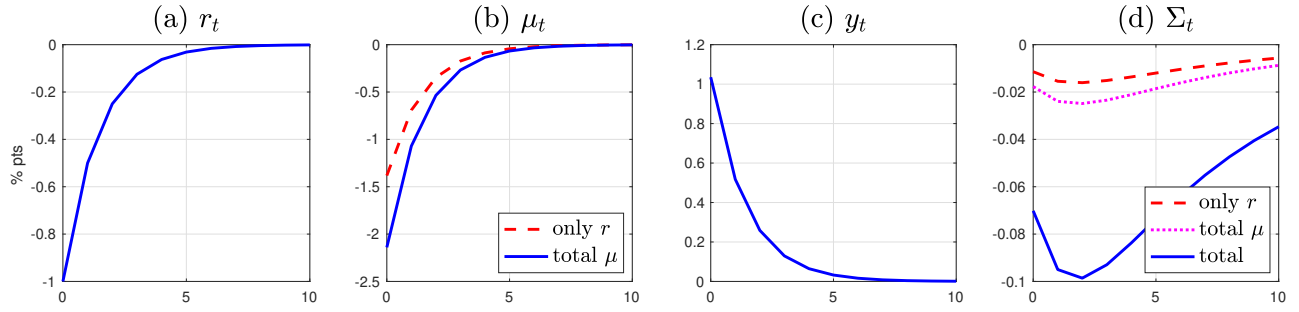


Figure 2: **The effect of monetary policy on consumption inequality** Σ_t . blue curves in panels b,c and d depict the reaction of μ_t , y_t and Σ_t respectively to a mean reverting cut in real interest rates depicted in panel a. The red-dashed line in panel b and d depict the effect of self-insurance on μ_t and Σ_t respectively, only through asset markets. The magenta-dotted line in panel d depicts the effect on Σ_t in which the income-risk channel is shut off. All panels plot log deviations from steady state $\times 100$.

lower passthrough would reduce consumption risk even if income risk was acyclical ($\Theta = 1$): the magenta-dotted line in panel d shows the effect of lower passthrough on consumption inequality, holding income risk fixed. Again, the red-dashed line depicts the effect on $\hat{\Sigma}_t$ due solely to the improvement in the household's ability to self-insure through the asset market, while the magenta-dotted line shows the total effect of lower passthrough. Finally, when risk is countercyclical, the higher output induced by lower interest rates also reduces income risk, lowering consumption inequality even further (the blue-solid line in panel d shows the total effect through all these channels).¹⁷

The effect of monetary policy on self-insurance via asset markets is absent in *zero liquidity* models (Bilbiie, 2008; Hansen et al., 2020; Challe, 2020) in which households do not borrow or lend in equilibrium. While the contribution of this channel (red-dashed line in panel d) on the overall effect of monetary policy on consumption inequality Σ (solid blue line in panel d) is relatively modest given our baseline calibration, this is because our CARA-Normal model features a relatively small MPC, implying that fluctuations in the MPC μ also have a small effect on consumption risk. In quantitative HANK models with a higher average MPC, the effect of monetary policy via the self-insurance channel can be much larger.

While monetary policy affects consumption risk through the multiple channels just described, we can summarize the overall effect through a single sufficient statistic. Since lower interest rates raise \hat{y}_t , the overall effect of monetary policy on consumption inequality can be summarized by a relationship between $\hat{\Sigma}_t$ and \hat{y}_t , using the IS equation (17) and the μ recursion (18) to eliminate $\hat{\mu}_t$ in (25). The coefficient on \hat{y}_t captures the net effect of a cut in interest rates, which raises output, on consumption risk.

Lemma 1 (Dynamics of consumption inequality). *Up to first-order, $\hat{\Sigma}_t$ evolves according to*

$$\hat{\Sigma}_t = \underbrace{-\gamma y (1 - \tilde{\beta}) \Omega (\hat{y}_t - \varkappa(\Omega) \hat{y}_t^e)}_{\text{consumption risk}} + \beta^{-1} \tilde{\beta} \hat{\Sigma}_{t-1} \quad (26)$$

where $\varkappa(\Omega) \in (0, 1)$ for $\Omega \geq \Omega^c$ and is defined in Appendix E.1.

Proof. See Appendix E.1. □

¹⁷If income risk is procyclical $\Theta < 1$, then higher output increases income risk, resulting in a smaller decline or even an increase in Σ_t relative to the acyclical income risk case.

Lemma 1 shows two things. First, the effect of monetary policy on consumption risk via both the self-insurance and income risk channels is summarized by the sufficient statistic $-\gamma y(1 - \tilde{\beta})\Omega$, the *cyclicality of consumption risk*. In other words, changes in interest rates affect output through the IS equation (17), and output in turn affects consumption risk through (26). When income risk is countercyclical ($\Theta > 1$), expansionary monetary policy reduces consumption inequality both by reducing passthrough and by reducing income risk. Thus, consumption risk is also countercyclical: $\partial \hat{\Sigma}_t / \partial \hat{y}_t = -\gamma y(1 - \tilde{\beta})\Omega < 0$. Even when income risk is acyclical ($\Theta = 1$), expansionary policy still reduces passthrough, i.e., consumption risk is still countercyclical, $-\gamma y(1 - \tilde{\beta})\Omega < 0$. When income risk is strongly procyclical ($\Theta \ll 1 \Rightarrow \Omega < 0$), consumption risk is also procyclical: higher output increases inequality as lower passthrough is outweighed by higher income risk, $\partial \hat{\Sigma}_t / \partial \hat{y}_t > 0$. Finally, when $\Omega = 0$, consumption risk is acyclical, and monetary policy cannot affect consumption risk up to first-order: higher output increases income risk but this is exactly balanced by lower passthrough.

Second, consumption risk would be perfectly stabilized by setting $\hat{y}_t = \varkappa(\Omega)\hat{y}_t^c$, where $\varkappa(\Omega) < 1$, i.e., by moving output less than one-for-one with the productively efficient level \hat{y}_t^c . Absent aggregate productivity shocks ($\hat{y}_t^c = 0$), consumption risk depends only on the level of output \hat{y}_t , and is perfectly stabilized by setting $\hat{y}_t = 0$. Stabilizing output perfectly keeps income risk constant; it also keeps real interest rates and wages constant, implying an unchanged passthrough $\hat{\mu}_t = 0$. Of course, since all inequality arises from idiosyncratic risk in our baseline, stabilizing risk is equivalent to stabilizing inequality.

In the presence of productivity shocks, it is no longer necessary to perfectly stabilize output in order to keep consumption risk constant. For example, following a negative productivity shock (\hat{y}_t^c), keeping output constant would require higher real wages to increase labor supply and compensate for the lower productivity. Higher wages would reduce passthrough $\hat{\mu}_t < 0$, *reducing* consumption risk. However, letting output \hat{y}_t fall as much as its productively efficient level \hat{y}_t^c would entail *lower* real wages as well as higher real interest rates, increasing passthrough in addition to increasing income risk (if income risk is countercyclical). Thus, keeping consumption risk constant still requires putting more weight on output stabilization – preventing output from fluctuating one-for-one with its productively efficient level \hat{y}_t^c – but does not require perfect output stabilization.

As we will see in Section 4, the desire to stabilize consumption inequality will lead the HANK planner to put more weight on stabilizing output relative to RANK.

4 Dynamics under optimal monetary policy

As is common in the NK literature, we characterize optimal policy by using a linear-quadratic (LQ) approach.¹⁸ Appendix E.2 shows that after some algebra we can write a quadratic approximation of the planner’s objective function in terms of output and inflation. The HANK planner chooses the sequences $\{\hat{y}_t, \pi_t\}_{t=0}^{\infty}$ to minimize the loss function subject to the linearized Phillips curve (24).

¹⁸The presence of consumption inequality means that the HANK economy is not at its first-best level in the zero inflation steady state. Consequently, as described by Benigno and Woodford (2005) in RANK, maximizing a quadratic approximation to the welfare objective subject to linear constraints will not yield a first-order accurate approximation to optimal policy owing to the presence of first-order terms in the quadratic approximation. Thus, in Appendix E.2, following Benigno and Woodford (2005), we eliminate these linear terms using a second order approximation of the implementability conditions.

Proposition 3 (Optimal Monetary Policy in HANK). *The LQ approximation of the planning problem described in Section 3.2 is given by*

$$\begin{aligned} \min_{\{\hat{y}_t, \pi_t\}_{t=0}^{\infty}} \quad & \frac{1}{2} \sum_{t=0}^{\infty} \beta^t \left\{ \Upsilon(\Omega) \left(\hat{y}_t - \delta(\Omega) \hat{y}_t^e \right)^2 + \frac{\varepsilon}{\kappa} \pi_t^2 \right\} \\ \text{s.t.} \quad & \pi_t = \beta \pi_{t+1} + \kappa (\hat{y}_t - \hat{y}_t^e) + \frac{\varepsilon}{\Psi} \hat{\varepsilon}_t, \end{aligned} \quad (27)$$

where $\Upsilon(\Omega)$ and $\delta(\Omega)$ are defined in Appendix E.2 and satisfy $\Upsilon(0) = \delta(0) = 1$. When income risk is acyclical or countercyclical $\left(\Theta \geq 1 \Rightarrow \Omega \geq \Omega^c = \frac{\Lambda}{(1-\beta)(1-\Lambda)} > 0 \right)$, $\Upsilon(\Omega) > 1$ and $\delta(\Omega) \in (0, 1)$.

Proof. See Appendices E.2 and E.3. □

To understand this LQ problem, it is useful to compare it to its RANK counterpart.

Corollary 1. *In the RANK economy without idiosyncratic risk ($\sigma = 0 \Rightarrow \Omega = 0$), $\Upsilon = \delta = 1$, i.e., the planner's problem becomes*

$$\begin{aligned} \min_{\{\hat{y}_t, \pi_t\}_{t=0}^{\infty}} \quad & \frac{1}{2} \sum_{t=0}^{\infty} \beta^t \left\{ \left(\hat{y}_t - \hat{y}_t^e \right)^2 + \frac{\varepsilon}{\kappa} \pi_t^2 \right\}, \\ \text{s.t.} \quad & \pi_t = \beta \pi_{t+1} + \kappa (\hat{y}_t - \hat{y}_t^e) + \frac{\varepsilon}{\Psi} \hat{\varepsilon}_t \end{aligned} \quad (28)$$

The HANK and RANK planners in (27) and (28), respectively, face the same constraint: the Phillips curve is unaffected by heterogeneity and idiosyncratic risk. Thus, all differences between HANK and RANK are summarized by the different weights in the planner's loss functions. The RANK loss function (28) is a special case of (27): absent idiosyncratic income risk $\sigma = 0$, $\Omega = 0$ 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 (28)) and price stability, which would be attained by setting $\pi_t = 0$ (last term in (28)).

The HANK planner has an additional third objective: stabilizing consumption inequality.¹⁹ When income risk is acyclical or countercyclical ($\Omega \geq \Omega^c > 0$), this motive leads to two key differences between the HANK and RANK loss functions. First, the HANK planner puts some weight on stabilizing the *level* of output rather than purely trying to minimize the output *gap*, in order to mitigate fluctuations in consumption risk. To see this, note that the first term in the loss function (27) can be written as

$$\Upsilon(\Omega) \left(\hat{y}_t - \delta(\Omega) \hat{y}_t^e \right)^2 = \Upsilon(\Omega) \left(\underbrace{[1 - \varpi(\Omega)] (\hat{y}_t - \varkappa(\Omega) \hat{y}_t^e)}_{\text{consumption risk}} + \underbrace{\varpi(\Omega) (\hat{y}_t - \hat{y}_t^e)}_{\text{productive efficiency}} \right)^2, \quad (29)$$

where $\varpi(\Omega) \in (0, 1)$ when income risk is acyclical or countercyclical (see Appendix E.3). Recall from Lemma 1 that consumption risk is proportional to $\hat{y}_t - \varkappa(\Omega) \hat{y}_t^e$ where $\varkappa(\Omega) < 1$. In other words, consumption risk would be perfectly stabilized at its steady state level by setting $\hat{y}_t = \varkappa(\Omega) \hat{y}_t^e$, adjusting output less than one-for-one with changes in its productively efficient level \hat{y}_t^e . The first term in the loss function

¹⁹Again, absent optimal fiscal policy, the HANK planner would seek to use monetary policy to reduce consumption risk, rather than merely stabilizing it at its steady state level. Given that fiscal policy optimally trades off consumption risk and productive efficiency in steady state, monetary policy has no further incentive to reduce consumption risk absent aggregate shocks, and instead seeks to stabilize consumption risk at its steady state level. Also, as explained earlier, since all inequality arises from idiosyncratic risk in our baseline, stabilizing inequality is equivalent to stabilizing risk.

reflects a compromise between this objective of stabilizing consumption risk and the RANK objective of maintaining productive efficiency: it depends on a convex combination of $\hat{y}_t - \varkappa(\Omega)\hat{y}_t^e$ and the output gap, $\hat{y}_t - \hat{y}_t^e$, with weights $1 - \varpi(\Omega)$ and $\varpi(\Omega)$ respectively. Since $\delta(\Omega) = \varpi(\Omega) + (1 - \varpi(\Omega))\varkappa(\Omega)$ is a convex combination of 1 and $\varkappa(\Omega) \in (0, 1)$, we have $\delta(\Omega) < 1$, i.e., this component of the loss function would be minimized by moving \hat{y}_t less than one-for-one with \hat{y}_t^e – a compromise between stabilizing the *level* of output and the *output gap*. δ can be thought of as the weight on output gap stabilization, relative to output stabilization:

$$\hat{y}_t - \delta(\Omega)\hat{y}_t^e = [1 - \delta(\Omega)] \underbrace{\hat{y}_t}_{\text{output level}} + \delta(\Omega) \underbrace{(\hat{y}_t - \hat{y}_t^e)}_{\text{output gap}}$$

In our calibration with countercyclical risk, $\delta = 0.6$, implying roughly equal weight on output and output gap stabilization.

Second, compared to the RANK planner, the HANK planner puts more weight on stabilizing economic activity relative to inflation, reflecting the fact that stabilizing economic activity now also mitigates fluctuations in consumption risk, in addition to fostering productive efficiency. The weight on the first term $\Upsilon(\Omega) (\hat{y}_t - \delta(\Omega)\hat{y}_t^e)^2$ is scaled up by a factor $\Upsilon(\Omega) > 1$. In our calibration, $\Upsilon = 1.76$, implying that the relative weight on price stability is almost halved relative to RANK. Thus, the HANK planner will tolerate higher fluctuations in inflation and smaller output fluctuations.




Figure 3: **Comparative Statics:** The blue curves denotes the values of Υ , the magenta-dashed curves denote δ and the black-dotted curves denote ϖ for different values of σ, φ, γ .

Figure 3 plots Υ (solid blue line) and δ (dashed-magenta line) as functions of income risk σ , risk aversion γ and cyclical income risk φ . When $\sigma = 0$, the HANK and RANK objective functions are trivially identical, $\Upsilon = \delta = 1$. As σ increases, the level of consumption risk also increases and so stabilizing risk becomes more important, warranting larger deviations from RANK (higher $\Upsilon > 1$ and lower $\delta < 1$; see panel (a)). Similarly, if households were risk neutral, consumption risk/inequality would not be costly and so the planner’s objective function would remain the same as in RANK ($\Upsilon = \delta = 1$). Higher risk aversion γ makes fluctuations in consumption risk more costly, again warranting larger deviations from RANK (Υ is increasing, while δ is decreasing in γ ; see panel (b)). Finally, panel (c) shows that more countercyclical income risk (more negative φ) tends to cause the HANK planner to put more weight on stabilizing the level of output relative to either the output gap or price level. Intuitively, when consumption risk is more sensitive to fluctuations in the level of output, output fluctuations are more costly because they

lead to larger fluctuations in consumption risk. In addition to δ (the weight on output gap stabilization relative to output level stabilization), Figure 3 also plots ϖ (the weight on productive efficiency relative to consumption risk). The comparative statics of δ and ϖ are very similar, reflecting the fact that stabilizing consumption risk requires close to perfect output stabilization ($\varkappa \approx 0.1$ in our calibration).

Our analytic approach uncovers that how much weight the HANK planner puts on output stabilization depends on Ω , which is proportional to our sufficient statistic $-\gamma y(1 - \tilde{\beta})\Omega$ in equation (26). This reveals that the planner puts more weight on output stabilization not merely because consumption inequality exists, but because fluctuations in consumption inequality depend on fluctuations in output. Indeed, in the special case where consumption risk is acyclical ($\Omega = 0$), even though consumption inequality exists and reduces welfare, its evolution does not depend on the level of output (cf. equation (26)). In this case fluctuations in output do not add to fluctuations in consumption inequality, so they are no more costly than in RANK, and the planner can continue to focus on productive efficiency and price stability ($\Upsilon = \delta = 1$). It follows that 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 the two economies.²⁰

Lemma 2. *In HANK with $\Omega = 0$, the planner’s objective function becomes (28) as in RANK. Consequently, optimal policy implements the same sequence $\{\hat{y}_t, \pi_t\}$ in both economies.*

4.1 Target Criterion

Since the HANK and RANK planners face the same constraint (the Phillips curve), differences in their objective functions directly translate into differences between the *target criteria* describing optimal policy in the two economies (again, except in the special case with $\Omega = 0$, where the objective functions and target criteria are the same in HANK and RANK). Specifically, since cyclical consumption risk makes fluctuations in output more costly (cf. equation (27)), it leads optimal policy to put more weight on stabilizing the level of output, relative to either the output gap or the price level.

Proposition 4 (Target Criterion). *In HANK, optimal policy is characterized by the following target criterion, for all dates $t \geq 0$:*

$$(1 - \delta(\Omega))\hat{y}_t + \delta(\Omega)(\hat{y}_t - \hat{y}_t^c) + \frac{\varepsilon}{\Upsilon(\Omega)}\hat{p}_t = 0 \quad (30)$$

while the target criterion in RANK ($\delta(\Omega) = \Upsilon(\Omega) = 1$) becomes:

$$(\hat{y}_t - \hat{y}_t^c) + \varepsilon\hat{p}_t = 0. \quad (31)$$

In RANK, optimal monetary policy takes the form of *flexible price level targeting*: the planner stabilizes a weighted average of the output gap and the price level. Reflecting the differences in the objective function

²⁰Werning (2015) has highlighted that the presence of idiosyncratic risk and incomplete markets does not *necessarily* change the *positive* properties of New Keynesian economies. We uncover a parallel irrelevance result regarding the *normative* properties of HANK economies: optimal policy does not differ from RANK simply because inequality exists, but because monetary policy can affect inequality. Werning (2015) “as-if” result obtains in a zero liquidity economy when income risk is acyclical, i.e., individual income is proportional to aggregate income. Because his economy features zero liquidity, acyclical income risk also implies acyclical consumption risk.

of the two planners, the HANK planner deviates from this in two ways, putting some weight on the *level* of output in addition to the output gap and price level, and putting a lower weight on the price level relative to economic activity. We now document how these differences affect the dynamic response to shocks.

4.2 Productivity shocks

As is well known, given our maintained assumption that the subsidy τ^* eliminates steady state monopolistic distortions, in response to productivity shocks in RANK optimal policy features a *divine coincidence*: it is both feasible and optimal to implement zero inflation ($\hat{p}_t = 0$) while closing the output gap ($\hat{y}_t - \hat{y}_t^e = 0$). Intuitively, maintaining output at its productively efficient level also keeps prices stable. This can be seen from the target criterion (31) along with the Phillips curve (24), which (given $\hat{\epsilon}_t = 0$) imply $\hat{p}_t = \hat{y}_t - \hat{y}_t^e = 0$. Figure 4 plots the optimal response to a date 0 productivity shock in RANK (red-dashed line) and HANK (blue-solid line). The red dashed-lines in panels (a) and (b) show that the RANK planner responds to a fall in productivity which decreases $\hat{y}_t^e = \frac{1+\rho/y}{1+\gamma\rho}\hat{z}_t < 0$ by tracking this level, $\hat{y}_t = \hat{y}_t^e < 0$, resulting in zero inflation and achieving both productive efficiency and price stability.

Since the HANK planner has an additional objective – stabilizing inequality – while they *could* implement $\hat{y}_t = \hat{y}_t^e$ and $\pi_t = 0$, they will 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 falling as much as the flexible-price level of output \hat{y}_t^e initially. This entails positive inflation initially. In contrast, the planner commits to mildly negative output gaps ($\hat{y}_t < \hat{y}_t^e < 0$) in the future, which in turn entail mild deflation in the future. This is formalized in the following Proposition.

Proposition 5. *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^e 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 $\hat{y}_t < \hat{y}_t^e$. Following an increase in productivity all these signs are reversed, i.e., π_t and $\hat{y}_t - \hat{y}_t^e$ are negative at date 0 and positive for all $t \in (T, \infty)$ for some $T > 0$.*

Proof. See Appendix F. □

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

Now suppose that following a negative productivity shock, monetary policy continued to set $\hat{y}_t = \hat{y}_t^e < 0$ $\forall t \geq 0$ (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 4. This raises the first-order benefit of marginally increasing output above \hat{y}_t^e to curtail the rise in inequality. Meanwhile, at $\hat{y}_t = \hat{y}_t^e$ the cost of marginally increasing output above \hat{y}_t^e , measured by the output gap, $\hat{y}_t - \hat{y}_t^e$, remains unchanged. Since the benefit of increasing \hat{y}_t above \hat{y}_t^e increases while the cost of doing so remains unchanged, the planner sets $0 > \hat{y}_t > \hat{y}_t^e$. Output still falls on impact, but by less than the flexible-price level of output \hat{y}_t^e , implying a positive output gap (blue curve

in panel (a)). This tradeoff is also reflected in the target criterion (30), which can be rewritten as:

$$\left(\hat{y}_t - \delta(\Omega)\hat{y}_t^e\right) + \frac{\varepsilon}{\Upsilon(\Omega)}\hat{p}_t = 0$$

Intuitively, rather than tracking \hat{y}_t^e one-for-one, which would stabilize the output gap, the planner seeks to minimize the gap between \hat{y}_t and $\delta(\Omega)\hat{y}_t^e$ (where $\delta(\Omega) < 1$). This reflects a compromise between the planner's goal of stabilizing inequality, which calls for stabilizing output, and fostering productive efficiency.

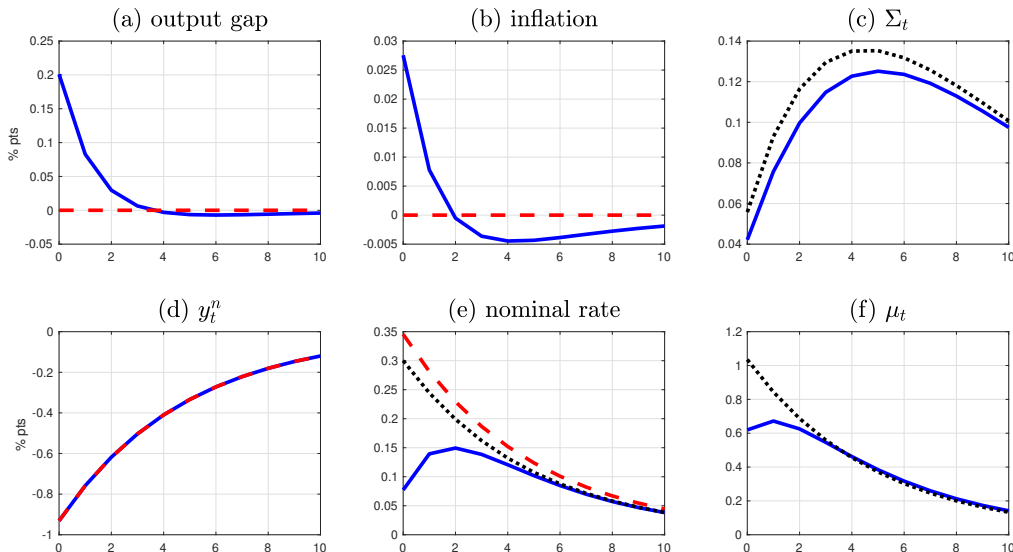


Figure 4: **Optimal policy in response to productivity shocks** in HANK (solid blue curves) and RANK (dashed red curves). Black-dotted curves denote outcomes in HANK under the non-optimal policy which sets $\hat{y}_t = \hat{y}_t^n, \pi_t = 0$ for all $t \geq 0$. All panels plot log-deviations from steady state $\times 100$.

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 passthrough from income to consumption risk (blue curve in panel (f)) than would occur if monetary policy set $\hat{y}_t = \hat{y}_t^e$ and $\pi_t = 0$ (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 $\hat{y}_t = \hat{y}_t^e$ and $\pi_t = 0$ for all $t \geq 0$ (black-dotted curve in panel (c)).

Implementing $\hat{y}_t > \hat{y}_t^e$ results in inflation early on. The planner tolerates higher inflation in order to cushion the fall in output, as can be seen from the lower weight on price stability in (30) (since $\Upsilon(\Omega) > 1$). Nonetheless, to mitigate this rise in inflation, the planner commits to set \hat{y}_t slightly below \hat{y}_t^e in the future, lowering both future and date 0 inflation because of the forward looking nature of the Phillips curve.

4.3 Markup shocks

We now discuss the optimal response to markup shocks. Absent productivity shocks, the output gap in the target criterion (30) is simply $\hat{y}_t - \hat{y}_t^e = \hat{y}_t$. Even in RANK, markup shocks 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 \hat{y}_t^e (which remains unchanged) and the level of output consistent with zero inflation i.e. $\hat{y}_t^n = -\frac{\varepsilon}{\kappa\Psi}\hat{\varepsilon}_t$. 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. Thus, the RANK planner responds to a positive markup shock by permitting some fall in output (red-dashed line in panel a, Figure 5) 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, a fall in output *improves* productive efficiency. Second, a fall in output increases consumption inequality, reducing welfare. Proposition 6 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 (30) to the case with only markup shocks (implying $\hat{y}_t^e = 0$):

$$\hat{y}_t + \frac{\varepsilon}{\Upsilon(\Omega)} \hat{p}_t = 0$$

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 acyclical or countercyclical income risk.

Proposition 6. *Consider a HANK economy with acyclical/countercyclical income risk, and a RANK economy where the median households in the RANK and HANK economies have the same coefficient of relative risk aversion (γ/y) and Frisch elasticity (ρ/y) in steady state. Under optimal policy in HANK, following an increase in firms' desired markup ($\hat{\varepsilon}_0 > 0$), at date 0, \hat{y}_0 falls (but less than \hat{y}_0^n) and $\pi_0 > 0$. Furthermore, the fall in output is smaller than under 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 5 plots IRFs following a positive markup shock under optimal policy in RANK (dashed-red curves) and HANK with acyclical or countercyclical income risk (solid blue 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 (dotted-black curve which lies above the solid blue 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 passthrough $\hat{\mu}_t$ (panel (f)). As in RANK, the HANK planner commits to modest deflation in the future to mitigate the initial rise in inflation.

4.4 Implementing optimal policy using an interest rate rule

Equation (30) 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

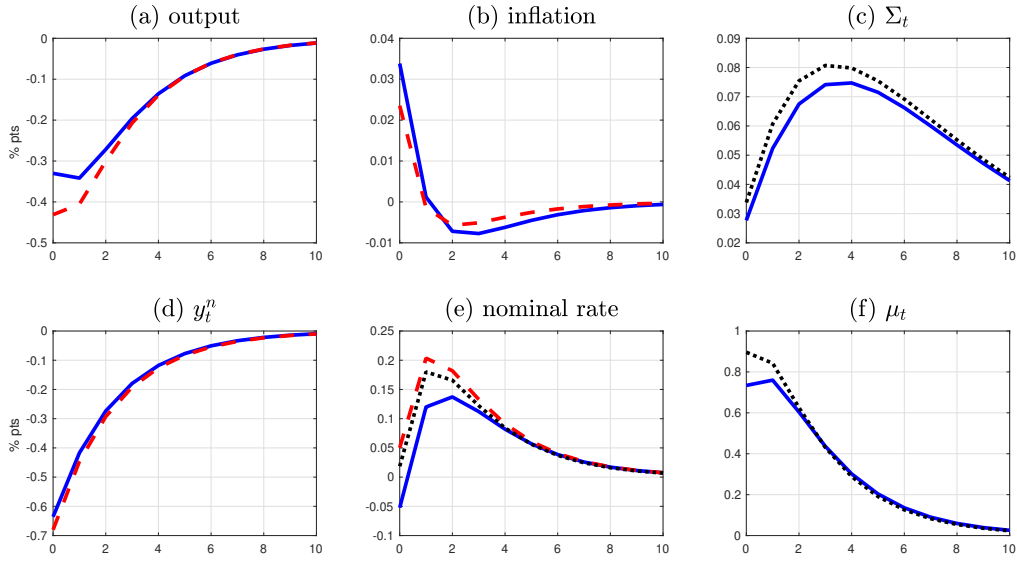


Figure 5: **Optimal policy in response to markup shocks** in HANK (solid blue lines) and RANK (dashed red lines). Black-dotted lines denote outcomes in HANK under non-optimal policy which implements the same $\{\hat{y}_t, \pi_t\}_{t=0}^{\infty}$ as in RANK. All panels plot log-deviations from steady state $\times 100$.

allocations and inflation. One such interest rate rule is

$$i_t = i_t^* + \phi\pi_t + \phi_{\text{gap}}(\Delta y_t - \Delta \hat{y}_t^e) + \phi_y \Delta y_t \quad (32)$$

where $\phi_{\text{gap}} = \phi \frac{\Upsilon(\Omega)}{\varepsilon} \delta(\Omega)$ is the weight on the change in output gap, $\phi_y = \phi \frac{\Upsilon(\Omega)}{\varepsilon} (1 - \delta(\Omega))$ is the weight on output growth and i_t^* , defined in Appendix F.4, denotes the equilibrium nominal interest rate under optimal policy. Appendix F.4 shows that for ϕ sufficiently large, this rule implements the optimal allocations as a unique equilibrium. With acyclical or countercyclical income risk (implying $\Upsilon > 1$), this rule reacts more strongly to changes in output growth and the output gap, relative to π_t , compared to the corresponding rule in RANK where $\Upsilon = 1, \delta = 0$. Again, notice that (32) does not require the policymaker to change nominal rates in response to changes in some measure of inequality; the concern for inequality is captured by a larger coefficients on stabilizing real activity relative to π_t .

5 Unequal exposure to aggregate shocks

In general, market incompleteness affects households' ability to insure against aggregate as well as idiosyncratic risk. When different households – borrowers vs lenders, stockholders vs non-stockholders – are unequally exposed to aggregate shocks, they would efficiently share this risk given access to complete asset markets. Market incompleteness prevents this, implying that monetary policy may be able to improve welfare by facilitating insurance against aggregate as well as idiosyncratic risk. In our baseline model, monetary policy has no such role since households are equally exposed to aggregate shocks: all households receive an equal share of profits and the utilitarian planner removes all pre-existing wealth inequality. This allowed us to focus on idiosyncratic risk in Section 4. We now relax these assumptions, allowing for unequally distributed profits and initial wealth inequality, and study how the planner's desire to compensate for missing markets to insure against aggregate risk affects optimal monetary policy.

5.1 Unequal distribution of profits

We now relax our baseline assumption of equally distributed profits by assuming that a fraction $\eta^d < 1$ in each cohort s receive an equal share of dividends (“stockholders”), while the remaining $1 - \eta^d$ households receive no dividends (“non-stockholders”). Both groups supply labor and face the same distribution of idiosyncratic shocks $\xi_t^s(i)$. Appendix G presents the utilitarian planner’s problem in this economy. In addition to the instruments available to the planner in the baseline, we allow the planner to levy a lump sum tax $J = \frac{1-\eta^d}{\eta^d}D$ on stockholders (where D denotes steady state dividends) and make a lump-sum transfer $\frac{\eta^d}{1-\eta^d}J$ to each non-stockholder, equalizing the average consumption of the two groups in steady state. This ensures that unequally distributed profits do not introduce an incentive for monetary policy to redistribute between stockholders and non-stockholders on average. However, the transfer cannot be adjusted to keep 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 $c_t^s(i; d) = \mathcal{C}_t^d + \mu_t x_t^s(i; d)$ while that of a non-stockholder is $c_t^s(i; nd) = \mathcal{C}_t^{nd} + \mu_t x_t^s(i; nd)$, where

$$\mathcal{C}_t^d = y_t + \left(\frac{1-\eta^d}{\eta^d}\right) \mu_t \mathcal{V}_t, \quad \mathcal{C}_t^{nd} = y_t - \mu_t \mathcal{V}_t, \quad \text{and} \quad \mathcal{V}_t = (D_t - D) + \frac{\vartheta}{R_t} \mathcal{V}_{t+1}. \quad (33)$$

\mathcal{V}_t denotes the present discounted value of dividends relative to their steady state value D . Appendix G shows that linearizing (33) yields the valuation equation

$$\tilde{\mathcal{V}}_t = \mathcal{D}_y \hat{y}_t + \mathcal{D}_z \hat{z}_t + \beta \tilde{\mathcal{V}}_{t+1} \quad \text{where} \quad \mathcal{D}_y = \frac{1}{\varepsilon} - \frac{\varepsilon - 1}{\varepsilon} \left(\frac{1 + \gamma \rho}{\rho/y}\right), \quad \mathcal{D}_z = \frac{1 + \rho/y}{(\rho/y)} \frac{\varepsilon - 1}{\varepsilon}, \quad (34)$$

where $\tilde{\mathcal{V}}_t = \hat{\mathcal{V}}_t/y$ denotes the deviation of \mathcal{V}_t in levels divided by steady state output. \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 $\frac{\varepsilon}{\varepsilon-1}$.

The consumption of the two groups is equalized in steady state since $\mathcal{V} = 0$. However, the two groups are unequally exposed to aggregate shocks which affect dividends. If households had access to complete markets for aggregate shocks, stockholders and non-stockholders would insure each other and the consumption of the two groups would not diverge in response to aggregate shocks. In our incomplete markets economy, such insurance is not possible and shocks which increase current or future dividends tend to increase the consumption of stockholders for a given aggregate income and reduce the consumption of non-stockholders (and conversely for shocks which reduce dividends). 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 $\frac{\gamma^2 \mu_t^2 w_t^2 \sigma_t^2}{2}$ increase inequality. Secondly, between-group consumption inequality arising from unequally distributed dividends increases Σ_t for a given level of risk (see Appendix G for details):

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

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.

The implications of this source of between-group inequality can be seen by inspecting the planner's quadratic loss function (36) in the Proposition below.

Proposition 7 (Optimal Policy with an unequal distribution of profits). *The utilitarian planner's LQ problem can be written as*

$$\min_{\{\hat{y}_t, \pi_t, \hat{\mathcal{V}}_t\}_{t=0}^{\infty}} \frac{1}{2} \sum_{t=0}^{\infty} \beta^t \left\{ \Upsilon(\Omega) \left(\hat{y}_0 - \delta(\Omega) \hat{y}_0^e \right)^2 + \frac{\varepsilon}{\kappa} \pi_0^2 \right\} + \frac{\mathbb{K}(\eta^d)}{2} \left\{ \hat{\mathcal{V}}_0^2 + \left(1 - \beta^{-1} \tilde{\beta}\right) \sum_{t=1}^{\infty} \beta^t \hat{\mathcal{V}}_t^2 \right\} \quad (36)$$

subject to the Phillips curve (24) and the valuation equation (34). $\mathbb{K}(\eta^d) \geq 0$ is defined in Appendix G and satisfies $\mathbb{K}(1) = 0$ and $\mathbb{K}'(\eta^d) < 0$, i.e. more concentrated wealth (lower η^d) increases \mathbb{K} . Optimal policy satisfies the following target criterion for $t = 0$:

$$\Upsilon(\Omega) x_0 + \varepsilon \hat{p}_0 + \mathbb{K}(\eta^d) \mathcal{D}_y \hat{\mathcal{V}}_0 = 0 \quad (37)$$

and for $t > 0$:

$$\Upsilon(\Omega) \left(x_t - \beta^{-1} \tilde{\beta} x_{t-1} \right) + \varepsilon \left(\hat{p}_t - \beta^{-1} \tilde{\beta} \hat{p}_{t-1} \right) + \mathbb{K}(\eta^d) \mathcal{D}_y \left(1 - \beta^{-1} \tilde{\beta} \right) \hat{\mathcal{V}}_t = 0 \quad (38)$$

where $x_t = \hat{y}_t - \delta(\Omega) \hat{y}_t^e$ and \mathcal{D}_y is the effect of higher output on dividends.

Proof. See Appendix G.3. □

When $\eta^d = 1$ and dividends are equally distributed, the last term in (36) vanishes and the loss function and target criteria are identical to (27) and (30) in the baseline, respectively. When $\eta^d < 1$, the last term is non-zero and 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. Stabilizing $\hat{\mathcal{V}}_t$ helps compensate for the absence of complete markets against aggregate shocks affecting the path of dividends.

The weight on stabilizing dividends in the loss function $\mathbb{K}(\eta^d)$ is increasing in the concentration of stockholdings since higher concentration amplifies the effect of a given change in dividends on the consumption gap. The sign of the coefficient on stabilizing dividends in the target criterion, however, also depends on the effect of higher output on dividends \mathcal{D}_y . If $\mathcal{D}_y < 0$, the planner seeks to implement higher output (even compared to the baseline) in response to a shock which raises $\hat{\mathcal{V}}_t$. Higher $\hat{\mathcal{V}}_t$ 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}_t 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. While compensating for missing markets against aggregate risk is conceptually different from facilitating insurance against idiosyncratic income risk, both motives lead optimal monetary policy to put more weight on output stabilization.

This motive for stabilizing dividends is particularly strong at date 0 since a change in $\hat{\mathcal{V}}_0$ generates consumption gaps between all stockholders and non-stockholders 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 weight on $\widetilde{\mathcal{V}}_t^2$ in the loss function and target criterion is not time invariant: the planner puts more weight on stabilizing \mathcal{V}_t at $t = 0$ than at all subsequent dates.

The motive to stabilize \mathcal{V}_t would be present even in an economy in which idiosyncratic risk is absent ($\sigma_t = 0$), or present but insensitive to monetary policy ($\sigma_t > 0$ but $\Omega = 0$), as long as $\eta^d < 1$. This would imply $\Upsilon = 1$ but $\mathbb{K}(\eta^d) \neq 0$. Even if idiosyncratic risk is absent or insensitive to monetary policy, unequally distributed dividends would leave households imperfectly insured against aggregate risk, allowing monetary policy to improve welfare by substituting for these missing markets as in BEGS. Reducing idiosyncratic risk and providing insurance against aggregate risk are two *distinct* motives which cause optimal policy in HANK to differ from RANK.

While the target criterion characterizes the optimal response to all shocks, we now focus on markup shocks to save space. Figure 6 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 and hence $\widetilde{\mathcal{V}}_t$ (panel d) driving the average consumption of stockholders \widehat{C}_t^d above that of non-stockholders \widehat{C}_t^{nd} (panel e). This effect is more severe, the more concentrated are stockholdings: 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 line denotes the baseline with equally distributed 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% pts. in response to a positive markup shock. This is in line with the numerical results of BEGS, whose HANK planner raises output by 0-0.1% pts.²¹ Similarly, both HANK planners allow inflation to increase on impact (by 0.1% pts in our economy, 0.2% pts in BEGS), followed by mild deflation – in contrast to RANK which features a fall in output and a smaller initial increase in inflation.

5.2 The non-utilitarian planner and the URE channel

Even with equally distributed dividends, initial wealth inequality would also leave households unequally exposed to aggregate shocks. As described in Section 3, a utilitarian planner optimally uses the wealth tax to eliminate initial wealth inequality, removing this source of unequal exposure. As we show next, a non-utilitarian (NU) planner chooses not to eliminate pre-existing wealth inequality, implying that savers and borrowers are unequally exposed to changes in interest rates.

The NU planner maximizes the Pareto weighted sum of households' lifetime utilities, assigning different weights to households with 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

²¹Unequally 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. BEGS take a similar approach: they introduce an unequal distribution of dividends, calibrate the tax rate on dividends in line with U.S. data, and calibrate Pareto weights so that absent shocks, this dividend tax is optimal.

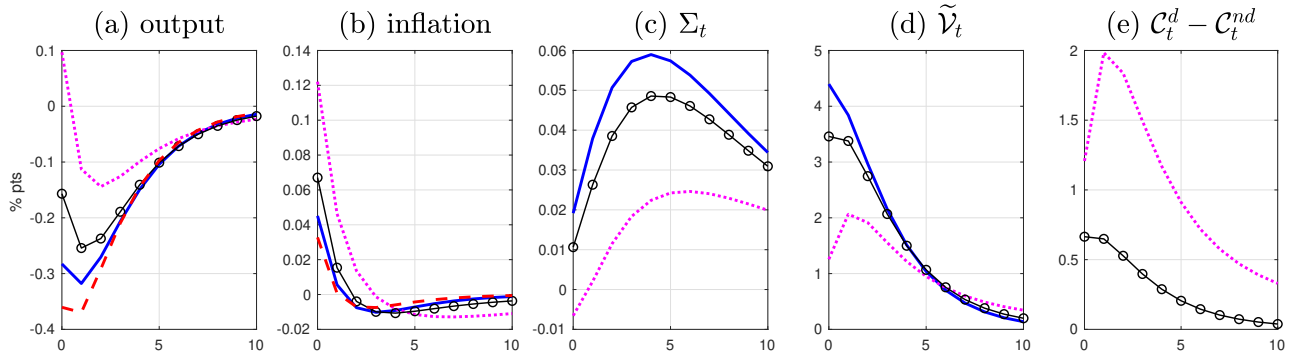


Figure 6: **Optimal policy in response to markup shocks** The red-dashed lines depict dynamics in RANK. All other curves depict dynamics in HANK in response to an increase in firms' desired markups. The blue curves depict the case with equal distribution of profits, the black lines with circle markers 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 $\times 100$, except d and e which plot $100 \times \tilde{\mathcal{V}}_t$ and $100 \times (\hat{C}_t^d - \hat{C}_t^{nd})/y$ respectively.

$e^{\gamma\alpha a_0^s(i)}$ to households with wealth $a_0^s(i)$ at date 0.²² $\alpha \geq 0$ indexes the planner's tolerance for pre-existing wealth inequality. When $\alpha = 0$, the planner is utilitarian and puts equal weights on all individuals alive at date 0. The larger α , the higher the relative weight on individuals with higher wealth at date 0. Given α , 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 welfare cost of consumption inequality Σ_t , which is now defined as

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

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 all inequality resulting from idiosyncratic shocks from date 0 onwards as undesirable. This is reflected in the fact that while (21) is unchanged, (22) is now given by

$$\ln \Sigma_0 = \frac{1}{2} \gamma^2 \mu_0^2 w_0^2 \sigma_0^2 + \underbrace{\ln \left[\frac{1 - \vartheta}{1 - \vartheta e^{\frac{\Lambda}{2} \left(\frac{\alpha - (1 - \tau_0^a) \mu_0}{\mu} \right)^2}} \right]}_{\text{welfare cost of date 0 wealth inequality}} \quad (39)$$

For $\alpha > 0$, completely eliminating wealth inequality (setting $\tau_0^a = 1$), no longer sets the last term on the RHS to zero. If the planner has access to a state contingent wealth tax (which can be changed in

²²Since 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 10, this weight is β^s .

response to the aggregate shock), they would still set this term to zero, eliminating any undesirable wealth inequality, but the level of wealth tax that accomplishes this is now $1 - \tau_0^a = \alpha/\mu_0$. 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 8 (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 a 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 aggregate shocks; the planner can only use the wealth tax to deliver the desired level of redistribution absent aggregate shocks. Appendix D.1 shows that the wealth tax that accomplishes this is $1 - \tau_0^{a*} = \alpha/\mu$. Absent shocks, this tax sets the second term on the RHS of (39) 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 given by $\ln \left[\frac{1-\vartheta}{1-\vartheta e^{\frac{\Lambda}{2} \left(\frac{\alpha}{\mu} \right)^2 \left(\frac{\mu-\mu_0}{\mu} \right)^2}} \right]$, which differs from zero unless $\alpha = 0$ (the planner is utilitarian) or $\mu_0 = \mu$ (there is no aggregate shock). Since some wealth inequality remains, a surprise change in interest rates still redistributes between savers and borrowers (the 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 passthrough 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.²³ Thus the effect of μ_t on Σ_t is different at date 0 than at subsequent dates.

If households had access to complete markets against aggregate shocks, the consumption gap would not respond to surprise changes in interest rates and the effects just described would be absent. With incomplete markets, monetary policy takes into account the URE channel when responding to an aggregate shock, compensating for missing insurance markets against aggregate risk. As in our baseline economy, optimal policy can be characterized in terms of an LQ problem; the effect of the URE channel is reflected in the fact that the date 0 loss function is different than at subsequent dates.

Proposition 9 (NU planner's LQ problem). *The LQ approximation to the NU planner's problem is*

$$\min_{\{\hat{y}_t, \pi_t\}_{t=0}^{\infty}} \frac{1}{2} \left\{ \Upsilon_0(\Omega) (\hat{y}_0 - \delta_0(\Omega) \hat{y}_0^e)^2 + \frac{\varepsilon}{\kappa} \pi_0^2 \right\} + \frac{1}{2} \sum_{t=1}^{\infty} \beta^t \left\{ \Upsilon(\Omega) (\hat{y}_t - \delta(\Omega) \hat{y}_t^e)^2 + \frac{\varepsilon}{\kappa} \pi_t^2 \right\}$$

subject to the Phillips curve (24). $\Upsilon_0(\Omega) \geq \Upsilon(\Omega)$ is increasing in α and $\Upsilon_0(\Omega) = \Upsilon(\Omega)$ when $\alpha = 0$. The

²³Since 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 $(\mu_0 - \mu)^2$.

solution to this problem is characterized by the following target criterion for $t = 0$

$$\left(1 - \delta_0(\Omega)\right)\widehat{y}_0 + \delta_0(\Omega)\left(\widehat{y}_0 - \widehat{y}_0^e\right) + \frac{\varepsilon}{\Upsilon_0(\Omega)}\widehat{p}_0 = 0, \quad (40)$$

while for $t > 0$, the target criterion is the same as that for the utilitarian planner, (30) in Proposition 3.

Proof. See Appendices E.2 and E.3. □

At dates $t > 0$, the loss functions and target criterion are the same as in our baseline with the utilitarian planner: the NU planner’s preference for wealth redistribution, and the extent of initial inequality, do not modify the tradeoff between price stability, productive efficiency and consumption insurance relative to Section 4. However at $t = 0$, the NU planner puts more weight on stabilizing output and the output gap, relative to inflation, than at subsequent dates: $\Upsilon_0(\Omega) > \Upsilon(\Omega)$. 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 her desired level. But the planner only has one instrument – the nominal interest rate – which affects both μ_0 and y_0 . Thus, stabilizing $\widehat{\mu}_0$ requires keeping y_0 closer to its steady state level, even if this comes at the cost of higher inflation. It is worth noting that this stabilization of y_0 occurs even when idiosyncratic consumption risk is acyclical ($\Omega = 0$) and insensitive to monetary policy. Thus, while the effect of monetary policy on consumption inequality via the URE channel is distinct from its effect via idiosyncratic risk, the qualitative implications for optimal policy are similar: monetary policy should put even more weight on stabilizing output relative to price stability.




Figure 7: Optimal policy in response to a negative productivity shock with the non-utilitarian planner Blue lines depict the utilitarian baseline, dotted-magenta lines depict the planner with $\alpha = \mu$ and the red-dashed curves depict RANK. All panels plot log-deviations from steady state $\times 100$.

The URE channel turns out to be weak under our baseline calibration, implying that the optimal path of output and inflation depends little on the planner’s Pareto weights. Figure 7 shows the optimal response to a negative productivity shock. Blue lines depict the utilitarian baseline ($\tau_0^{a*} = 100\%$) and dotted-magenta lines 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 . Qualitatively, the NU planner sets lower interest rates at date 0 (panel c), implementing an even smaller decline in date 0 output, in order to prevent the MPC out of wealth from rising sharply (panel d). Quantitatively, however,

the differences between the solid blue and dotted magenta lines is small.²⁴ The URE channel has little bite under our calibration because our CARA-Normal model features a small average MPC; it could be more powerful in quantitative models featuring a larger or heterogeneous MPCs and more wealth dispersion.

6 Extensions and some discussion

Our versatile framework can be extended in many directions to study how additional channels and shocks affect optimal monetary policy in HANK. We present three such extensions in the appendix. Appendix H extends our analysis to include MPC heterogeneity by incorporating a fraction of hand-to-mouth (HtM) households into our baseline economy. While this does not qualitatively change our results, adding high MPC households makes consumption risk and inequality quantitatively more sensitive to changes in output induced by monetary policy. Thus, the differences between optimal monetary policy in HANK and RANK are magnified and the HANK planner stabilizes output fluctuations even more than in our baseline.

Appendix I extends our baseline by allowing for persistent idiosyncratic income risk. Similarly to the extension with HtM households, introducing persistent income risk does not qualitatively change our results, but quantitatively magnifies the sensitivity of consumption risk to monetary policy. Thus, introducing persistence also increases the differences between optimal monetary policy in HANK and RANK, leading the HANK planner to stabilize output even more than in our baseline.

Finally, Appendix J studies the optimal monetary policy response to demand shocks, i.e., shocks which do not affect the flexible-price level of output. Since these shocks do not induce a tradeoff between productive efficiency and price stability, optimal policy under RANK features divine coincidence in response to these shocks, implementing $\hat{y}_t - \hat{y}_t^n = \pi_t = 0$. We show that this divine coincidence policy is in general not optimal in HANK, even though it always remains feasible. This is because perfectly stabilizing prices and productive efficiency can cause demand shocks to create excessive fluctuations in inequality. Instead, optimal policy deviates from price stability and productive efficiency in order to mitigate these fluctuations in inequality.

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 stabilize consumption inequality in HANK; our analytical framework sharply distinguishes between two ways in which monetary policy can do this. First, monetary policy can reduce fluctuations in idiosyncratic consumption risk, compensating for the absence of markets to insure against idiosyncratic shocks. Second, monetary policy can reduce fluctuations in between-group inequality arising from unequal exposures to aggregate shocks and policy, compensating for missing markets against aggregate shocks. When consumption risk is countercyclical, both idiosyncratic risk and unequal exposures lead optimal monetary policy to put some weight on stabilizing output, and correspondingly less weight on productive efficiency and price stability, in response to aggregate productivity and markup shocks.

As the extensions mentioned in Section 6 illustrate, our tractable framework can be extended in many

²⁴The effect of the URE channel on the NU planner's optimal response to a markup shock (not plotted here) is also small.

directions to study how other shocks or features of HANK economies affect optimal policy. It can also be used as a framework to think about what features of quantitative HANK models affect optimal policy in these environments. In quantitative HANK models, the same broad motives – reducing idiosyncratic consumption risk and reducing fluctuations in between-group inequality – still shape the differences between optimal policy in HANK and RANK. However, the quantitative importance of the various channels we identify could be different. For example, while we find that monetary policy’s effect on passthrough plays a relatively modest role, relative to its effect on income risk, this is because our CARA-Normal framework delivers small MPCs. The effect on passthrough could be more important in a quantitative model with CRRA preferences and binding borrowing constraints, which deliver higher average MPCs.

Our results also help identify which features of quantitative HANK models would cause optimal policy to differ from RANK. For example, we showed that countercyclical income risk tends to increase the difference between optimal policy in HANK and RANK. This would still be true in quantitative HANK models, but the relevant definition of cyclical income risk would be different. In our model with CARA utility, the relevant measure is the cyclical income risk. In the CARA-Normal framework, it is the cyclical income risk which determines whether, for example, there is compounding or discounting in the aggregate Euler equation (Acharya and Dogra, 2020). In contrast, in models with CRRA utility (e.g., the zero-liquidity models in Werning 2015 and Bilbiie 2021), it is the cyclical income risk that affects differences between the positive properties of HANK and RANK models.²⁵ Thus, when using our results to assess what optimal policy would be in a quantitative HANK model with CRRA utility, one should check the cyclical income risk.

Finally, a practical implication of our analysis is that monetary policymakers who are concerned with inequality need not 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.

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.

²⁵In Bilbiie (2021), for example, the cyclical income risk determines whether there is compounding or discounting in the aggregate Euler equation; in Werning (2015) the slope of the aggregate Euler equation is identical in HANK and RANK when *log* income risk is acyclical.

- Benigno, Pierpaolo and Michael Woodford**, “Inflation Stabilization and Welfare: The Case of a Distorted Steady State,” *Journal of the European Economic Association*, 2005, 3 (6), 1185–1236.
- Bhandari, Anmol, David Evans, Mikhail Golosov, and Thomas J Sargent**, “Inequality, Business Cycles and Monetary-Fiscal- Policy,” *Econometrica*, 2021, 89 (6), 2559–2599.
- Bilbiie, Florin O.**, “Limited Asset Markets Participation, Monetary Policy and (Inverted) Aggregate Demand Logic,” *Journal of Economic Theory*, May 2008, 140 (1), 162–196.
- , “Monetary Policy and Heterogeneity: An Analytical Framework,” Working Paper, November 2021.
- Blanchard, Olivier Jean and Jordi Galí**, “Real Wage Rigidities and the New Keynesian Model,” *Journal of Money, Credit and Banking*, 2007, 39, 35 – 65.
- 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.
- Davila, Eduardo and Andreas Schaab**, “Optimal Monetary Policy with Heterogeneous Agents: A Timeless Ramsey Approach,” Technical Report 2022.
- 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, Fatih Karahan, Serdar Ozkan, and Jae Song**, “What do data on millions of US workers reveal about lifecycle earnings dynamics?,” *Econometrica*, 2021, 89 (5), 2303–2339.
- , **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.
- Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante**, “The Macroeconomic Implications of Rising Wage Inequality in the United States,” *Journal of Political Economy*, 2010, 118 (4), 681–722.
- 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 2022.
- 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 Christian K Wolf**, “Optimal Policy Rules in HANK,” Technical Report, Working Paper 2022.
- 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), 3133–3158.
- 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 Redistributive Inflation,” *Annals of Economics and Statistics*, 2022, 146, 3–64.
- 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: Proofs

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 Appendix J. The optimal labor supply decision of household i is given by:

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

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

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

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-measured cash-on-hand and so, $x_{t+1}^s(i)$ is given by

$$x_{t+1}^s(i) = (1 - \tau_{t+1}^a) a_{t+1}^s(i) + w_{t+1} (\xi_{t+1}^s(i) - \bar{\xi})$$

Substituting out for $a_{t+1}^s(i)$ and $\ell_t^s(i)$ using (A.1) and (A.2), and using the definition of $x_t^s(i)$, the above expression can be written as

$$x_{t+1}^s(i) = \frac{(1 - \tau_{t+1}^a) R_t}{\vartheta} \left\{ x_t^s(i) + w_t (\rho \ln w_t + \bar{\xi}) + D_t - T_t - (1 + \gamma \rho w_t) c_t^s(i) \right\} + w_{t+1} (\xi_{t+1}^s(i) - \bar{\xi})$$

Since $x_{t+1}^s(i)$ is normally distributed, given (A.4), $c_{t+1}^s(i)$ is also normally distributed with mean:

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

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:

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

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

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

Using $c_t^s(i) = \mathcal{C}_t + \mu_t x_t^s(i)$, we have:

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

Matching coefficients on $x_t^s(i)$, we have for all $t \geq 0$:

$$\mu_t^{-1} = 1 + \gamma \rho w_t + \frac{\vartheta}{(1 - \tau_{t+1}^a) R_t} \mu_{t+1}^{-1} \quad (\text{A.5})$$

Notice that (A.5) is the same as (12) in the main text once we use the fact that $\tau_{t+1}^a = 0$ for all $t \geq 0$. Next, since the expression above must hold for all values of $x_t^s(i)$ including $x_t^s(i) = 0$, we have

$$\begin{aligned} \mathcal{C}_t &= -\frac{\vartheta \mu_t}{\mu_{t+1} (1 - \tau_{t+1}^a) R_t \gamma} \ln[\beta \zeta_t (1 - \tau_{t+1}^a) R_t] + \frac{\vartheta \mu_t}{\mu_{t+1} (1 - \tau_{t+1}^a) R_t} \mathcal{C}_{t+1} + \mu_t [w_t (\rho \ln w_t + \bar{\xi}) + D_t - T_t] \\ &\quad - \frac{\vartheta}{(1 - \tau_{t+1}^a) R_t} \frac{\mu_t}{\mu_{t+1}} \frac{\gamma \mu_{t+1}^2 w_{t+1}^2 \sigma_{t+1}^2}{2} \end{aligned} \quad (\text{A.6})$$

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.6) together with $\mathcal{C}_t = y_t$ yields

$$\begin{aligned} [1 - \mu_t(1 + \gamma \rho w_t)] y_t &= -\frac{\vartheta \mu_t}{\mu_{t+1}(1 - \tau_{t+1}^a) R_t} \frac{1}{\gamma} \ln[\beta \zeta_t (1 - \tau_{t+1}^a) R_t] + \frac{\vartheta \mu_t}{\mu_{t+1}(1 - \tau_{t+1}^a) R_t} y_{t+1} \\ &\quad - \frac{\vartheta}{(1 - \tau_{t+1}^a) R_t} \frac{\mu_t}{\mu_{t+1}} \frac{\gamma \mu_{t+1}^2 w_{t+1}^2 \sigma_{t+1}^2}{2} \end{aligned}$$

Next, (A.5) implies that $1 - \mu_t(1 + \gamma \rho w_t) = \frac{\vartheta \mu_t}{\mu_{t+1}(1 - \tau_{t+1}^a) R_t}$ so dividing both sides of the equation above by $1 - \mu_t(1 + \gamma \rho w_t)$ yields

$$y_t = -\frac{1}{\gamma} \ln[\beta \zeta_t (1 - \tau_{t+1}^a) R_t] + y_{t+1} - \frac{\gamma \mu_{t+1}^2 w_{t+1}^2 \sigma_{t+1}^2}{2}$$

which is the same as equation (11) in the main text once we have used the fact that $\tau_{t+1}^a = 0$ for all $t \geq 0$.

B Derivation of Σ recursion

B.1 Evolution of cash-on-hand within cohort

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 forwards 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 $\mathcal{C}_t = y_t$. 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:

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

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:

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

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 for dates $t > 0$ 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 $a_0^s(i) \sim N(0, -sw^2\sigma^2)$ for $s \leq 0$, this can be rewritten as:

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

Taking logs, rewriting and using the definition $\Lambda = \gamma^2\mu^2 w^2 \sigma^2$, this is the same as (39) 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:

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

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 $\bar{\sigma}_1 = \sqrt{\frac{2\rho^2 \ln \vartheta^{-1}}{\left(\frac{\gamma\rho}{1+\gamma\rho}\right)^2 \left(\frac{2\left(1-\frac{\varphi}{\gamma}\right) \ln \vartheta^{-1}}{1+2 \ln \vartheta} + (1-\beta)\right)^2}}$ 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.:

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

Rearranging:

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

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 $\Lambda < \Lambda^*$, $\tilde{\beta} = \beta \vartheta e^{\frac{\Lambda}{2}} < 1$. \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 expressions for wages in steady state, using $\vartheta e^{\frac{\Lambda}{2}} = 1$ we have:

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

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

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

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

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

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 aggregate shocks ($z_t = 1$ and $\varepsilon_t = \varepsilon \forall t$). This problem can be written as:

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

s.t.

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

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

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

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

Let $M_{1,t}$ denote the multiplier on the date t aggregate Euler equation, $M_{2,t}$ that on the date t Phillips curve, $M_{3,t}$ that on the date t Σ recursion and $M_{4,t}$ that on the relationship between y_t, w_t and Π_t . The

necessary conditions for optimality are as follows.

First-order condition with respect to w_t :

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

FOC wrt y_t :

$$\begin{aligned} & -\gamma \mathbb{U}_t - \gamma M_{1,t} + \beta^{-1} M_{1,t-1} \left\{ \gamma - \varphi \gamma^2 \mu_t^2 w^2 \sigma^2 e^{2[\varphi(y_t - y)]} \right\} + M_{2,t-1} \left(\frac{w_t}{y_{t-1} w_{t-1}} \right) (\Pi_t - 1) \Pi_t \\ & - \beta M_{2,t} \left(\frac{y_{t+1} w_{t+1}}{y_t^2 w_t} \right) (\Pi_{t+1} - 1) \Pi_{t+1} + M_{3,t} \varphi \gamma^2 \mu_t^2 w^2 \sigma^2 e^{2\varphi(y_t - y)} + M_{4,t} = 0 \end{aligned} \quad (\text{D.6})$$

FOC wrt μ_t :

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

FOC wrt 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 wrt Π_t :

$$\left[\left(\frac{y_t w_t}{y_{t-1} w_{t-1}} \right) M_{2,t-1} - M_{2,t} \right] (2\Pi_t - 1) + \Psi M_{4,t} \frac{\rho \ln w_t + \bar{\xi}}{\left[1 + \gamma \rho + \frac{\Psi}{2} (\Pi_t - 1)^2 \right]^2} (\Pi_t - 1) = 0 \quad (\text{D.9})$$

FOC wrt τ_0^a :

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

FOC wrt 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 $M_{3,0} \neq 0$. In particular, if $\alpha = 0$, i.e., the planner is utilitarian, we have $\tau_0^a = 1$. 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 forwards, we get $M_{3,t}/\mathbb{U}_t = (1 - \tilde{\beta})^{-1} \neq 0$.

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

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

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

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

τ^w , the variables y_t, Π_t, μ_t, w_t remain at their steady state level absent aggregate 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 $m_i = M_i/\bar{U}$ 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:

$$\max_{\{w_t, y_t, \mu_t, \Sigma_t, \Pi_t\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t \left(\prod_{k=0}^{t-1} \zeta_k \right) \left\{ -\frac{1}{\gamma} (1 + \gamma \rho w_t) e^{-\gamma y_t \Sigma_t} \right\} \quad (\text{D.20})$$

s.t.

$$\gamma y_t = \gamma y_{t+1} - \ln \beta \vartheta - \ln \zeta_t + \ln \mu_{t+1} + \ln [\mu_t^{-1} - (1 + \gamma \rho w_t)] - \frac{\gamma^2 \mu_{t+1}^2 w^2 \sigma^2}{2} e^{2\varphi(y_{t+1} - y) + 2\zeta_{t+1}} \quad (\text{D.21})$$

$$(\Pi_t - 1) \Pi_t = \frac{\varepsilon_t}{\Psi} \left[1 - \frac{\varepsilon_t - 1}{\varepsilon_t} \frac{(1 - \tau^w) z_t}{(1 - \tau^*) w_t} \right] + \beta \left(\frac{z_t w_{t+1} y_{t+1}}{z_{t+1} w_t y_t} \right) (\Pi_{t+1} - 1) \Pi_{t+1} \quad (\text{D.22})$$

$$\ln \Sigma_t = \frac{\gamma^2 \mu_t^2 w^2 \sigma^2}{2} e^{2\varphi(y_t - y) + \zeta_t} + \ln [1 - \vartheta + \vartheta \Sigma_{t-1}] + \mathbb{I}(t=0) \ln \left[\frac{1 - \vartheta e^{\frac{\Lambda}{2}}}{1 - \vartheta e^{\left(\frac{\alpha}{\mu}\right)^2 \frac{\Lambda}{2} \left(\frac{\mu - \mu_0}{\mu}\right)^2}} \right] \quad (\text{D.23})$$

$$y_t = \frac{z_t (\rho \ln w_t + \bar{\xi})}{1 + \gamma \rho z_t + \frac{\Psi}{2} (\Pi_t - 1)^2} \quad (\text{D.24})$$

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

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

The optimal decisions satisfy:

FOC wrt w_t (multiplied through by w_t):

$$\begin{aligned}
& \mathbb{U}_t \frac{\gamma \rho w_t}{1 + \gamma \rho w_t} + \zeta_{t-1}^{-1} M_{2,t-1} \left(\frac{z_{t-1} y_t w_t}{z_t y_{t-1} w_{t-1}} \right) (\Pi_t - 1) \Pi_t - M_{1,t} \frac{\gamma \rho w_t}{\mu_t^{-1} - (1 + \gamma \rho w_t)} \\
& + M_{2,t} \left\{ \frac{\varepsilon_t - 1}{\Psi} \frac{(1 - \tau^w) z_t}{(1 - \tau^*) w_t} - \beta \left(\frac{z_t y_{t+1} w_{t+1}}{z_{t+1} y_t w_t} \right) (\Pi_{t+1} - 1) \Pi_{t+1} \right\} - \frac{M_{4,t}}{\gamma} z_t \frac{\gamma \rho}{1 + \gamma \rho z_t + \frac{\Psi}{2} (\Pi_t - 1)^2} = 0
\end{aligned} \tag{D.25}$$

FOC wrt y_t :

$$-\gamma \mathbb{U}_t - \gamma M_{1,t} + \beta^{-1} \zeta_{t-1}^{-1} M_{1,t-1} \left\{ \gamma - \varphi \gamma^2 \mu_t^2 w^2 \sigma^2 e^{2[\varphi(y_t-y)+\varsigma_t]} \right\} + \zeta_{t-1}^{-1} M_{2,t-1} \left(\frac{z_{t-1} w_t}{z_t y_{t-1} w_{t-1}} \right) (\Pi_t - 1) \Pi_t \tag{D.26}$$

$$-\beta M_{2,t} \left(\frac{z_t y_{t+1} w_{t+1}}{z_{t+1} y_t^2 w_t} \right) (\Pi_t - 1) \Pi_t + M_{3,t} \varphi \gamma^2 \mu_t^2 w^2 \sigma^2 e^{2\varphi(y_t-y)+2\varsigma_t} + M_{4,t} = 0 \tag{D.27}$$

FOC wrt μ_t :

$$\begin{aligned}
& -M_{1,t} \frac{\mu_t^{-1}}{\mu_t^{-1} - 1 - \gamma \rho w_t} + \beta^{-1} \zeta_{t-1}^{-1} M_{1,t-1} \left[1 - \gamma^2 \sigma^2 w^2 \mu_t^2 e^{2\varphi(y_t-y)+2\varsigma_t} \right] + M_{3,t} \gamma^2 \sigma^2 w^2 \mu_t^2 e^{2\varphi(y_t-y)+2\varsigma_t} \\
& - \mathbb{I}(t=0) M_{3,0} \frac{\vartheta e^{\frac{\Lambda}{2} \left(\frac{\alpha}{\mu}\right)^2 \left(\frac{\mu - \mu_0}{\mu}\right)^2}}{1 - \vartheta e^{\frac{\Lambda}{2} \left(\frac{\alpha - (1 - \tau_0^a) \mu_0}{\mu}\right)^2}} \Lambda \left(\frac{\alpha}{\mu} \right)^2 \left(\frac{\mu - \mu_0}{\mu} \right) \frac{\mu_0}{\mu} = 0
\end{aligned} \tag{D.28}$$

FOC wrt Σ_t :

$$\mathbb{U}_t - M_{3,t} + \beta \zeta_t M_{3,t+1} \frac{\vartheta \Sigma_t}{1 - \vartheta + \vartheta \Sigma_t} = 0 \quad (\text{D.29})$$

FOC wrt Π_t :

$$\zeta_{t-1}^{-1} M_{2,t-1} \left(\frac{z_{t-1} w_t y_t}{z_t w_{t-1} y_{t-1}} \right) (2\Pi_t - 1) - M_{2,t} (2\Pi_t - 1) + \Psi M_{4,t} (\Pi_t - 1) = 0 \quad (\text{D.30})$$

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 tax is optimally set to

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

Consequently, (D.3) becomes

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

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 variables is the same as that chosen by the utilitarian planner.

E Local approximation

E.1 Log-linearized dynamic equations

All hatted variables denote log-deviations of from steady state, except for the hatted multipliers which denote deviations in levels. In the baseline model with all four shocks, the log-linearized equations describing aggregate dynamics are:

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

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

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

$$\pi_t = \kappa (\hat{y}_t - \hat{y}_t^e) + \beta \pi_{t+1} + \frac{\varepsilon}{\Psi} \hat{\varepsilon}_t \quad (\text{E.4})$$

where $\kappa = \frac{\varepsilon(1+\gamma\rho)}{\Psi(\rho/y)}$. Using (E.2) and (E.3) to substitute out i_t and \hat{w}_t and using the fact that $\Omega = \frac{w-1}{1+\gamma\rho w}$ and $1 + \left(1 - \tilde{\beta}\right) \Omega = \frac{\Theta}{1-\Lambda}$, the IS equation (E.1) can be written as

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

Solving this equation forwards yields:

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

where

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

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 log-linearized Σ_t recursion is

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

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

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

where Γ_t is defined in (E.6). Then we can write the log-linearized Σ_t recursion as

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

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

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

Restricting attention to the case without demand shocks ($\hat{\zeta}_t = \hat{\varsigma}_t = 0$) as in our baseline model

$$\hat{\Sigma}_t = -\gamma y (1 - \tilde{\beta}) \Omega \left[\hat{y}_t - \underbrace{\left(\frac{1 + \Omega}{\Omega} \right) \frac{\Lambda}{1 - \tilde{\beta} (1 - \Lambda) \varrho_z} \frac{1}{1 + \rho/y}}_{=\kappa(\Omega)} \hat{y}_t^e \right] + \beta^{-1} \tilde{\beta} \hat{\Sigma}_{t-1}$$

where $\widehat{y}_t^c = \frac{1+\rho/y}{1+\gamma\rho} \widehat{z}_t$. When $\Omega \geq \Omega^c = \frac{\Lambda}{(1-\widetilde{\beta})(1-\Lambda)}$, we clearly have $\varkappa(\Omega) > 0$; we also have

$$\begin{aligned} \varkappa(\Omega) &= \left(\frac{1+\Omega}{\Omega} \right) \frac{\Lambda}{1-\widetilde{\beta}(1-\Lambda)} \frac{1}{\varrho_z} \frac{1}{1+\rho/y} \\ &\leq \varkappa(\Omega^c) \\ &= \frac{1-\widetilde{\beta}(1-\Lambda)}{1-\widetilde{\beta}(1-\Lambda)} \frac{1}{\varrho_z} \frac{1}{1+\rho/y} \\ &< \frac{1}{1+\rho/y} < 1 \end{aligned}$$

Thus, for $\Omega \geq \Omega^c$ we have $\varkappa(\Omega) \in (0, 1)$, as Lemma 1 claims.

E.2 Derivation of the Quadratic Loss function

As is well known, in the presence of a distorted steady state, maximizing a second-order approximation to the objective function (D.20) subject to first-order approximations of constraints (D.21)-(D.24), will not generally lead to a solution to the optimal policy problem which is accurate up to first-order. But following Benigno and Woodford (2005) and others, we obtain a valid linear-quadratic (LQ) approximation to the non-linear planning problem described in Appendix D.3 by using a second-order approximation of the constraints to eliminate the linear terms in the second-order approximations of the objective function.

Taking a second-order approximation to the planner's objective function \mathbb{W}_0 , we have:²⁶

$$\begin{aligned} \mathbb{W}_0 &\approx \frac{\mathbb{U}}{1-\beta} \\ &+ \mathbb{U} \sum_{t=0}^{\infty} \beta^t \left\{ \frac{\gamma\rho w}{1+\gamma\rho w} \widehat{w}_t - \gamma y \widehat{y}_t + \widehat{\Sigma}_t + \frac{1}{2} (\gamma y)^2 \widehat{y}_t^2 - \gamma y \widehat{y}_t \widehat{\Sigma}_t - \gamma y \frac{\gamma\rho(1+\Omega)}{1+\gamma\rho} \widehat{y}_t \widehat{w}_t + \frac{\gamma\rho(1+\Omega)}{1+\gamma\rho} \widehat{w}_t \widehat{\Sigma}_t \right\} \end{aligned} \quad (\text{E.7})$$

The second-order approximation to the IS curve at date t can be written as:

$$\begin{aligned} g_t^{\text{IS}} &= \gamma y \Theta \widehat{y}_{t+1} + (1-\Lambda) \widehat{\mu}_{t+1} - \frac{1}{\widetilde{\beta}} \widehat{\mu}_t - \left(\frac{1-\widetilde{\beta}}{\widetilde{\beta}} \right) \left(\frac{\gamma\rho w}{1+\gamma\rho w} \right) \widehat{w}_t - \gamma y \widehat{y}_t \\ &- (\gamma y)^2 \frac{(1-\Theta)^2}{\Lambda} \widehat{y}_{t+1}^2 - 2\gamma y (1-\Theta) \widehat{\mu}_{t+1} \widehat{y}_{t+1} - \left(\frac{1+\Lambda}{2} \right) \widehat{\mu}_{t+1}^2 + \frac{1}{2} \left(2 - \frac{1}{\widetilde{\beta}} \right) \frac{1}{\widetilde{\beta}} \widehat{\mu}_t^2 \\ &- \frac{1}{2} \left(\frac{1-\widetilde{\beta}}{\widetilde{\beta}} \right)^2 \left(\frac{\gamma\rho(1+\Omega)}{1+\gamma\rho} \right)^2 \widehat{w}_t^2 - \frac{(1-\widetilde{\beta})}{\widetilde{\beta}^2} \frac{\gamma\rho(1+\Omega)}{1+\gamma\rho} \widehat{w}_t \widehat{\mu}_t \end{aligned} \quad (\text{E.8})$$

where we have used $\frac{\vartheta}{R_t} = \mu_{t+1} [\mu_t^{-1} - (1+\gamma\rho w_t)]$ to eliminate R_t . Next, since the steady state multiplier on the Phillips curve $M_2 = 0$, we can skip taking a second-order approximation of the Phillips curve. So, we proceed by taking a second-order approximation of the Σ_t recursion, we have:

$$\begin{aligned} g_t^{\Sigma} &\approx \Lambda \widehat{\mu}_t + \gamma y (1-\Theta) \widehat{y}_t + \beta^{-1} \widetilde{\beta} \widehat{\Sigma}_{t-1} - \widehat{\Sigma}_t + \frac{1}{2} \widehat{\Sigma}_t^2 - \frac{1}{2} \left(\beta^{-1} \widetilde{\beta} \right)^2 \widehat{\Sigma}_{t-1}^2 \\ &+ (\gamma y)^2 \frac{(1-\Theta)^2}{\Lambda} \widehat{y}_t^2 + 2\gamma y (1-\Theta) \widehat{\mu}_t \widehat{y}_t + \frac{\Lambda}{2} \widehat{\mu}_t^2 + \mathbb{I}(t=0) \frac{\vartheta}{1-\vartheta} \left(\frac{\alpha}{\mu} \right)^2 \frac{\Lambda}{2} \widehat{\mu}_0^2 \end{aligned} \quad (\text{E.9})$$

²⁶This approximation is valid for all specifications of Pareto weights considered in Sections 3, 4 and 5.2.

Finally, we can write the second-order approximation of (D.24) as:

$$\mathbf{g}_t^y \approx y\hat{y}_t - \frac{y}{1+\gamma\rho}\hat{z}_t - \frac{\rho}{1+\gamma\rho}\hat{w}_t + \frac{1}{2}\frac{\rho}{1+\gamma\rho}\hat{w}_t^2 - \frac{\rho}{(1+\gamma\rho)^2}\hat{w}_t\hat{z}_t + \frac{1}{2}\frac{\Psi y}{1+\gamma\rho}\pi_t^2 \quad (\text{E.10})$$

Note that (E.8)-(E.10) equal 0 for any allocation satisfying the constraints up to second-order. Thus, we can use these equations together with the FOCs from the planner's problem absent shocks to eliminate first-order terms from the objective function (E.7). This yields the purely second-order approximation to (E.7):

$$\mathbb{W}_0 \approx \frac{\mathbb{U}}{1-\tilde{\beta}} + \mathbb{U} \sum_{t=0}^{\infty} \beta^t \tilde{\mathbb{U}}_t$$

where

$$\begin{aligned} \tilde{\mathbb{U}}_t = & \frac{1}{2} (\gamma y)^2 \hat{y}_t^2 - \gamma y \hat{y}_t \hat{\Sigma}_t - \gamma y \frac{\gamma \rho (1 + \Omega)}{1 + \gamma \rho} \hat{y}_t \hat{w}_t + \frac{\gamma \rho (1 + \Omega)}{1 + \gamma \rho} \hat{w}_t \hat{\Sigma}_t \\ & + m_1 \left\{ -\beta^{-1} (\gamma y)^2 \frac{(1 - \Theta)^2}{\Lambda} \hat{y}_t^2 - 2\beta^{-1} \gamma y (1 - \Theta) \hat{\mu}_t \hat{y}_t - \beta^{-1} \left(\frac{1 + \Lambda}{2} \right) \hat{\mu}_t^2 + \frac{1}{2} \left(2 - \frac{1}{\tilde{\beta}} \right) \frac{1}{\tilde{\beta}} \hat{\mu}_t^2 \right\} \\ & + m_1 \left\{ -\frac{1}{2} \left(\frac{1 - \tilde{\beta}}{\tilde{\beta}} \right)^2 \left(\frac{\gamma \rho (1 + \Omega)}{1 + \gamma \rho} \right)^2 \hat{w}_t^2 - \frac{(1 - \tilde{\beta})}{\tilde{\beta}^2} \frac{\gamma \rho (1 + \Omega)}{1 + \gamma \rho} \hat{w}_t \hat{\mu}_t \right\} \\ & + m_3 \left\{ \frac{1 - \beta^{-1} \tilde{\beta}^2}{2} \hat{\Sigma}_t^2 + (\gamma y)^2 \frac{(1 - \Theta)^2}{\Lambda} \hat{y}_t^2 + 2\gamma y (1 - \Theta) \hat{\mu}_t \hat{y}_t + \frac{\Lambda}{2} \hat{\mu}_t^2 + \mathbb{I}(t=0) \frac{\vartheta}{1 - \vartheta} \left(\frac{\alpha}{\mu} \right)^2 \frac{\Lambda}{2} \hat{\mu}_0^2 \right\} \\ & + m_{4y} \left\{ \frac{1}{2} \frac{\rho/y}{1 + \gamma \rho} \hat{w}_t^2 - \frac{\rho/y}{(1 + \gamma \rho)^2} \hat{w}_t \hat{z}_t + \frac{1}{2} \frac{\Psi}{1 + \gamma \rho} \pi_t^2 \right\} \end{aligned} \quad (\text{E.11})$$

where $m_i = M_i/\mathbb{U}$ denote the normalized steady state multipliers as above. Clearly maximizing \mathbb{W}_0 is equivalent to minimizing $\sum_{t=0}^{\infty} \beta^t \tilde{\mathbb{U}}_t$ since $\mathbb{U} < 0$.

Using the expressions derived above for steady state multipliers and substituting out for \hat{w}_t using $\hat{w}_t = \frac{1+\gamma\rho}{\rho/y}\hat{y}_t - \frac{1}{\rho/y}\hat{z}_t$ and $\hat{\mu}_t$ using $\hat{\mu}_t = \Gamma_z \hat{z}_t - \gamma y \left[1 + (1 - \tilde{\beta}) \Omega \right] \hat{y}_t$, we can obtain a loss function in \hat{y}_t, π_t, z_t and Σ_t (ignoring terms independent of policy) for $t > 0$:

$$\begin{aligned} \tilde{\mathbb{U}}_t = & \frac{1}{2} \gamma y \left[\left(\frac{y}{\rho} \right) \frac{(1 - \beta^{-1} \tilde{\beta}) (1 - \Lambda)}{1 - \beta^{-1} \tilde{\beta} (1 - \Lambda)} (1 + \Omega) \Upsilon(\Omega) + (\gamma y) (1 - \tilde{\beta}) \Omega^2 \right] \hat{y}_t^2 \\ & - \frac{\gamma y}{1 + \gamma \rho} \left[\frac{\gamma y (1 - \tilde{\beta}) \Omega}{1 - \tilde{\beta} \rho z (1 - \Lambda)} + \delta(\Omega) \Upsilon(\Omega) \left(\frac{y}{\rho} + 1 \right) \frac{(1 - \beta^{-1} \tilde{\beta}) (1 - \Lambda)}{1 - \beta^{-1} \tilde{\beta} (1 - \Lambda)} \right] (1 + \Omega) \hat{y}_t \hat{z}_t \\ & + \frac{1}{2} \left(\frac{1 - \beta^{-1} \tilde{\beta}^2}{1 - \tilde{\beta}} \right) \hat{\Sigma}_t^2 + \gamma y \Omega \hat{y}_t \hat{\Sigma}_t - \frac{\gamma y}{1 + \gamma \rho} (1 + \Omega) \hat{z}_t \hat{\Sigma}_t + \frac{1}{2} \frac{(1 - \beta^{-1} \tilde{\beta}) (1 - \Lambda)}{1 - \beta^{-1} \tilde{\beta} (1 - \Lambda)} (1 + \Omega) \frac{\Psi \gamma y}{1 + \gamma \rho} \pi_t^2 \end{aligned} \quad (\text{E.12})$$

where $\Upsilon(\Omega)$ and $\delta(\Omega)$ are given by

$$\Upsilon(\Omega) = 1 + \gamma\rho \frac{\Omega}{1+\Omega} \left\{ (1-\tilde{\beta}) \Omega \left(\frac{2}{\Lambda(1-\Lambda)} - 1 \right) - 1 \right\} \quad (\text{E.13})$$

$$\delta(\Omega) = \frac{1}{\Upsilon(\Omega)} \left[1 + \left(\frac{1+\Lambda}{1-\Lambda} \right) \frac{\gamma\rho(1-\tilde{\beta})\Omega}{1-\tilde{\beta}\varrho_z(1-\Lambda)} \frac{1}{1+\rho/y} \right] \quad (\text{E.14})$$

For $t = 0$ we have:

$$\begin{aligned} \tilde{\mathbb{U}}_0 &= \frac{1}{2}\gamma y \left[\left(\frac{y}{\rho} \right) \frac{(1-\beta^{-1}\tilde{\beta})(1-\Lambda)}{1-\beta^{-1}\tilde{\beta}(1-\Lambda)} (1+\Omega) \Upsilon_0(\Omega) + (\gamma y) (1-\tilde{\beta}) \Omega^2 \right] \hat{y}_0^2 \\ &\quad - \frac{\gamma y}{1+\gamma\rho} \left[\frac{\gamma y (1-\tilde{\beta}) \Omega}{1-\tilde{\beta}\varrho_z(1-\Lambda)} + \delta_0(\Omega) \Upsilon_0(\Omega) \left(\frac{y}{\rho} + 1 \right) \frac{(1-\beta^{-1}\tilde{\beta})(1-\Lambda)}{1-\beta^{-1}\tilde{\beta}(1-\Lambda)} \right] (1+\Omega) \hat{y}_0 \hat{z}_0 \\ &\quad + \frac{1}{2} \left(\frac{1-\beta^{-1}\tilde{\beta}^2}{1-\tilde{\beta}} \right) \hat{\Sigma}_0^2 + \gamma y \Omega \hat{y}_0 \hat{\Sigma}_0 - \frac{\gamma y}{1+\gamma\rho} (1+\Omega) \hat{z}_0 \hat{\Sigma}_0 + \frac{1}{2} \frac{(1-\beta^{-1}\tilde{\beta})(1-\Lambda)}{1-\beta^{-1}\tilde{\beta}(1-\Lambda)} (1+\Omega) \frac{\Psi\gamma y}{1+\gamma\rho} \pi_0^2 \end{aligned} \quad (\text{E.15})$$

where $\Upsilon_0(\Omega)$ and $\delta_0(\Omega)$ are given by

$$\Upsilon_0(\Omega) = \Upsilon(\Omega) + \left[1 + (1-\tilde{\beta}) \Omega \right] \mathbb{G} \quad (\text{E.16})$$

$$\delta_0(\Omega) = \frac{\Upsilon(\Omega)}{\Upsilon_0(\Omega)} \delta(\Omega) + \frac{1}{\Upsilon_0(\Omega)} \frac{(1-\tilde{\beta})(1+\Omega)}{1-\tilde{\beta}\varrho_z(1-\Lambda)} \frac{1}{1+\rho/y} \mathbb{G} \quad (\text{E.17})$$

where

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

Notice that when $\alpha = 0$, $\mathbb{G} = 0$, $\Upsilon(\Omega) = \Upsilon_0(\Omega)$ and $\delta(\Omega) = \delta_0(\Omega)$, and there is no difference between the two expressions above. In principle, one could derive optimal policy by minimizing $\sum_{t=0}^{\infty} \beta^t \tilde{\mathbb{U}}_t$ subject to the subject to the linearized Phillips curve (19) and the linearized Σ recursion (26). However, it is useful to use the the linearized Σ recursion (26) to substitute out for $\hat{\Sigma}_t$ and obtain a loss function purely in terms of \hat{y}_t, π_t and \hat{z}_t . The terms involving $\hat{\Sigma}_t$ in the objective function can be written as

$$L_{\Sigma} = \sum_{t=0}^{\infty} \beta^t \left\{ \frac{1}{2} \left(\frac{1-\beta^{-1}\tilde{\beta}^2}{1-\tilde{\beta}} \right) \hat{\Sigma}_t^2 + \gamma y \Omega \hat{y}_t \hat{\Sigma}_t - \frac{\gamma y}{1+\gamma\rho} (1+\Omega) \hat{z}_t \hat{\Sigma}_t \right\} \quad (\text{E.18})$$

Next, solving (26) back to date -1 , we have

$$\hat{\Sigma}_t = -\gamma y (1-\tilde{\beta}) \Omega \sum_{k=0}^t \left(\frac{\tilde{\beta}}{\beta} \right)^{t-k} \hat{y}_k + \Lambda \frac{(1+\Omega)(1-\tilde{\beta})}{1-\tilde{\beta}\varrho_z(1-\Lambda)} \frac{\gamma y}{1+\gamma\rho} \sum_{k=0}^t \left(\frac{\tilde{\beta}}{\beta} \right)^{t-k} \hat{z}_k + \left(\frac{\tilde{\beta}}{\beta} \right)^{t+1} \hat{\Sigma}_{-1} \quad (\text{E.19})$$

Substituting (E.19) into (E.18) yields, after some algebra

$$L_{\Sigma} = -\frac{1}{2} (\gamma y)^2 (1 - \tilde{\beta}) \Omega^2 \sum_{t=0}^{\infty} \beta^t \hat{y}_t^2 + \frac{(\gamma y)^2}{1 + \gamma \rho} (1 + \Omega) \Omega \frac{(1 - \tilde{\beta})}{1 - \tilde{\beta} \rho_z (1 - \Lambda)} \sum_{t=0}^{\infty} \beta^t \hat{y}_t \hat{z}_t$$

Substituting this expression into (E.12) and (E.15) yields the expression

$$\begin{aligned} \tilde{U}_t &= \frac{\gamma y}{2} \left[\left(\frac{y}{\rho} \right) \frac{(1 - \beta^{-1} \tilde{\beta}) (1 - \Lambda)}{1 - \beta^{-1} \tilde{\beta} (1 - \Lambda)} (1 + \Omega) \Upsilon_0(\Omega) \right] \hat{y}_0^2 \\ &+ \frac{\gamma y}{2} \left[\left(\frac{y}{\rho} \right) \frac{(1 - \beta^{-1} \tilde{\beta}) (1 - \Lambda)}{1 - \beta^{-1} \tilde{\beta} (1 - \Lambda)} (1 + \Omega) \Upsilon(\Omega) \right] \sum_{t=1}^{\infty} \beta^t \hat{y}_t^2 \\ &- \frac{\gamma y}{1 + \gamma \rho} \left[\delta_0(\Omega) \Upsilon_0(\Omega) \left(\frac{y}{\rho} + 1 \right) \frac{(1 - \beta^{-1} \tilde{\beta}) (1 - \Lambda)}{1 - \beta^{-1} \tilde{\beta} (1 - \Lambda)} \right] (1 + \Omega) \hat{y}_0 \hat{z}_0 \\ &- \frac{\gamma y}{1 + \gamma \rho} \left[\delta(\Omega) \Upsilon(\Omega) \left(\frac{y}{\rho} + 1 \right) \frac{(1 - \beta^{-1} \tilde{\beta}) (1 - \Lambda)}{1 - \beta^{-1} \tilde{\beta} (1 - \Lambda)} \right] (1 + \Omega) \sum_{t=1}^{\infty} \beta^t \hat{y}_t \hat{z}_t \\ &+ \frac{\gamma y}{2} \frac{(1 - \beta^{-1} \tilde{\beta}) (1 - \Lambda)}{1 - \beta^{-1} \tilde{\beta} (1 - \Lambda)} (1 + \Omega) \frac{\Psi}{1 + \gamma \rho} \sum_{t=0}^{\infty} \beta^t \pi_t^2 \end{aligned} \quad (\text{E.20})$$

Dividing by $\frac{(1 - \beta^{-1} \tilde{\beta})(1 - \Lambda)}{1 - \beta^{-1} \tilde{\beta} (1 - \Lambda)} (1 + \Omega) \gamma y \left(\frac{y}{\rho} \right)$, using the fact that $\varepsilon/\kappa = \frac{\Psi(\rho/y)}{1 + \gamma \rho}$ and using the definition of $\hat{y}_t^e = \frac{1 + \rho/y}{1 + \gamma \rho} \hat{z}_t$, yields the objective function in the main text in Proposition 9

$$\frac{1}{2} \left\{ \Upsilon_0(\Omega) \left(\hat{y}_0 - \delta_0(\Omega) \hat{y}_0^e \right)^2 + \frac{\varepsilon}{\kappa} \pi_0^2 \right\} + \frac{1}{2} \sum_{t=1}^{\infty} \beta^t \left\{ \Upsilon(\Omega) \left(\hat{y}_t - \delta(\Omega) \hat{y}_t^e \right)^2 + \frac{\varepsilon}{\kappa} \pi_t^2 \right\} \quad (\text{E.21})$$

For the utilitarian planner, $\Upsilon_0(\Omega) = \Upsilon(\Omega)$ and $\delta_0(\Omega) = \delta(\Omega)$ and the expression simplifies to the expression in Proposition 3:

$$\frac{1}{2} \sum_{t=0}^{\infty} \beta^t \left\{ \Upsilon(\Omega) \left(\hat{y}_t - \delta(\Omega) \hat{y}_t^e \right)^2 + \frac{\varepsilon}{\kappa} \pi_t^2 \right\} \quad (\text{E.22})$$

The optimal policy problem can now simply be specified as minimizing (E.21) subject to the linearized Phillips curve (24). In Lagrangian form:

$$\begin{aligned} \mathcal{L} &= \frac{1}{2} \left\{ \Upsilon_0(\Omega) \left(\hat{y}_0 - \delta_0(\Omega) \hat{y}_0^e \right)^2 + \frac{\varepsilon}{\kappa} \pi_0^2 \right\} + \frac{1}{2} \sum_{t=1}^{\infty} \beta^t \left\{ \Upsilon(\Omega) \left(\hat{y}_t - \delta(\Omega) \hat{y}_t^e \right)^2 + \frac{\varepsilon}{\kappa} \pi_t^2 \right\} \\ &+ \sum_{t=0}^{\infty} \beta^t F_t \left\{ \beta \pi_{t+1} + \kappa (\hat{y}_t - \hat{y}_t^e) + \frac{\varepsilon}{\Psi} \hat{\varepsilon}_t - \pi_t \right\} \end{aligned}$$

The FOC w.r.t. \hat{y}_t can be written as:

$$\begin{aligned}\Upsilon_0(\Omega) (\hat{y}_0 - \delta_0(\Omega)\hat{y}_0^e) + \kappa F_0 &= 0 && \text{for } t = 0 \\ \Upsilon(\Omega) (\hat{y}_t - \delta(\Omega)\hat{y}_t^e) + \kappa F_t &= 0 && \text{for } t > 0\end{aligned}$$

The FOC w.r.t. π_t can be written as

$$\frac{\varepsilon}{\kappa}\pi_t - F_t + F_{t-1} = 0 \quad \Leftrightarrow \quad F_t = \frac{\varepsilon}{\kappa}\hat{p}_t$$

where $\kappa = \frac{\varepsilon}{\Psi} \frac{1+\gamma\rho}{\rho/y}$. Combining the two FOCs we can derive the target criterion:

$$\begin{aligned}\hat{y}_0 - \delta_0(\Omega)\hat{y}_0^e + \frac{\varepsilon}{\Upsilon_0(\Omega)}\hat{p}_0 &= 0 && \text{for } t = 0 \\ \hat{y}_t - \delta(\Omega)\hat{y}_t^e + \frac{\varepsilon}{\Upsilon(\Omega)}\hat{p}_t &= 0 && \text{for } t > 0\end{aligned}$$

E.3 Properties of loss function weights

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

Proof.

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

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:

$$\delta(\Omega) = \frac{1 + \Omega + (\Omega + \Omega^2) \frac{\gamma\rho(1-\tilde{\beta})}{1-\tilde{\beta}\rho_z(1-\Lambda)} \frac{1}{1+(\rho/y)} \left[\frac{1+\Lambda}{1-\Lambda} \right]}{1 + (1-\gamma\rho)\Omega + \gamma\rho\Omega^2 \left(\frac{2}{\Lambda(1-\Lambda)} - 1 \right) (1-\tilde{\beta})} \quad (\text{E.23})$$

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

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

This expression can be simplified to yield:

$$1 + \frac{(1 - \tilde{\beta})}{1 - \tilde{\beta}\rho_z(1 - \Lambda)} \frac{y}{\rho + y} \left(\frac{1 + \Lambda}{1 - \Lambda} \right) < \Omega (1 - \tilde{\beta}) \left[\left(\frac{2}{\Lambda(1 - \Lambda)} - 1 \right) - \frac{1}{1 - \tilde{\beta}\rho_z(1 - \Lambda)} \frac{y}{\rho + y} \left(\frac{1 + \Lambda}{1 - \Lambda} \right) \right] \quad (\text{E.24})$$

First, we show that the term in the square brackets on the RHS of (E.24) 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 is 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.24) is positive. Next, to show that (E.24) 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.24), 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. □

Claim 3. $\Upsilon(0) = \delta(0) = 1$ when $\alpha = 0$.

Proof. True by inspection of equations (E.13), (E.14). □

Claim 4. $\Upsilon_0(\Omega) > \Upsilon(\Omega)$ when $\alpha \neq 0$

Proof. The claim that $\Upsilon_0 > \Upsilon$ is true by inspection of equations (E.16) since $\mathbb{G} > 0$ for $\Omega \geq \Omega^c > 0$. □

Claim 5. $\Upsilon_0(\Omega)$ is increasing in α for $\alpha > 0$.

Proof. Substituting the definition of \mathbb{G} into (E.16):

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

This is clearly increasing in α for $\alpha > 0$. □

Claim 6. $\varpi(\Omega) \in (0, 1)$ for $\Omega \geq \Omega^c$

Proof. Define:

$$\varpi(\Omega) = \frac{\delta(\Omega) - \varkappa(\Omega)}{1 - \varkappa(\Omega)}$$

where $\varkappa(\Omega) = \left(\frac{1+\Omega}{\Omega}\right) \frac{\Lambda}{1-\tilde{\beta}(1-\Lambda)\varrho_z} \frac{1}{1+\rho/y} \in (0, 1)$ for $\Omega \geq \Omega^c$. Recall that Claim 2 above showed that $\delta(\Omega) < 1$ for $\Omega \geq \Omega^c$. Thus, we have $\varpi(\Omega) < 1$ for $\Omega \geq \Omega^c$. It remains to show that $\varpi(\Omega) > 0$, i.e., $\delta(\Omega) > \varkappa(\Omega)$. Recall from (E.14) that $\delta(\Omega)$ is given by:

$$\begin{aligned} \delta(\Omega) &= \frac{1}{\Upsilon(\Omega)} \left[1 + \left(\frac{1+\Lambda}{1-\Lambda}\right) \frac{\gamma\rho(1-\tilde{\beta})\Omega}{1-\tilde{\beta}\varrho_z(1-\Lambda)} \frac{1}{1+\rho/y} \right] \\ &= \frac{1}{\Upsilon(\Omega)} \left[1 + \frac{\gamma\rho}{1+\Omega^{-1}} \left(\frac{1+\Lambda}{1-\Lambda}\right) \frac{(1-\tilde{\beta})\Omega}{\Lambda} \varkappa(\Omega) \right] \end{aligned}$$

So we have:

$$\begin{aligned} \delta(\Omega) - \varkappa(\Omega) &= \frac{1}{\Upsilon(\Omega)} \left[1 + \frac{\gamma\rho}{1+\Omega^{-1}} \left(\frac{1+\Lambda}{1-\Lambda}\right) \frac{(1-\tilde{\beta})\Omega}{\Lambda} \varkappa(\Omega) - \varkappa(\Omega)\Upsilon(\Omega) \right] \\ &= \frac{1}{\Upsilon(\Omega)} \left(1 - \frac{\gamma\rho(1-\tilde{\beta})(1-\Lambda)\Omega + (1+\Omega^{-1})\Lambda - \gamma\rho\Lambda}{1-\tilde{\beta}(1-\Lambda)\varrho_z} \frac{1}{1+\rho/y} \right), \end{aligned}$$

where we have used the definitions of $\varkappa(\Omega)$ and $\Upsilon(\Omega)$. Since $\Upsilon(\Omega) > 0$, the expression above is positive if

$$1 \geq \frac{\gamma\rho(1-\tilde{\beta})(1-\Lambda)\Omega + (1+\Omega^{-1})\Lambda - \gamma\rho\Lambda}{1-\tilde{\beta}(1-\Lambda)\varrho_z} \frac{1}{1+\rho/y}$$

or

$$(1+\rho/y) \left[1 - \tilde{\beta}\varrho_z(1-\Lambda) \right] > \Lambda(1+\Omega^{-1}) - \gamma\rho\Lambda + \gamma\rho(1-\tilde{\beta})(1-\Lambda)\Omega \equiv \beth(\Omega)$$

Clearly, $\beth(\Omega)$ is a convex function of Ω . Since $\Omega = \frac{w-1}{1+\gamma\rho w} < \lim_{w \rightarrow \infty} \frac{w-1}{1+\gamma\rho w} = \frac{1}{\gamma\rho}$, Ω is contained on the interval $[\Omega^c, (\gamma\rho)^{-1}]$. Thus,

$$\beth(\Omega) \leq \max \left\{ \beth(\Omega^c), \beth\left(\frac{1}{\gamma\rho}\right) \right\}$$

where $\beth(\Omega^c) = \beth\left(\frac{1}{\gamma\rho}\right) = 1 - \tilde{\beta}(1-\Lambda)$. Thus, $\beth(\Omega) \leq 1 - \tilde{\beta}(1-\Lambda)$. Clearly, we have

$$(1+\rho/y) \left[1 - \tilde{\beta}\varrho_z(1-\Lambda) \right] > 1 - \tilde{\beta}(1-\Lambda) \geq \beth(\Omega),$$

since $\rho/y > 0$ and $\varrho_z \in [0, 1)$. Thus, $\delta(\Omega) - \varkappa(\Omega) > 0$ and $\varpi(\Omega) > 0$ when $\Omega \geq \Omega^c$. \square

E.4 Deriving the target-criterion allowing for demand shocks

To derive a more general target criterion which allows for demand shocks in addition to aggregate productivity and markup shocks, we proceed by linearizing the first-order conditions of the non-linear planner's

problem rather than adopting an LQ approach. Linearizing the first-order conditions (D.25)-(D.30) and constraints (D.1)-(D.4) around the steady state described in Appendix D.2 yields the following FOC wrt w :

$$\begin{aligned}
& -(\gamma y)(1+\Omega)\hat{y}_t + (1+\Omega)\hat{\Sigma}_t - \left(\frac{1-\tilde{\beta}}{\tilde{\beta}}\right)(1+\Omega)\hat{m}_{1,t} - m_1 \left(\frac{1-\tilde{\beta}}{\tilde{\beta}}\right)^2 \frac{\gamma\rho}{1+\gamma\rho}(1+\Omega)^2\hat{w}_t \\
& - \left(\frac{1-\tilde{\beta}}{\tilde{\beta}^2}\right)(1+\Omega)m_1\hat{\mu}_t + \left(\frac{1+\gamma\rho}{\gamma\rho}\right)\frac{\varepsilon}{\Psi}\hat{m}_{2,t} - \frac{\hat{m}_{4,t}}{\gamma} + \frac{m_4}{\gamma}\hat{w}_t - \frac{m_4}{\gamma}\frac{1}{1+\gamma\rho}\hat{z}_t = 0 \quad (\text{E.25})
\end{aligned}$$

FOC wrt y :

$$\begin{aligned}
& -\frac{\gamma\rho(1+\Omega)}{1+\gamma\rho}\hat{w}_t + (\gamma y) \left[1 + 2\frac{(1-\Theta)^2}{\Lambda} \left(m_3 - \frac{m_1}{\beta} \right) \right] \hat{y}_t - \hat{\Sigma}_t - \hat{m}_{1,t} + \frac{\Theta}{\beta} (\hat{m}_{1,t-1} - m_1\hat{\zeta}_{t-1}) \\
& + 2(1-\Theta) \left(m_3 - \frac{m_1}{\beta} \right) \hat{\mu}_t + (1-\Theta)\hat{m}_{3,t} + \frac{\hat{m}_{4,t}}{\gamma} + 2(1-\Theta) \left(m_3 - \frac{m_1}{\beta} \right) \hat{\zeta}_t = 0 \quad (\text{E.26})
\end{aligned}$$

FOC wrt π :

$$\Delta\hat{m}_{2,t} = \frac{(1-\beta^{-1}\tilde{\beta})(1-\Lambda)}{1-\beta^{-1}\tilde{\beta}(1-\Lambda)}(1+\Omega)\frac{(\gamma y)\Psi}{1+\gamma\rho}\pi_t \quad \Rightarrow \quad \hat{m}_{2,t} = \frac{(1-\beta^{-1}\tilde{\beta})(1-\Lambda)}{1-\beta^{-1}\tilde{\beta}(1-\Lambda)}(1+\Omega)\frac{(\gamma y)\Psi}{1+\gamma\rho}\hat{p}_t \quad (\text{E.27})$$

FOC wrt μ :

$$\begin{aligned}
& -\left(\frac{1-\tilde{\beta}}{\tilde{\beta}^2}\right)\frac{\gamma\rho(1+\Omega)}{1+\gamma\rho}m_1\hat{w}_t + \left[2\Lambda \left(m_3 - \frac{m_1}{\beta} \right) - \frac{1-\tilde{\beta}}{\tilde{\beta}^2}m_1 \right] \hat{\mu}_t + \Lambda\hat{m}_{3,t} \\
& + 2(\gamma y)(1-\Theta) \left(m_3 - \frac{m_1}{\beta} \right) \hat{y}_t - \frac{1}{\tilde{\beta}} \left(\hat{m}_{1,t} - \frac{\tilde{\beta}}{\beta}(1-\Lambda)(\hat{m}_{1,t-1} - m_1\hat{\zeta}_{t-1}) \right) + 2\Lambda \left(m_3 - \frac{m_1}{\beta} \right) \hat{\zeta}_t \\
& + \mathbb{I}(t=0)\Lambda \left(\frac{\alpha}{\mu} \right)^2 m_3 \left(\frac{\vartheta}{1-\vartheta} \right) \hat{\mu}_0 = 0 \quad (\text{E.28})
\end{aligned}$$

FOC wrt Σ :

$$\frac{\gamma\rho}{1+\gamma\rho}(1+\Omega)\hat{w}_t - (\gamma y)\hat{y}_t - \hat{m}_{3,t} + \tilde{\beta}\hat{m}_{3,t+1} + \frac{1-\beta^{-1}\tilde{\beta}^2}{1-\tilde{\beta}}\hat{\Sigma}_t + \tilde{\beta}m_3\hat{\zeta}_t = 0 \quad (\text{E.29})$$

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

E.4.1 Deriving the target criterion

Add the FOC wrt w (E.25) to the FOC wrt y (E.26) to obtain:

$$\begin{aligned}
& -(\gamma y)(1+\Omega)\widehat{y}_t + \Omega\widehat{\Sigma}_t - \frac{\gamma\rho}{1+\gamma\rho}(1+\Omega)\left[m_1\left(\frac{1-\widetilde{\beta}}{\widetilde{\beta}}\right)^2(1+\Omega)+1\right]\widehat{w}_t + \frac{m_4}{\gamma}\widehat{w}_t \\
& + \left(\frac{1+\gamma\rho}{\gamma\rho}\right)\frac{\varepsilon}{\Psi}\widehat{m}_{2,t} - \frac{m_4}{\gamma}\frac{1}{1+\gamma\rho}\widehat{z}_t + \frac{\Theta}{\beta}\left(\widehat{m}_{1,t-1} - m_1\widehat{\zeta}_{t-1}\right) - \left[\left(\frac{1-\widetilde{\beta}}{\widetilde{\beta}}\right)(1+\Omega)+1\right]\widehat{m}_{1,t} \\
& + (\gamma y)\left[1+2\frac{(1-\Theta)^2}{\Lambda}\left(m_3 - \frac{m_1}{\beta}\right)\right]\widehat{y}_t - \left[\left(\frac{1-\widetilde{\beta}}{\widetilde{\beta}^2}\right)(1+\Omega)m_1 - 2(1-\Theta)\left(m_3 - \frac{m_1}{\beta}\right)\right]\widehat{\mu}_t \\
& + 2(1-\Theta)\left(m_3 - \frac{m_1}{\beta}\right)\widehat{\varsigma}_t + (1-\Theta)\widehat{m}_{3,t} \tag{E.30}
\end{aligned}$$

Combine with (E.28):

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

Next, use the GDP definition (E.3) to substitute out for \widehat{w}_t :

$$\begin{aligned}
& (\gamma y)\left\{-\Omega+2\frac{(1-\Theta)^2}{\Lambda}\left(m_3 - \frac{m_1}{\beta}\right) - 2(1-\Theta)\left[1+\Omega(1-\widetilde{\beta})\right]\left(m_3 - \frac{m_1}{\beta}\right)\right\}\widehat{y}_t + \Omega\widehat{\Sigma}_t \\
& + \left\{m_1\left(\frac{1-\widetilde{\beta}}{\widetilde{\beta}^2}\right)\frac{\gamma\rho(1+\Omega)}{1+\gamma\rho}\left[1+\Omega(1-\widetilde{\beta})\right] - \frac{\gamma\rho(1+\Omega)}{1+\gamma\rho}\left[m_1\left(\frac{1-\widetilde{\beta}}{\widetilde{\beta}}\right)^2(1+\Omega)+1\right] + \frac{m_4}{\gamma}\right\}\frac{1+\gamma\rho}{\rho/y}\widehat{y}_t \\
& - \left\{m_1\left(\frac{1-\widetilde{\beta}}{\widetilde{\beta}^2}\right)\frac{\gamma\rho(1+\Omega)}{1+\gamma\rho}\left[1+\Omega(1-\widetilde{\beta})\right] - \frac{\gamma\rho(1+\Omega)}{1+\gamma\rho}\left[m_1\left(\frac{1-\widetilde{\beta}}{\widetilde{\beta}}\right)^2(1+\Omega)+1\right] + \frac{m_4}{\gamma}\right\}\frac{1}{\rho/y}\widehat{z}_t \\
& - (1-\widetilde{\beta})\left[2\left(m_3 - \frac{m_1}{\beta}\right) + \frac{1}{\beta}m_1\right]\Omega\widehat{\mu}_t - (1-\widetilde{\beta})\Omega\widehat{m}_{3,t} + \left(\frac{1+\gamma\rho}{\gamma\rho}\right)\frac{\varepsilon}{\Psi}\widehat{m}_{2,t} - \frac{m_4}{\gamma}\frac{1}{1+\gamma\rho}\widehat{z}_t \\
& - 2(1-\widetilde{\beta})\Omega\left(m_3 - \frac{m_1}{\beta}\right)\widehat{\varsigma}_t - \mathbb{I}(t=0)\Lambda\left(\frac{\alpha}{\mu}\right)^2\left[1+\Omega(1-\widetilde{\beta})\right]m_3\left(\frac{\vartheta}{1-\vartheta}\right)\widehat{\mu}_0 = 0 \tag{E.31}
\end{aligned}$$

Next, using (E.5) to substitute out for $\widehat{\mu}_t$, using (E.27) to eliminate $\widehat{m}_{2,t}$ and using the definitions of m_1, m_3

and m_4 , (E.31) becomes

$$\begin{aligned}
& (\gamma y) \Omega \left\{ -1 - 2 \left(\frac{1 - \Theta}{\Lambda} \right) \left[\frac{1 - \beta^{-1} \tilde{\beta}}{1 - \beta^{-1} \tilde{\beta} (1 - \Lambda)} \right] + \left[\frac{2(1 - \beta^{-1} \tilde{\beta}) + \Lambda}{1 - \beta^{-1} \tilde{\beta} (1 - \Lambda)} \right] \frac{\Theta}{1 - \Lambda} \right\} \hat{y}_t + \Omega \hat{\Sigma}_t \\
& \quad + \frac{1}{\rho/y} \frac{(1 + \Omega)(1 - \beta^{-1} \tilde{\beta})(1 - \Lambda)}{1 - \beta^{-1} \tilde{\beta} (1 - \Lambda)} \left(\hat{y}_t - \frac{1 + \rho/y}{1 + \gamma \rho} \hat{z}_t \right) \\
& - \Omega \left[\frac{2(1 - \beta^{-1} \tilde{\beta}) + \Lambda}{1 - \beta^{-1} \tilde{\beta} (1 - \Lambda)} \right] \Gamma_t - (1 - \tilde{\beta}) \Omega \hat{m}_{3,t} - 2\Omega \left[\frac{1 - \beta^{-1} \tilde{\beta}}{1 - \beta^{-1} \tilde{\beta} (1 - \Lambda)} \right] \hat{\varsigma}_t \\
& \quad + \mathbb{I}(t=0) (\gamma y) \Lambda \left(\frac{\alpha}{\mu} \right)^2 \left[1 + (1 - \tilde{\beta}) \Omega \right]^2 \frac{1}{1 - \tilde{\beta}} \left(\frac{\vartheta}{1 - \vartheta} \right) \hat{y}_0 \\
& - \mathbb{I}(t=0) \Lambda \left(\frac{\alpha}{\mu} \right)^2 \left[1 + \Omega (1 - \tilde{\beta}) \right] \frac{1}{1 - \tilde{\beta}} \left(\frac{\vartheta}{1 - \vartheta} \right) \Gamma_0 + \left(\frac{\varepsilon y}{\rho} \right) \frac{(1 - \beta^{-1} \tilde{\beta})(1 - \Lambda)}{1 - \beta^{-1} \tilde{\beta} (1 - \Lambda)} (1 + \Omega) \hat{p}_t = 0
\end{aligned} \tag{E.32}$$

Guess that

$$\hat{m}_{3,t} = \frac{1}{1 - \tilde{\beta}} \hat{\Sigma}_t + \gamma y \Omega \hat{y}_t + a_z \hat{z}_t + a_\zeta \hat{\zeta}_t + a_\varsigma \hat{\varsigma}_t \tag{E.33}$$

and use this in (E.29) with \hat{w}_t substituted out using the definition of GDP:

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

using the fact that $\hat{z}_{t+1} = \varrho_z \hat{z}_t$, $\hat{\zeta}_{t+1} = \varrho_\zeta \hat{\zeta}_t$ and $\hat{\varsigma}_{t+1} = \varrho_\varsigma \hat{\varsigma}_t$. Using the expression for $\hat{\Sigma}_{t+1}$ in (26) in the equation above, we have

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

which implies that a_z, a_ζ and a_ς must satisfy:

$$a_z = -\frac{\gamma y}{1 + \gamma \rho} \left[\frac{1 + \Omega}{1 - \tilde{\beta} \varrho_z (1 - \Lambda)} \right] \quad (\text{E.34})$$

$$a_\zeta = \frac{\tilde{\beta}}{1 - \tilde{\beta}} \left[\frac{1}{1 - \tilde{\beta} (1 - \Lambda) \varrho_\zeta} \right] \quad (\text{E.35})$$

$$a_\varsigma = \frac{1}{1 - \tilde{\beta}} \left[\frac{\tilde{\beta} \varrho_\varsigma \Lambda}{1 - \tilde{\beta} (1 - \Lambda) \varrho_\varsigma} \right] \quad (\text{E.36})$$

Using the expression (E.33) for $\hat{m}_{3,t}$ in (E.32):

$$\hat{y}_0 - \delta_0(\Omega) \left(\frac{1 + \rho/y}{1 + \gamma \rho} \right) \hat{z}_t + \chi_0(\Omega) \hat{\zeta}_t - \Xi_0(\Omega) \hat{\varsigma}_t + \frac{\varepsilon}{\Upsilon_0(\Omega)} \hat{p}_t = 0 \quad (\text{E.37})$$

$$\hat{y}_t - \delta(\Omega) \left(\frac{1 + \rho/y}{1 + \gamma \rho} \right) \hat{z}_t + \chi(\Omega) \hat{\zeta}_t - \Xi(\Omega) \hat{\varsigma}_t + \frac{\varepsilon}{\Upsilon(\Omega)} \hat{p}_t = 0 \quad \text{for } t > 0 \quad (\text{E.38})$$

where $\Upsilon(\Omega), \delta(\Omega), \Upsilon_0(\Omega)$ and $\delta_{t0}(\Omega)$ are the same as in (E.13), (E.14), (E.16) and (E.17) and

$$\chi(\Omega) = \frac{1}{\Upsilon(\Omega)} \frac{\Omega}{1 + \Omega} \left(\frac{1 + \Lambda}{1 - \Lambda} \right) \left[\frac{\tilde{\beta}(\rho/y)}{1 - \tilde{\beta} \varrho_\zeta (1 - \Lambda)} \right] \quad (\text{E.39})$$

$$\Xi(\Omega) = \frac{1}{\Upsilon(\Omega)} (\rho/y) \frac{\Omega}{1 + \Omega} \frac{2(1 - \tilde{\beta} \varrho_\varsigma) + \tilde{\beta} \varrho_\varsigma \Lambda (1 - \Lambda)}{(1 - \Lambda) [1 - \tilde{\beta} \varrho_\varsigma (1 - \Lambda)]} \quad (\text{E.40})$$

$$\chi_0(\Omega) = \frac{\Upsilon(\Omega)}{\Upsilon_0(\Omega)} \chi(\Omega) + \frac{1}{\Upsilon_0(\Omega)} \frac{1}{\gamma y} \frac{\tilde{\beta}}{1 - \tilde{\beta} \varrho_\zeta (1 - \Lambda)} \mathbb{G} \quad (\text{E.41})$$

$$\Xi_0(\Omega) = \frac{\Upsilon(\Omega)}{\Upsilon_0(\Omega)} \Xi(\Omega) - \frac{1}{\Upsilon_0(\Omega)} \frac{1}{\gamma y} \frac{\tilde{\beta} \varrho_\varsigma \Lambda}{1 - \tilde{\beta} \varrho_\varsigma (1 - \Lambda)} \mathbb{G} \quad (\text{E.42})$$

where

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

Note that in the baseline with the utilitarian planner ($\alpha = 0$), we have $\mathbb{G} = 0$ and $\Xi_0(\Omega) = \Xi(\Omega)$ and $\chi(\Omega) = \chi_0(\Omega)$. 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 9 for the non-utilitarian planner (again, setting $\hat{\zeta}_t = \hat{\varsigma}_t = 0$), i.e., it yields the same target criterion as the LQ approach. It can also be specialized to yield Proposition 10 for demand shocks (setting $\alpha = 0$ and $\hat{z}_t = \hat{e}_t = 0$).

Claim 7. $\chi(\Omega) > 0$ 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 8. $\Xi(\Omega) > 0$ with countercyclical risk.

Proof. For $\Omega \geq \Omega^c > 0$, it is clear from the expression for $\Xi(\Omega)$ that $\Xi(\Omega) > 0$. \square

F Optimal Dynamics

As shown in Appendix E.4.1, the dynamics of x_t and π_t are given by the target criterion (30):

$$x_t - x_{t-1} + \bar{\varepsilon}\pi_t = 0 \quad (\text{F.1})$$

and the Phillips curve

$$\pi_t = \beta\pi_{t+1} + \kappa \left(x_t - [1 - \delta(\Omega)] \frac{1 + (\rho/y)}{1 + \gamma\rho} \hat{z}_t - \chi(\Omega)\hat{\zeta}_t + \Xi(\Omega)\hat{\varsigma}_t + \frac{\rho/y}{1 + \gamma\rho} \hat{\varepsilon}_t \right) \quad (\text{F.2})$$

where

$$x_t = \hat{y}_t - \delta(\Omega) \left(\frac{1 + \rho/y}{1 + \gamma\rho} \right) \hat{z}_t + \chi(\Omega)\hat{\zeta}_t - \Xi(\Omega)\hat{\varsigma}_t,$$

$\bar{\varepsilon} = \varepsilon/\Upsilon(\Omega)$ and $\kappa = \frac{\varepsilon}{\Psi} \frac{1 + \gamma\rho}{\rho/y}$. Substituting the target criterion into the Phillips curve, we get a second-order difference equation:

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

The solution to this system has the form:

$$x_t = \mathcal{A}_x x_{t-1} + \mathcal{A}_z \hat{z}_t + \mathcal{A}_\zeta \hat{\zeta}_t + \mathcal{A}_\varepsilon \hat{\varepsilon}_t + \mathcal{A}_\varsigma \hat{\varsigma}_t \quad (\text{F.3})$$

$$\pi_t = \mathcal{B}_x x_{t-1} + \mathcal{B}_z \hat{z}_t + \mathcal{B}_\zeta \hat{\zeta}_t + \mathcal{B}_\varepsilon \hat{\varepsilon}_t + \mathcal{B}_\varsigma \hat{\varsigma}_t \quad (\text{F.4})$$

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

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

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

$$\mathcal{A}_x = \frac{1}{2} \left(1 + \frac{\kappa\bar{\varepsilon} + 1}{\beta} - \sqrt{\left[1 + \frac{\kappa\bar{\varepsilon} + 1}{\beta} \right]^2 - \frac{4}{\beta}} \right) \in (0, 1) \quad (\text{F.6})$$

$$\mathcal{A}_z = \frac{\kappa\beta^{-1}\bar{\varepsilon}}{\kappa\beta^{-1}\bar{\varepsilon} + (1 - \mathcal{A}_x) + \left(\frac{1}{\beta} - \varrho_z \right)} [1 - \delta(\Omega)] \frac{1 + (\rho/y)}{1 + \gamma\rho} \quad (\text{F.7})$$

$$\mathcal{A}_\zeta = \frac{\kappa\beta^{-1}\bar{\varepsilon}}{\kappa\beta^{-1}\bar{\varepsilon} + (1 - \mathcal{A}_x) + \left(\frac{1}{\beta} - \varrho_\zeta\right)} \chi(\Omega) \quad (\text{F.8})$$

$$\mathcal{A}_\varepsilon = -\frac{\kappa\beta^{-1}\bar{\varepsilon}}{\kappa\beta^{-1}\bar{\varepsilon} + (1 - \mathcal{A}_x) + \left(\frac{1}{\beta} - \varrho_\varepsilon\right)} \frac{\rho/y}{1 + \gamma\rho} \quad (\text{F.9})$$

$$\mathcal{A}_\varsigma = -\frac{\kappa\beta^{-1}\bar{\varepsilon}}{\kappa\beta^{-1}\bar{\varepsilon} + (1 - \mathcal{A}_x) + \left(\frac{1}{\beta} - \varrho_\varsigma\right)} \Xi(\Omega) \quad (\text{F.10})$$

$$\mathcal{B}_x = \frac{1 - \mathcal{A}_x}{\bar{\varepsilon}} \quad (\text{F.11})$$

$$\mathcal{B}_i = -\frac{1}{\bar{\varepsilon}} \mathcal{A}_i \quad \text{for } i \in \{z, \zeta, \varepsilon, \varsigma\} \quad (\text{F.12})$$

Claim 9. *The following statements are true:*

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

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

Proof. To see that $\frac{\kappa\beta^{-1}\bar{\varepsilon}}{\kappa\beta^{-1}\bar{\varepsilon} + (1 - \mathcal{A}_x) + \left(\frac{1}{\beta} - \varrho_i\right)} \in (0, 1)$ for $i \in \{z, \zeta, \varepsilon, \varsigma\}$, 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 $i \in \{z, \zeta, \varepsilon, \varsigma\}$ 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 $\mathcal{A}_\varepsilon < 0$. □

F.1 Proof of Propositions 5, 6, 11 and 12

F.1.1 Impact effects following a productivity shock

Since $x_{-1} = 0$, it follows from equations (F.3) and (F.4) that the impact effect of a productivity shock is:

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

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

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

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

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.3) and (F.4) that:

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

Following a markup shock, the dynamics of π_t and x_t are described by the same equations (F.1) and (F.2) except that $\bar{\varepsilon}(\Omega) = \frac{\varepsilon}{\Upsilon(\Omega)}$ 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 $\frac{\partial \mathcal{A}_\varepsilon}{\partial \bar{\varepsilon}} < 0$ (output falls less on impact when $\bar{\varepsilon}$ is lower) and $\frac{\partial \mathcal{B}_\varepsilon}{\partial \bar{\varepsilon}} < 0$ (inflation increases more on impact when $\bar{\varepsilon}$ is lower). We have:

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

where we have plugged in the expression for \mathcal{A}_x from (F.6) in the second line. Since

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

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

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

Clearly, the denominator is increasing in ε so \mathcal{B}_ε is decreasing in $\bar{\varepsilon}$.

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:

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

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

F.2 Impact effects following a risk shock

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

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

for $\Omega \geq \Omega^c$. Similarly, the impact response of π_0 is given by $\mathcal{B}_\varsigma = -\frac{1}{\bar{\varepsilon}}\mathcal{A}_\varsigma > 0$.

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

The following Lemma characterizes the behavior of $\hat{y}_t - \hat{y}_t^n$ and π_t following a generic shock \mathcal{S}_0 where $\mathcal{S}_0 \in \{\hat{z}_0, \hat{\varepsilon}_0, \hat{\zeta}_0, \hat{\varsigma}_0\}$ for large t . In doing so, the Lemma provides a proof of the claims made in Propositions 5, 6, 11 and 12 about long-run behavior of $\hat{y}_t - \hat{y}_t^n$ and π_t .

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

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

Proof. We know that $\frac{\partial\pi_0}{\partial\mathcal{S}_0} = \mathcal{B}_\mathcal{S} = -\frac{1}{\bar{\varepsilon}}\mathcal{A}_\mathcal{S}$. Thus, we need to show that for large enough t , π_t and $\hat{y}_t - \hat{y}_t^n$ 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 $\varrho_\mathcal{S} \neq \mathcal{A}_x$. Using this in (30), the dynamics of inflation can then be written as:

$$\pi_t = -\frac{\mathcal{A}_\mathcal{S}}{\bar{\varepsilon}} \left(\frac{\varrho_\mathcal{S}^{t+1} - \mathcal{A}_x^{t+1}}{\varrho_\mathcal{S} - \mathcal{A}_x} - \frac{\varrho_\mathcal{S}^t - \mathcal{A}_x^t}{\varrho_\mathcal{S} - \mathcal{A}_x} \right) \mathcal{S}_0$$

where $\bar{\varepsilon} > 0$, $\mathcal{A}_\mathcal{S} > 0$ 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 $\max\{\mathcal{A}_x, \varrho_\mathcal{S}\}$. If $\varrho_\mathcal{S} < \mathcal{A}_\mathcal{S}$, dividing expression for π_t above by \mathcal{A}_x^t and taking the limit $t \rightarrow \infty$, we have:

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

which has the same sign as $\mathcal{A}_\mathcal{S} \times \mathcal{S}_0$. 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 $\mathcal{A}_\mathcal{S} \times \mathcal{S}_0$ for large t . Instead if $\varrho_\mathcal{S} > \mathcal{A}_x$, dividing by $\varrho_\mathcal{S}^t$ and

taking limits, we have

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

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:

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

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

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

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 – equations (F.3) and (F.4). Substituting (E.5) into the linearized IS equation (17):

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

We can use this equation along with equations (F.3) and (F.4) to express i_t in terms of x_{t-1} and the shocks:

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

where

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

Next, we show that (32) implements the optimal allocations uniquely. First, note that first-differencing the target criterion (30) and multiplying by $\phi_x \equiv \phi \frac{\Upsilon(\Omega)}{\varepsilon}$ yields:

$$\phi \pi_t + \phi_x \Delta x_t = \phi \pi_t + \phi_{\text{gap}} (\Delta \hat{y}_t - \Delta \hat{y}_t^e) + \phi_y \Delta \hat{y}_t = 0$$

where $x_t = \hat{y}_t - \delta(\Omega) \frac{1 + (\rho/y)}{1 + \gamma \rho} \hat{z}_t$, $\phi > 0$ is a constant, $\phi_{\text{gap}} = \phi \frac{\Upsilon(\Omega)}{\varepsilon} \delta(\Omega)$ is the weight on the change in output gap and $\phi_y = \phi \frac{\Upsilon(\Omega)}{\varepsilon} (1 - \delta(\Omega))$. Here, instead of writing the rule in terms of \hat{y}_t and the output gap $\hat{y}_t - \hat{y}_t^e$, it is more convenient to write it in terms of a single variable x_t ; the two formulations are equivalent. Since by definition, we have $i_t = i_t^*$ under optimal policy, it follows that the rule (32) is satisfied at the optimal allocation. To see that this rule implements optimal allocations uniquely, it suffices to look at the determinacy properties of the system comprised by the IS curve, the Phillips curve and the interest rate rule absent shocks. This system can be written as:

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

In matrix-form, this can be written as:

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

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

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

Notice that

$$\mathcal{P}(0) = \frac{\phi \Upsilon(\Omega) / \varepsilon - \Phi_x}{\beta \gamma y \left[1 + (1 - \tilde{\beta}) \Omega \right]} \quad \mathcal{P}(1) = \frac{\frac{\kappa(1 - \phi)}{\beta} - \left(\frac{1}{\beta} - 1 \right) \left[\Phi_x - \gamma y (1 - \tilde{\beta}) \Omega \right]}{\gamma y \left[1 + (1 - \tilde{\beta}) \Omega \right]}$$

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:

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

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 2 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 nonstockholder (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 nonstockholder $\mathbb{T}_t(i) = -T_t + \frac{\eta^d}{1 - \eta^d} J$. 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:

$$c_t^s(i; d) = C_t^d + \mu_t x_t^s(i; d)$$

and for nonstockholders:

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

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

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

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

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

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:

$$\eta^d \mathcal{C}_t^d + (1 - \eta^d) \mathcal{C}^{nd} = y_t$$

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):

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

Iterating forwards:

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

In other words, we have $\mathcal{C}_t^d = y_t + \frac{1 - \eta^d}{\eta^d} \mu_t \mathcal{V}_t$ as in the main text. Market clearing, then implies that $\mathcal{C}_t^{nd} = y_t - \mu_t \mathcal{V}_t$. 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 nonstockholders is the same $\mathcal{C}^d = \mathcal{C}^{nd}$. 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 aggregate dividends $D_t = y_t - (1 - \tau^*) \tilde{w}_t n_t$ can be written as:

$$D_t = y_t - \frac{(\varepsilon - 1) w_t}{\varepsilon (1 - \tau^w)} \frac{y_t}{z_t} \left[1 + \frac{\Psi}{2} (\Pi_t - 1)^2 \right] = \left(1 - \frac{\varepsilon - 1}{\varepsilon (1 - \tau^w)} \frac{w_t}{z_t} \right) y_t - \frac{\varepsilon - 1}{\varepsilon (1 - \tau^w)} \frac{w_t}{z_t} \frac{\Psi}{2} (\Pi_t - 1)^2 y_t,$$

we can write the level-deviation \widehat{D}_t as:

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

Using this it is straightforward to derive $\widetilde{V}_t = \mathcal{D}_y \widehat{y}_t + \mathcal{D}_z \widehat{z}_t + \beta \widetilde{V}_{t+1}$, where $\widetilde{V}_t = \widehat{V}_t / y$ and \widehat{V}_t denotes the level deviation of \mathcal{V}_t from its steady state value of 0.

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 nonstockholders, 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:

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

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

G.2 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.

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

where $1 - \tau^w$ is given by (23). 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 $\mathbb{B}'(0) = 0$, 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 nonstockholders is equal, we have the following.

FOC wrt w_t :

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

FOC wrt y_t :

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

FOC wrt Σ_t :

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

FOC wrt Π_t :

$$\hat{m}_{2,t} = \frac{(1 - \beta^{-1} \tilde{\beta})(1 - \Lambda)}{1 - \beta^{-1} \tilde{\beta}(1 - \Lambda)} (1 + \Omega) \frac{(\gamma y) \Psi}{1 + \gamma \rho} \hat{p}_t \quad (\text{G.12})$$

FOC wrt μ_t :

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

FOC wrt \mathcal{V}_t :

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

where $\hat{m}_{5,t} = \widehat{M}_{5,t}/\mathbb{U}$. Following the same steps as in Appendix E.4.1, we can arrive at the following expression which is the analog of equations (E.37)-(E.38) in that Appendix:

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

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 $t = 0$:

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

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

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

which is the same as in Proposition 7. Clearly, $\mathbb{K}(1) = 0$ and $\mathbb{K}'(\eta^d) = -\frac{\rho}{m_4} \frac{\gamma^2 \mu^2}{1-\beta} \left(\frac{1}{\eta^d} \right)^2 < 0$.

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

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

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 $\mathbb{K} \neq 0$.

G.3 LQ representation

Relative to the derivation of the LQ problem in our baseline model in Appendix E.2, the only difference is that unequally distributed profits introduce an additional term in the second-order Σ_t recursion, which can now be written as:

$$\begin{aligned} g_t^\Sigma &\approx \Lambda \hat{\mu}_t + \gamma y (1 - \Theta) \hat{y}_t + \beta^{-1} \tilde{\beta} \hat{\Sigma}_{t-1} - \hat{\Sigma}_t + \frac{1}{2} \hat{\Sigma}_t^2 - \frac{1}{2} (\beta^{-1} \tilde{\beta})^2 \hat{\Sigma}_{t-1}^2 + (\gamma y)^2 \frac{(1 - \Theta)^2}{\Lambda} \hat{y}_t^2 \\ &\quad + 2\gamma y (1 - \Theta) \hat{\mu}_t \hat{y}_t + \frac{\Lambda}{2} \hat{\mu}_t^2 + \frac{1}{2} (\gamma y)^2 \left(\frac{1 - \eta^d}{\eta^d} \right) \mu^2 \left[\mathbb{I}(t = 0) \tilde{\mathcal{V}}_0^2 + \mathbb{I}(t > 0) (1 - \beta^{-1} \tilde{\beta}) \tilde{\mathcal{V}}_t^2 \right] \end{aligned} \quad (\text{G.15})$$

The rest of the equations remain unchanged. Thus, the purely second-order approximation to the planner's objective is as described in (E.22) plus the additional terms involving $\tilde{\mathcal{V}}_t$ (multiplied by m_3). Thus, following the same steps above, we can arrive at the same expression as in Proposition 7:

$$\frac{1}{2} \sum_{t=0}^{\infty} \beta^t \left\{ \Upsilon(\Omega) \left(\hat{y}_t - \delta(\Omega) \hat{y}_t^e \right)^2 + \frac{\varepsilon}{\kappa} \pi_t^2 \right\} + \frac{\mathbb{K}(\eta^d)}{2} \left\{ \tilde{\mathcal{V}}_0^2 + \sum_{t=1}^{\infty} \beta^t (1 - \beta^{-1} \tilde{\beta}) \tilde{\mathcal{V}}_t^2 \right\} \quad (\text{G.16})$$

The optimal policy problem can now simply be specified as minimizing (G.16) subject to the linearized Phillips curve (24) and valuation equation (34). In Lagrangian form:

$$\begin{aligned} \mathcal{L} = & \frac{1}{2} \sum_{t=0}^{\infty} \beta^t \left\{ \Upsilon(\Omega) \left(\hat{y}_t - \delta(\Omega) \hat{y}_t^e \right)^2 + \frac{\varepsilon}{\kappa} \pi_t^2 \right\} + \frac{\mathbb{K}(\eta^d)}{2} \left\{ \tilde{\mathcal{V}}_0^2 + \sum_{t=1}^{\infty} \beta^t \left(1 - \beta^{-1} \tilde{\beta} \right) \tilde{\mathcal{V}}_t^2 \right\} \\ & + \sum_{t=0}^{\infty} \beta^t F_{1,t} \left\{ \beta \pi_{t+1} + \kappa \left[\hat{y}_t - \hat{y}_t^e + \frac{\rho/y}{1 + \gamma \rho} \hat{\varepsilon}_t \right] - \pi_t \right\} \\ & + \sum_{t=0}^{\infty} \beta^t F_{2,t} \left\{ \mathcal{D}_y \hat{y}_t + \frac{\varepsilon - 1}{\varepsilon} \frac{1 + \gamma \rho}{\rho/y} \hat{y}_t^e + \tilde{\beta} \tilde{\mathcal{V}}_{t+1} - \tilde{\mathcal{V}}_t \right\} \end{aligned}$$

The FOC w.r.t. \hat{y}_t can be written as:

$$\Upsilon(\Omega) \left(\hat{y}_t - \delta(\Omega) \hat{y}_t^e \right) + \kappa F_{1,t} + F_{2,t} \mathcal{D}_y = 0$$

The FOC w.r.t. π_t can be written as:

$$\frac{\varepsilon}{\kappa} \pi_t - F_{1,t} + F_{1,t-1} = 0 \quad \Leftrightarrow \quad F_{1,t} = \frac{\varepsilon}{\kappa} \hat{p}_t$$

where $\kappa = \frac{\varepsilon}{\Psi} \frac{1 + \gamma \rho}{\rho/y}$. Finally the FOC w.r.t. \mathcal{V}_t can be written as:

$$\begin{aligned} \mathbb{K}(\eta^d) \tilde{\mathcal{V}}_0 - F_{2,0} &= 0 & \text{for } t = 0 \\ \left(1 - \beta^{-1} \tilde{\beta} \right) \mathbb{K}(\eta^d) \tilde{\mathcal{V}}_0 - F_{2,t} + \beta^{-1} \tilde{\beta} F_{2,t-1} &= 0 & \text{for } t > 0 \end{aligned}$$

Combining these three FOCs, we can derive the target criterion in Proposition 7:

$$\Upsilon(\Omega) x_0 + \varepsilon \hat{p}_0 + \mathbb{K}(\eta^d) \mathcal{D}_y \hat{\mathcal{V}}_0 = 0 \quad (\text{G.17})$$

and for $t > 0$:

$$\Upsilon(\Omega) \left(x_t - \beta^{-1} \tilde{\beta} x_{t-1} \right) + \varepsilon \left(\hat{p}_t - \beta^{-1} \tilde{\beta} \hat{p}_{t-1} \right) + \mathbb{K}(\eta^d) \mathcal{D}_y \left(1 - \beta^{-1} \tilde{\beta} \right) \hat{\mathcal{V}}_t = 0 \quad (\text{G.18})$$

where $x_t = \hat{y}_t - \delta(\Omega) \hat{y}_t^e$.

H Hand to Mouth households

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 who cannot trade bonds and consume their after tax-income. These households are otherwise identical to the remaining $1 - \eta^h$ unconstrained households who 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 that of the unconstrained households, i.e. $\tilde{\mu}_t > \mu_t$.

Appendix H.1 and H.2 show that the presence of HtM households does not change the dynamics of aggregate variable, conditional on a given a path of interest rates. These dynamics are still given by (17)-(19) – 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.²⁷ 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 $U_t = u(c_t, n_t; \xi) \times \Sigma_t$, 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 tradeoffs 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 6. *The effect of a one-time increase in output engineered by monetary policy reduces inequality Σ_t by a larger amount, the larger the fraction of HtM households η^h :*

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

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 Shock Figure 8 shows the dynamics under optimal policy following a negative productivity shock in RANK (dashed red curves), HANK with no HtMs (solid blue curves) and HANK with 30% HtMs (dot-dashed magenta curves).²⁸ 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 who cannot self-insure using the bond market. This can be seen by comparing the dot-dashed magenta curves in panel c), which plots consumption inequality amongst 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,

²⁷This is for the same reasons as in Bilbiie (2008); Werning (2015); Acharya and Dogra (2020).

²⁸ $\eta^h = 0.3$ is in line with Kaplan et al. (2014) who find that approximately 30% of U.S. 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.

compare the dot-dashed magenta curve in panel e), which plots inequality under optimal policy with 30% HtMs, to the dotted-black 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 amongst 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.




Figure 8: **Optimal policy in response to productivity shocks** In panels a and b, solid blue 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 negative productivity 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 dotted-black 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 $\times 100$.

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 amongst HtMs, monetary policy stabilizes output more (dot-dashed magenta curve relative to solid blue curve in panel a), Figure 9) at the cost of higher inflation (dot-dashed magenta curve relative to solid blue 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 tradeoffs analyzed in our baseline. In fact, it accentuates the differences relative to RANK: with higher MPCs, i.e., higher passthrough from income to consumption risk, consumption inequality is even more sensitive to monetary policy. Consequently, policy deviates even further from RANK to stabilize inequality. This suggests that the tradeoffs we study analytically would be even more important in quantitative HANK economies with a substantial fraction of high MPC households.




Figure 9: **Optimal policy in response to markup shocks** In panels a and b, solid blue 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 dotted-black 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 $\times 100$.

H.1 Decision problem of HtM households

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

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

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:

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

Since $c_t^s(i; h) = y_t + \tilde{\mu}_t w_t (\xi_t^s(i) - \bar{\xi})$, 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:

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

where (25) describes the evolution of $\hat{\Sigma}_t$. 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:

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

where we have used the equilibrium relationship between wages and output (E.3) (we have also set all shocks to zero without loss of generality). Thus, we have:

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

We consider a one-time change in $\widehat{y}_t > 0$ engineered by monetary policy. Since equations (17)-(19) 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)-(E.3) to zero:

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

where the first equation is (17), the second is (18) and the last equation is (E.3). Combining the three equations and eliminating \widehat{w}_t yields

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

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

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

Taking the derivative w.r.t η^h , we get:

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

which is negative for countercyclical and acyclical risk ($\Theta \geq 1$) for β sufficiently close to 1.²⁹ Thus, a higher fraction of HTMs (η^h) implies that Σ_t falls more in response to the same increase in output.

²⁹To see this, note that $\Sigma^h \frac{1}{(1-\tilde{\beta})^2} - \Sigma^{nh} = \frac{e^{\frac{\Lambda}{2}(1-\beta\vartheta e^{\frac{\Lambda}{2}})^{-2}}}{(1-\beta\vartheta e^{\frac{\Lambda}{2}})^2} - \frac{1-\vartheta}{1-\vartheta e^{\frac{\Lambda}{2}}} e^{\frac{\Lambda}{2}}$ is increasing in β , negative at $\beta = 0$ and positive at $\beta = 1$ for any ϑ, Λ satisfying $\vartheta e^{\Lambda/2} < 1$.

H.4 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.

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

Fiscal policy sets τ^w such that the planner finds it optimal to implement $\Pi = 1$ in steady state, as in our baseline. To plot Figures 8 and 9, we first solve for τ^w numerically, then we linearize the first order conditions and compute the optimal dynamics to shocks numerically.

I Persistent income risk

Our baseline model described in the main paper featured i.i.d. idiosyncratic income risk, whereas empirical studies find that idiosyncratic income risk is highly persistent (Heathcote et al., 2010; Guvenen et al., 2021). We now relax this assumption by allowing for persistent idiosyncratic disutility shocks. Specifically, we assume that

$$\xi_t^s(i) - \bar{\xi} = \sigma_t e_t^s(i) \quad \text{where} \quad e_t^s(i) = \varrho_\xi e_{t-1}^s(i) + v_t^s(i), \quad v_t^s(i) \sim i.i.d. N(0, 1), \quad e_{s-1}^s(i) = 0 \quad (\text{I.1})$$

We allow for $0 \leq \varrho_\xi \leq 1$. Setting $\varrho_\xi = 0$ corresponds to the baseline model. As in the baseline model, we allow for a flexible specification for the cyclicity of income risk by assuming that $w_t \sigma_t = w \sigma e^{\varphi(y_t - y)}$. Appendix I.1 shows that the optimal consumption decision rule of a household is described by

$$c_t^s(i) = C_t + \mu_t (a_t^s(i) + h_t^s(i)) \quad (\text{I.2})$$

and the aggregate Euler equation is now given by

$$C_t = -\frac{1}{\gamma} \ln \beta R_t + C_{t+1} - \frac{\gamma \mu_{t+1}^2 \sigma_{h,t+1}^2}{2} \quad (\text{I.3})$$

where $a_t^s(i)$ is the household's financial wealth and $h_t^s(i) \equiv \sigma_{h,t} e_t^s(i)$ denotes the household's human wealth, defined as the expected present-discounted value of their labor endowment

$$h_t^s(i) = \mathbb{E}_t \sum_{\tau=0}^{\infty} Q_{t+\tau|t} w_{t+\tau} (\xi_{t+\tau}^s(i) - \bar{\xi}) = \underbrace{\left[\sum_{\tau=0}^{\infty} Q_{t+\tau|t} w_{t+\tau} \sigma_{t+\tau} \varrho_{\xi}^{\tau} \right]}_{\sigma_{h,t}} e_t^s(i) \quad (\text{I.4})$$

where $Q_{t+\tau|t} = \prod_{k=0}^{\tau-1} \frac{\vartheta}{R_{t+k}}$. As in the baseline model, the MPC out of household financial and human wealth, μ_t is still given by (12). The consumption risk faced by households, the last term in (I.3) depends on the passthrough from human wealth to consumption (measured by μ_{t+1}^2) and the variance of shocks to human wealth $\sigma_{h,t+1}^2$. In our baseline model ($\varrho_{\xi} = 0$), human wealth $h_t^s(i)$ is simply $w_t(\xi_t^s - \bar{\xi})$, making (I.2) identical to (9), and the variance of shocks to human wealth is simply $\sigma_{h,t+1}^2 = w_{t+1}^2 \sigma_{t+1}^2$. However, with persistent idiosyncratic income, a positive shock to the household's current labor endowment also increases the expected value of their endowment in the future. This is reflected in the fact that $\sigma_{h,t}$ depends on not just $w_t \sigma_t$, but the whole future path $\{w_{t+k} \sigma_{t+k}\}_{k=0}^{\infty}$.

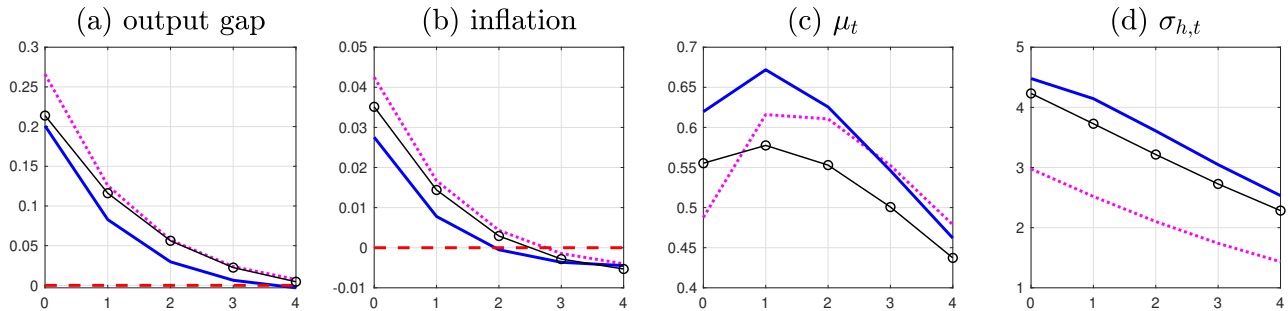


Figure 10: **Optimal policy in response to productivity shocks** In all panels, red-dashed curves depict dynamics in RANK; solid blue curves depicts dynamics in HANK with $\varrho_{\xi} = 0$; black lines with circle markers depicts dynamics in HANK with $\varrho_{\xi} = 0.5$ and the magenta dotted line depicts dynamics in HANK with $\varrho_{\xi} = 1$. All panels plot log-deviations from steady state $\times 100$.

Appendix I.2 shows that, as in our baseline, a utilitarian planner's felicity function can be decomposed into the flow utility of a notional representative agent and a welfare-relevant measure of consumption inequality Σ_t , which now evolves according to

$$\ln \Sigma_t = \frac{\gamma^2 \mu_t^2 \sigma_{h,t}^2}{2} + \ln [1 - \vartheta + \vartheta \Sigma_{t-1}] \quad (\text{I.5})$$

It is worth noting that with $\varrho_{\xi} > 0$, our economy features not one, but two dimensions of persistent wealth inequality: financial and human wealth inequality. In principle, this means that the planner must forecast the evolution of the joint distribution of financial and human wealth, not just the distribution of financial wealth as in the baseline. However, as (I.5) indicates, the evolution of this joint distribution can still be summarized by a single scalar Σ_t which depends on its own lagged value. This highlights the analytical tractability of our framework.

Equation (I.5) along with the definition of σ_h in (I.4) reveals that persistence ($\varrho_{\xi} > 0$) modifies the effect of monetary policy on consumption inequality in two ways. First, lower real interest rates, holding the path of aggregate output and wages fixed, now tend to increase the variance of human wealth σ_h^2 ,

putting more weight on the value of future labor endowments. Thus, while the effect of interest rates on passthrough μ_t remain unchanged (relative to the baseline i.i.d. case), $\varrho_\xi > 0$ tends to weaken the overall effect of interest rates on consumption risk, given the level of output. But this is not the only effect of higher persistence. Lower real interest rates also increase output, which reduces human capital risk (in the countercyclical income risk case) as in our baseline model. This effect becomes more pronounced, the higher the level of human capital risk σ_h . Higher ϱ_ξ tends to increase the level of human capital risk (for the same sequence of $\{w_{t+k}, \sigma_{t+k}\}_{k=0}^\infty$), since the same shock to current income has a larger effect on lifetime income: $\sigma_{h,t}(\varrho_\xi > 0) > \sigma_{h,t}(\varrho_\xi = 0)$. Thus, higher ϱ_ξ amplifies the effect of monetary policy on Σ_t via the level of output. Overall, this second effect dominates and higher persistence increases the sensitivity of Σ_t to changes in output induced by monetary policy. This effect is itself long-lived— $\frac{\partial \widehat{\Sigma}_{t+k}}{\partial \widehat{y}_t}$ is larger in absolute value at all horizons $k > 0$ when ϱ_ξ is higher— because consumption inequality is only slow to revert to its mean value following an increase in consumption risk (cf. equation (I.5)).

Lemma 7. *The effect of a one-time increase in output engineered by monetary policy reduces inequality Σ_{t+k} at all horizons $k \geq 0$ by a larger amount, the larger the persistence of idiosyncratic income ϱ_ξ :*

$$\frac{\partial}{\partial \varrho_\xi} \left(\frac{\partial \widehat{\Sigma}_t}{\partial \widehat{y}_t} \right) < 0 \quad \text{with acyclical/countercyclical income risk, } \Theta \geq 1$$

and

$$\frac{\partial}{\partial \varrho_\xi} \left(\frac{\partial \widehat{\Sigma}_{t+k}}{\partial \widehat{y}_t} \right) = \left(\beta^{-1} \widetilde{\beta} \right)^k \frac{\partial}{\partial \varrho_\xi} \left(\frac{\partial \widehat{\Sigma}_t}{\partial \widehat{y}_t} \right) \quad \forall k > 0$$

Proof. See Appendix I.4. □

Consequently, since the sensitivity of consumption risk to monetary policy is the main force leading optimal monetary policy to differ in HANK and RANK, introducing persistent idiosyncratic income risk magnifies these differences. Figure 10 shows the dynamics under optimal policy following a negative productivity shock in RANK (dashed red curves), and HANK with $\varrho_\xi = 0$ (blue line), $\varrho_\xi = 0.5$ (black line with circle markers) and $\varrho_\xi = 1$ (magenta dotted line). Recall that in our baseline with $\varrho_\xi = 0$, the HANK planner already cushions the fall in output relative to the RANK planner, resulting in higher inflation on impact. The black line with circle markers and magenta dotted line indicate that higher ϱ_ξ leads the HANK planner to cushion the fall in output even more, leading to higher inflation on impact. To understand why, note that the steady state *level* of human capital risk σ_h is the highest for the economy with $\varrho_\xi = 1$ and the lowest when $\varrho_\xi = 0$. Thus, panel (d) shows that by curtailing the fall in output, the HANK planner permits a smaller *proportional* increase in $\sigma_{h,t}$ when the level of σ_h is already high, i.e., in the economy with $\varrho_\xi = 1$ (compare the magenta dotted and blue lines). The planner does not allow a large increase in the level of $\sigma_{h,t}$, even temporarily, since doing so would persistently increase consumption inequality Σ_t (cf. Lemma 7). This more moderate decline in output (and smaller proportional increase in $\sigma_{h,t}$) also results in a smaller increase in passthrough μ_t (panel (c)). The case with $\varrho_\xi = 0.5$ lies between the i.i.d and random walk extremes. A higher ϱ_ξ modifies the optimal response to a markup shock in a similar fashion; we omit the results for the sake of brevity.

Overall, persistent income risk, like MPC heterogeneity, does not change the tradeoff facing the planner qualitatively. In fact, it also accentuates the difference relative to RANK, compared to the case with i.i.d.

income risk. Again this suggests that the tradeoffs we study analytically would be even more important in quantitative HANK models with realistic income processes.

I.1 Derivation of household decision rules

The date s problem of an individual i born at date s is now

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

subject to

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

where $\xi_t^s(i) - \bar{\xi} = \sigma_t e_t^s(i)$ and

$$e_t^s(i) = \varrho_\xi e_{t-1}^s(i) + v_t^s(i) \quad v_{i,t} \sim N(0, 1)$$

The derivation of the consumption function follows that in Appendix A. Guess that the consumption function takes the form:

$$c_t^s(i) = \mathcal{C}_t + \mu_t (a_t^s(i) + h_t^s(i))$$

where $h_t^s(i)$ denotes the expected present-discounted value of the household's labor endowment:

$$h_t^s(i) = \mathbb{E}_t \sum_{k=0}^{\infty} Q_{t+k|t} w_{t+k} (\xi_{t+k}^s(i) - \bar{\xi}) \equiv \sigma_{h,t} e_t^s(i) \quad Q_{t+k|t} = \prod_{k=0}^{\tau} \frac{\vartheta}{R_{t+k}}$$

Using the budget constraint, labor supply and the household's Euler equation, we have:

$$\begin{aligned} y_t + \left\{ \mu_t - \mu_{t+1} \frac{R_t}{\vartheta} [1 - (1 + \rho\gamma w_t) \mu_t] \right\} a_t^s(i) &= -\frac{1}{\gamma} \ln \beta R_t + y_{t+1} \\ &+ \left\{ \mu_{t+1} \frac{R_t}{\vartheta} [\sigma_t w_t - (1 + \rho\gamma w_t) \mu_t \sigma_{h,t}] + \varrho_\xi \mu_{t+1} \sigma_{h,t+1} - \mu_t \sigma_{h,t} \right\} e_t^s(i) \\ &- \frac{\gamma}{2} \mu_{t+1}^2 \sigma_{h,t+1}^2 \end{aligned}$$

Matching coefficients yields the standard μ_t recursion:

$$\mu_t^{-1} = 1 + \gamma \rho w_t + \frac{\vartheta}{R_t} \mu_{t+1}^{-1} \tag{I.6}$$

In addition, we have the following equation describing $\sigma_{h,t}$

$$\sigma_{h,t} = \sigma_t w_t + \frac{\vartheta}{R_t} \varrho_\xi \sigma_{h,t+1} \tag{I.7}$$

and the aggregate Euler equation is now given by

$$y_t = y_{t+1} - \frac{1}{\gamma} \ln \beta R_t - \frac{\gamma \mu_{t+1}^2 \sigma_{h,t+1}^2}{2} \quad (\text{I.8})$$

where we have used $C_t = y_t$ from market clearing.

I.2 Deriving the Σ recursion

As in our baseline, we assume that the planner is utilitarian and puts identical weight (equal to 1) on the welfare of all individuals on individual i both at date $s \leq 0$ and β^s on the welfare of individuals who will be born at date $s > 0$. Recall that in our baseline we allow the planner to set a date 0 tax on financial wealth to focus on the role of monetary policy in providing insurance, rather than redistribution between borrowers and lenders. But when $\varrho_\xi > 0$, households alive at the beginning of date 0 differ not only in financial wealth but also in terms of human wealth. To remove the planner's incentive to use monetary policy to redistribute between individuals with high and low human wealth, we allow the planner to tax individuals on their total wealth at the beginning of date 0. At the beginning of date 0, when the date 0 idiosyncratic shock $v_0^s(i)$ to household i 's time endowment has not yet been realized, the household's total wealth is given by $a_0^s(i) + \sigma_{h,0} \varrho_\xi e_{-1}^s(i)$ where we have used the fact that

$$h_0^s(i) = \sigma_{h,0} e_0^s(i) = \sigma_{h,0} \varrho_\xi e_{-1}^s(i) + \sigma_{h,0} v_0^s(i)$$

The planner levies a tax τ_0^a on this total amount implying that the household's post-tax human wealth after the realization of their date 0 idiosyncratic shock $v_0^s(i)$ is given by $(1 - \tau_0^a) [a_0^s(i) + \sigma_{h,0} \varrho_\xi e_{-1}^s(i)] + \sigma_{h,0} v_0^s(i)$. This also implies that the date 0 tax on financial wealth $\tau_0^a = 1$. However, δ now measures the extent to which the planner is willing to tolerate pre-existing human wealth inequality. $\delta = 0$ implies that the planner is also utilitarian towards human wealth inequality at date 0 while a higher δ implies that the planner assigns higher weights to the welfare of those with higher human wealth as of date -1. As in the baseline, the planner's objective function can be written as:

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

where Σ_t is now defined as

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

Next, subtracting the aggregate Euler equation from a household's Euler equation for all dates $t \geq 0$, we get

$$c_{t+1}^s(i) - c_{t+1} = c_t^s(i) - c_t + \mu_{t+1} \sigma_{h,t+1} v_{t+1}^s(i)$$

Using this in the definition of Σ_t for $t \geq 1$

$$\begin{aligned}
\Sigma_t &= (1 - \vartheta) \sum_{s=-\infty}^{t-1} \vartheta^{t-s} \int e^{-\gamma(c_{t-1}^s(i) - c_{t-1} + \mu_t \sigma_{h,t} v_t^s(i))} di + (1 - \vartheta) \int e^{-\gamma(\mu_t \sigma_{h,t} v_t^t(i))} di \\
&= \vartheta e^{\frac{1}{2} \gamma^2 \mu_t^2 \sigma_{h,t}^2} \left\{ (1 - \vartheta) \sum_{s=-\infty}^{t-1} \vartheta^{t-1-s} \int e^{-\gamma(c_{t-1}^s(i) - c_{t-1})} di \right\} + (1 - \vartheta) e^{\frac{1}{2} \gamma^2 \mu_t^2 \sigma_{h,t}^2} \\
&= e^{\frac{1}{2} \gamma^2 \mu_t^2 \sigma_{h,t}^2} [1 - \vartheta + \vartheta \Sigma_{t-1}]
\end{aligned}$$

Taking logs, we get

$$\ln \Sigma_t = \frac{\gamma^2 \mu_t^2 \sigma_{h,t}^2}{2} + \ln [1 - \vartheta + \vartheta \Sigma_{t-1}]$$

Next, for $t = 0$, we have

$$\begin{aligned}
\Sigma_0 &= (1 - \vartheta) \sum_{s=-\infty}^0 \vartheta^{-s} \int e^{-\gamma(c_0^s(i) - c_0)} di \\
&= (1 - \vartheta) \sum_{s=-\infty}^0 \vartheta^{-s} \int e^{-\gamma \mu_0 (1 - \tau_0^a) (a_0^s(i) + \sigma_{h,0} \varrho_\xi e_{-1}^s(i)) - \gamma \mu_0 \sigma_{h,0} v_0^s(i)} di \\
&= e^{\frac{1}{2} \gamma^2 \mu_0^2 \sigma_{h,0}^2} \left\{ (1 - \vartheta) \sum_{s=-\infty}^0 \vartheta^{-s} \int e^{-\gamma \mu_0 (1 - \tau_0^a) (a_0^s(i) + \sigma_{h,0} \varrho_\xi e_{-1}^s(i))} di \right\}
\end{aligned}$$

Clearly, since $a_0^s(i) + \sigma_{h,0} \varrho_\xi e_{-1}^s(i)$ has zero mean, the planner chooses $\tau_0^a = 1$ to minimize this expression, implying that the date 0 Σ recursion is the same as at all future dates:

$$\ln \Sigma_0 = \frac{1}{2} \gamma^2 \mu_0^2 \sigma_{h,0}^2 + \ln [1 - \vartheta + \vartheta \Sigma_{-1}] \quad \text{where} \quad \Sigma_{-1} = 1$$

I.3 Planning Problem

The planning problem can be written as:

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

s.t.

$$\gamma y_t = \gamma y_{t+1} - \ln \beta \vartheta + \ln \mu_{t+1} + \ln [\mu_t^{-1} - (1 + \gamma \rho w_t)] - \frac{\gamma^2 \mu_{t+1}^2 \sigma_{h,t+1}^2}{2} \quad (\text{I.9})$$

$$(\Pi_t - 1) \Pi_t = \frac{\varepsilon_t}{\Psi} \left[1 - \frac{\varepsilon_t - 1}{\varepsilon_t} \frac{(1 - \tau^w) z_t}{(1 - \tau^*) w_t} \right] + \beta \left(\frac{z_t y_{t+1} w_{t+1}}{z_{t+1} y_t w_t} \right) (\Pi_{t+1} - 1) \Pi_{t+1} \quad (\text{I.10})$$

$$\ln \Sigma_t = \frac{\gamma^2 \mu_t^2 \sigma_{h,t}^2}{2} + \ln [1 - \vartheta + \vartheta \Sigma_{t-1}] \quad (\text{I.11})$$

$$y_t = z_t \frac{\rho \ln w_t + \bar{\xi}}{1 + \rho\gamma z_t + \frac{\Psi}{2} (\Pi_t - 1)^2} \quad (\text{I.12})$$

$$\sigma_{h,t} = \sigma_t w_t + \varrho_\xi \mu_{t+1} [\mu_t^{-1} - 1 - \gamma\rho w_t] \sigma_{h,t+1} \quad (\text{I.13})$$

$$\Sigma_{-1} = 1 \quad (\text{I.14})$$

This can be expressed as a Lagrangian:

$$\begin{aligned} \mathcal{L} = & \sum_{t=0}^{\infty} \beta^t \left\{ -\frac{1}{\gamma} (1 + \gamma\rho w_t) e^{-\gamma y_t} \Sigma_t \right\} \\ & + \sum_{t=0}^{\infty} \beta^t M_{1,t} \left\{ \gamma y_{t+1} - \ln \beta \vartheta + \ln \mu_{t+1} + \ln [\mu_t^{-1} - (1 + \gamma\rho w_t)] - \frac{\gamma^2 \mu_{t+1}^2 \sigma_{h,t+1}^2}{2} - \gamma y_t \right\} \\ & + \sum_{t=0}^{\infty} \beta^t M_{2,t} \left\{ \frac{\varepsilon_t}{\Psi} \left[1 - \frac{\varepsilon_t - 1}{\varepsilon_t} \frac{(1 - \tau^w) z_t}{(1 - \tau^*) w_t} \right] + \beta \left(\frac{z_t w_{t+1} y_{t+1}}{z_{t+1} w_t y_t} \right) (\Pi_{t+1} - 1) \Pi_{t+1} - (\Pi_t - 1) \Pi_t \right\} \\ & + \sum_{t=0}^{\infty} \beta^t M_{3,t} \left\{ \frac{\gamma^2 \mu_t^2 \sigma_{h,t}^2}{2} + \ln [1 - \vartheta + \vartheta \Sigma_{t-1}] - \ln \Sigma_t \right\} \\ & + \sum_{t=0}^{\infty} \beta^t M_{4,t} \left\{ y_t - z_t \frac{\rho \ln w_t + \bar{\xi}}{1 + \gamma\rho z_t + \frac{\Psi}{2} (\Pi_t - 1)^2} \right\} \\ & + \sum_{t=0}^{\infty} \beta^t M_{5,t} \left\{ \sigma w e^{\varphi(y_t - y)} + \varrho_\xi \mu_{t+1} [\mu_t^{-1} - 1 - \gamma\rho w_t] \sigma_{h,t+1} - \sigma_{h,t} \right\} \end{aligned}$$

FOC wrt y_t (equation is divided by $-\gamma$):

$$\begin{aligned} \mathbb{U}_t + M_{1,t} - \beta^{-1} M_{1,t-1} + \beta \gamma M_{2,t} \left(\frac{z_t w_{t+1} y_{t+1}}{z_{t+1} w_t y_t^2} \right) (\Pi_{t+1} - 1) \Pi_{t+1} - \gamma M_{2,t-1} \left(\frac{z_{t-1} w_t}{z_t w_{t-1} y_{t-1}} \right) (\Pi_t - 1) \Pi_t \\ - \frac{M_{4,t}}{\gamma} - \frac{\varphi}{\gamma} M_{5,t} \sigma e^{\varphi(y_t - y)} = 0 \end{aligned}$$

FOC wrt w_t (equation is multiplied by w_t):

$$\begin{aligned} 0 = & \frac{\gamma\rho w_t}{1 + \gamma\rho w_t} \mathbb{U}_t - M_{1,t} \frac{\gamma\rho w_t}{\mu_t^{-1} - (1 + \gamma\rho w_t)} + M_{2,t} \frac{\varepsilon_t - 1}{\Psi} \frac{(1 - \tau^w) z_t}{(1 - \tau^*) w_t} \\ & - \beta M_{2,t} \left(\frac{z_t w_{t+1} y_{t+1}}{z_{t+1} w_t y_t} \right) (\Pi_{t+1} - 1) \Pi_{t+1} + M_{2,t-1} \left(\frac{z_{t-1} w_t y_t}{z_t w_{t-1} y_{t-1}} \right) (\Pi_t - 1) \Pi_t \\ & - \frac{M_{4,t}}{\gamma} \frac{\gamma\rho z_t}{1 + \rho\gamma z_t + \frac{\Psi}{2} (\Pi_t - 1)^2} - M_{5,t} \varrho_\xi \mu_{t+1} \gamma\rho w_t \sigma_{h,t+1} \end{aligned}$$

FOC wrt Σ_t (equation is multiplied by Σ_t):

$$M_{3,t} = -\frac{1}{\gamma} (1 + \gamma\rho w_t) e^{-\gamma y_t \Sigma_t} + \beta \frac{\vartheta \Sigma_t}{1 - \vartheta + \vartheta \Sigma_t} M_{3,t+1}$$

FOC wrt μ_t (equation is multiplied by μ_t):

$$\begin{aligned} & -M_{1,t} \frac{\mu_t^{-1}}{\mu_t^{-1} - (1 + \gamma\rho w_t)} + \beta^{-1} M_{1,t-1} - (\beta^{-1} M_{1,t-1} - M_{3,t}) \gamma^2 \mu_t^2 \sigma_{h,t}^2 - \varrho_\xi \frac{\mu_{t+1}}{\mu_t} M_{5,t} \sigma_{h,t+1} \\ & + \beta^{-1} \varrho_\xi \mu_t [\mu_{t-1}^{-1} - 1 - \gamma\rho w_{t-1}] M_{5,t-1} \sigma_{h,t} = 0 \end{aligned}$$

FOC wrt $\sigma_{h,t}$ (equation is multiplied by $\sigma_{h,t}$):

$$0 = -(\beta^{-1} M_{1,t-1} - M_{3,t}) \gamma^2 \mu_t^2 \sigma_{h,t}^2 - M_{5,t} \sigma_{h,t} + \beta^{-1} \varrho_\xi \mu_t [\mu_{t-1}^{-1} - 1 - \gamma\rho w_{t-1}] M_{5,t-1} \sigma_{h,t}$$

FOC wrt Π_t :

$$\left[M_{2,t} - \left(\frac{z_{t-1} w_t y_t}{z_t w_{t-1} y_{t-1}} \right) M_{2,t-1} \right] (2\Pi_t - 1) = M_{4,t} z_t \frac{y_t}{1 + \gamma\rho z_t + \frac{\Psi}{2} (\Pi_t - 1)^2} \Psi (\Pi_t - 1)$$

Fiscal policy sets τ^w such that the planner finds it optimal to implement $\Pi = 1$ in steady state, as in our baseline. We solve this system numerically, taking a first-order approximation of the first-order conditions and the constraints (linearizing the multipliers and log-linearizing all other variables).

I.4 Proof of Lemma 7

We consider a one time change in output $\hat{y}_t > 0$ engineered by monetary policy. Since the equations (I.9), (I.12) and (I.13) are forward looking, monetary policy can implement this with a change in nominal interest rates only at date t without affecting macroeconomic aggregates in the future. Thus, the response of the other variables to a one time change in \hat{y}_t are given by the solution to the following linearized equations, where we have imposed that all variables return to their steady state values at date $t+1$ (except for $\hat{\Sigma}_{t+1}$):

$$\gamma y \hat{y}_t = -\frac{\mu^{-1}}{\mu^{-1} - (1 + \gamma\rho w)} \hat{\mu}_t - \frac{\gamma\rho w}{\mu^{-1} - (1 + \gamma\rho w)} \hat{w}_t$$

$$\hat{w}_t = \frac{1 + \gamma\rho}{\rho/y} \hat{y}_t$$

$$\hat{\sigma}_{h,t} = \frac{\sigma w}{\sigma_h} \varphi y \hat{y}_t - \varrho_\xi \hat{\mu}_t - \varrho_\xi \mu \gamma \rho w \hat{w}_t$$

Using the steady state relationships between these variables, we have:

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

$$\hat{\sigma}_{h,t} = \gamma y \left[(1 - \tilde{\beta} \varrho_\xi) \frac{\varphi}{\gamma} + \varrho_\xi \left[\tilde{\beta} + (1 - \tilde{\beta}) \frac{(1 + \gamma\rho) w}{1 + \gamma\rho w} \right] - (1 - \tilde{\beta}) \varrho_\xi \frac{(1 + \gamma\rho) w}{1 + \gamma\rho w} \right] \hat{y}_t$$

Finally, log-linearizing (I.11)

$$\begin{aligned}\widehat{\Sigma}_t &= \gamma^2 \mu^2 \sigma_h^2 (\widehat{\mu}_t + \widehat{\sigma}_{h,t}) + \beta^{-1} \widetilde{\beta} \widehat{\Sigma}_{t-1} \\ &= -(\gamma y) \gamma^2 \mu^2 w^2 \sigma^2 \left(\frac{1}{1 - \widetilde{\beta} \varrho_\xi} \left(1 - \frac{\varphi}{\gamma} \right) + \left(\frac{1 - \widetilde{\beta}}{1 - \widetilde{\beta} \varrho_\xi^2} \right) \frac{w - 1}{1 + \gamma \rho w} \right) \widehat{y}_t + \beta^{-1} \widetilde{\beta} \widehat{\Sigma}_{t-1}\end{aligned}$$

Thus, we have:

$$\frac{\partial}{\partial \varrho_\xi} \left(\frac{\partial \widehat{\Sigma}_t}{\partial \widehat{y}_t} \right) = -\widetilde{\beta} \gamma y \left[(\Theta - 1 + \Lambda) + 2 \left(\frac{1 - \widetilde{\beta}}{1 - \widetilde{\beta} \varrho_\xi} \right) \Lambda \Omega \right]$$

where $\Lambda = \frac{\gamma^2 \mu^2 w^2 \sigma^2}{1 - \widetilde{\beta} \varrho_\xi^2}$ and $\Theta = 1 - \frac{\varphi \Lambda}{\gamma}$ and $\Omega = \frac{w-1}{1+\gamma \rho w}$. With countercyclical risk and $w > 1$, this derivative is negative, implying that higher ϱ_ξ increases the sensitivity of $\widehat{\Sigma}_t$ to \widehat{y}_t (in absolute value). Given that $\widehat{y}_{t+k} = 0$ for $k > 0$ in the experiment considered, we have

$$\begin{aligned}\frac{\partial}{\partial \varrho_\xi} \left(\frac{\partial \widehat{\Sigma}_{t+k}}{\partial \widehat{y}_t} \right) &= \frac{\partial}{\partial \varrho_\xi} \left(\frac{\partial \widehat{\Sigma}_{t+k}}{\partial \widehat{\Sigma}_t} \frac{\partial \widehat{\Sigma}_t}{\partial \widehat{y}_t} \right) \\ &= \left(\beta^{-1} \widetilde{\beta} \right)^k \left\{ -\widetilde{\beta} \gamma y \left[(\Theta - 1 + \Lambda) + 2 \left(\frac{1 - \widetilde{\beta}}{1 - \widetilde{\beta} \varrho_\xi} \right) \Lambda \Omega \right] \right\}\end{aligned}$$

J Optimal 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 which 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 tradeoff 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 tradeoff 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 internalised 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 $\sigma_t^2 w_t^2 = \sigma^2 w^2 \exp \{2 [\varphi(y_t - y) + \varsigma_t]\}$. 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 Σ_t recursion (26) which now becomes:³⁰

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

An increase in $\widehat{\zeta}_t$ directly affects income risk and thus persistently affects consumption inequality. More subtly, a fall in households discount factor $\widehat{\zeta}_t < 0$ increase the natural rate of interest, which in our economy is given by $r_t^* = -\frac{1 - \widetilde{\beta} \varrho_\zeta}{1 - \widetilde{\beta} (1 - \Lambda) \varrho_\zeta} \widehat{\zeta}_t - \frac{(1 - \widetilde{\beta} \varrho_\zeta) \Lambda}{1 - \widetilde{\beta} (1 - \Lambda) \varrho_\zeta} \widehat{\zeta}_{t+1}$. Thus, if monetary policy keeps output unchanged in response to a fall in $\widehat{\zeta}_t$, this entails a rise in interest rates which increases the passthrough μ_t . For a given level of income risk, higher passthrough 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 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.³¹

Proposition 10. *In response to demand shocks, the planner sets nominal interest rates so that the following target criterion holds at all dates $t \geq 0$:*

$$(\widehat{y}_t - y_t^*) + \frac{\varepsilon}{\Upsilon(\Omega)} \widehat{p}_t = 0 \quad (\text{J.2})$$

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.4.1 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 (J.2) indicates that the planner seeks to minimize fluctuations of the price level while also keeping output close to its desired level y_t^* . When risk is acyclical or countercyclical, demand shocks which tend to increase consumption inequality – higher ς_t or lower ζ_t – increase y_t^* .

³⁰See Appendix E.1 for a derivation. We have implicitly set $\widehat{z}_t = 0$ throughout this section.

³¹For this section, we do not derive a quadratic loss function but derive the target criterion by linearizing the non-linear first order conditions of the planner's problem. The target criterion in Proposition 10 is a generalization of the target criterion in (30) to include demand shocks but abstracting from productivity shocks ($\widehat{z}_t = 0$). Appendix E.4.1 derives a general target criterion which is valid in the presence of all four shocks that we study.

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 11. *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 11 plots the optimal response to a 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 $\pi_t = 0$ (dashed red 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 (26)). Lower interest rates and higher output (which implies higher wages) also makes it easier for households to self insure, lowering the passthrough 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, monetary policy implements $\widehat{y}_t = \widehat{y}_t^n = 0$ and $\pi_t = 0$ (which was optimal under RANK), this would result in higher inequality (dotted black curve in panel c).

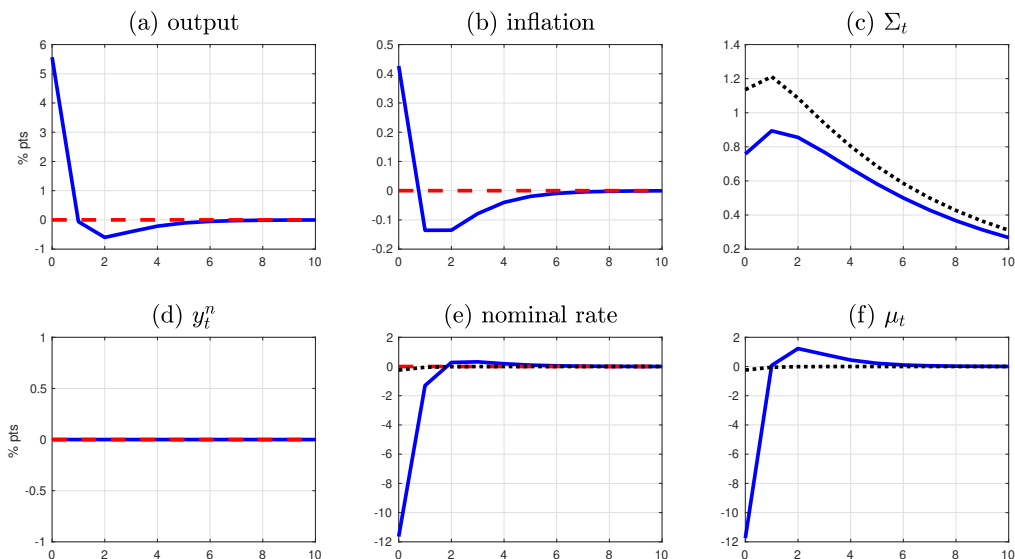


Figure 11: **Optimal policy in response to risk shocks** in HANK with $\Omega > 0$ (solid blue curves) and RANK (dashed red curves). Black-dotted lines denote outcomes in HANK under non-optimal policy which sets $\widehat{y}_t = \widehat{y}_t^n = 0$, $\pi_t = 0 \forall t \geq 0$. All panels plot log-deviations from steady state $\times 100$.

Discount factor shock A decrease in households' discount factor ($\widehat{\zeta}_t < 0$) increases r_t^* , the interest rate consistent with $\widehat{y}_t = \widehat{y}_t^n = 0$ 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 passthrough μ_t and hence consumption inequality. Thus, as with a positive risk shock, monetary policy deviates from the flexible-price allocation ($\widehat{y}_t = \widehat{y}_t^n = \pi_t = 0$) to mitigate this rise in inequality.

Proposition 12. *Under optimal policy with acyclical or countercyclical income risk, following an decrease in households' discount factor ($\widehat{\zeta}_0 < 0$), \widehat{y}_0 and π_0 both increase. In addition, $\exists T > 0$ such that for all $t \in (T, \infty)$, $\pi_t < 0$ and $\widehat{y}_t < 0$. Following a rise in households' discount factor, all these signs are reversed.*

Figure 12 plots the optimal dynamics following a negative discount factor shock. As in RANK, the HANK planner raises rates (panel e), increasing passthrough 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 12: **Optimal policy in response to discount factor shock** in HANK with $\Omega > 0$ (solid blue curves) and RANK (dashed red 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 $\times 100$.

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 passthrough 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 tradeoff 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 as well. 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.

K Irrelevance of government debt in the baseline model

Since we assume that new-born households receive a transfer from the government equal to average wealth B_t/ϑ , our baseline economy features a form of Ricardian 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:

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

which is the same as equation (A.6) in Appendix A. We also know that since $c_t^s(i) = \mathcal{C}_t + \mu_t a_t^s(i)$, aggregate consumption is given by:

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

Next, aggregating households' labor supply (10):

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

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 (K.1) along with $c_t = C_t + \mu_t \frac{B_t}{\vartheta}$:

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

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 new-born 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.