Optimal Monetary Policy According to HANK

Sushant Acharya, Edouard Challe, Keshav Dogra

Optimal Monetary Policy According to HANK*

Sushant Acharya† Edouard Challe‡ Keshav Dogra§

January 25, 2023

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 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, Moll and Violante 2018; Auclert, Rognlie and Straub 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, Martin-Baillon and Ragot 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

*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, PSE, Université Paris-Dauphine, Université Paris 8, University of St. Gallen, Banque de France, Drexel University, University of Essex, University of Melbourne and CREST, as well as from conference participants at the Barcelona Summer Forum, the NBER SI (Micro Data 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

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 cyclicality 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 *cyclicality 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 may not need to 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, Martin-Baillon and Ragot (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.

Like us, McKay and Wolf (2022) use an LQ approach to characterize optimal policy in HANK. As is well known, maximizing a quadratic approximation to welfare subject to a linear approximation of constraints does not deliver first-order accurate approximations to optimal policy when the steady state is *inefficient* – which is naturally the case in HANK models with incomplete markets and steady state consumption inequality. Following Benigno and Woodford (2005), we address this by using a second-order approximation to the constraints to eliminate first-order terms in the second-order approximation to the welfare-based loss function, allowing us to characterize *Pareto-optimal* allocations. McKay and Wolf (2022) instead eliminate linear terms by considering the problem of a planner who attaches quasi-Pareto weights to households' flow utilities based on their individual histories of idiosyncratic shocks. As we discuss in Section 4, while their approach yields a first-order accurate approximation to the solution of their planning problem, the solution to this planning problem is not generally Pareto-optimal, since their planner's objective function does not respect individual preferences.

Nuño and Thomas (2022) study how URE and unexpected inflation 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

[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.

households trade real (i.e., inflation-indexed) bonds, an earlier version of this paper did study this channel; its effect on optimal policy is similar to the URE channel discussed in Section 5.2. 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, Lin and Mano, 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, consuming all their income. 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, Nakamura and Steinsson, 2016; Acharya and Dogra, 2020; Bilbiie, 2021), social insurance (McKay and Reis, 2016, 2021; Kekre, 2022), and fiscal policy (Auclert, Rognlie and Straub, 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 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. All proofs are in the Online Appendix.

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 proba-

[2] See also Nisticò (2016), who generalizes the Two-Agent New Keynesian (TANK) model of Galí, Lopez-Salido and Valles (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.

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

$$\max_{\{c_t^s(i),\ell_t^s(i),a_t^s(i)\}} \quad \mathbb{E}_s \sum_{t=s}^{\infty} (\beta\vartheta)^{t-s}\, u\Big(c_t^s(i),\ell_t^s(i);\xi_t^s(i)\Big) \tag{1}$$

subject to

$$c_t^s(i)+q_t a_{t+1}^s(i) = (1-\tau^w)\widetilde{w}_t\ell_t^s(i)+(1-\tau_t^a)a_t^s(i)+D_t^s(i)-T_t \tag{2}$$

$$a_s^s(i) = 0 \tag{3}$$

Agents have CARA preferences over both consumption c and labor ℓ:

$$u\Big(c_t^s(i),\ell_t^s(i);\xi_t^s(i)\Big) = -\frac{1}{\gamma}e^{-\gamma c_t^s(i)} - \rho e^{\frac{1}{\rho}\left(\ell_t^s(i)-\xi_t^s(i)\right)} \tag{4}$$

Each agent i saves in riskless real actuarial bonds which have a pre-tax payoff of one unit of the consumption good at $t+1$ if the agent survives. These are issued by financial intermediaries at a price q_t. 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] 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 $\widetilde{w}_t$ and tax rate τ^w, a household supplies labor $\ell_t^s(i)$ at the post-tax real wage $w_t=(1-\tau^w)\widetilde{w}_t$. Households face uninsurable shocks $\xi_t^s(i)\sim N\left(\overline{\xi},\sigma_t^2\right)$ 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 (4) as $-e^{-\gamma c_t^s(i)}/\gamma-\rho e^{-l_t^s(i)/\rho}$ and the budget constraint as

$$c_t^s(i)+w_t l_t^s(i)+q_t a_{t+1}^s(i) = w_t\xi_t^s(i)+(1-\tau_t^a)a_t^s(i)+D_t^s(i)-T_t \tag{5}$$

The LHS of (5) 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.

[3] As we discuss in Section 2.1, if we had infinitely lived agents, our model would not feature a stationary wealth distribution.
[4] Appendix H introduces a fraction of households (the Hand to Mouth) who cannot access asset markets.
[5] We thank Gianluca Violante for suggesting this interpretation.

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:

$$\max_{a_{t+1},B_{t+1}} -\vartheta a_{t+1} + \frac{B_{t+1}}{P_{t+1}} \quad \text{s.t.} \quad -P_t q_t a_{t+1} + \frac{B_{t+1}}{1+i_t} \leq 0 \tag{6}$$

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

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:

$$y_t(j) = \left(\frac{P_t(j)}{P_t}\right)^{-\varepsilon_t} y_t \tag{7}$$

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

$$y_t(j) = z_t n_t(j) - \frac{\Psi}{2}\left(\frac{P_t(j)}{P_{t-1}(j)} - 1\right)^2 y_t \tag{8}$$

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 $\tau^\star$, so that firm j solves

$$\max_{\{P_t^j, n_t^j, y_t^j\}_{t=0}^\infty} \sum_{t=0}^\infty \beta^t \left\{\frac{P_t(j)}{P_t} y_t(j) - (1-\tau^\star)\widetilde{w}_t n_t(j)\right\} \tag{9}$$

subject to (7) and (8). This yields the standard Phillips curve:

$$(\Pi_t - 1)\Pi_t = \frac{\varepsilon_t}{\Psi}\left[1 - \frac{(\varepsilon_t - 1)z_t}{\varepsilon_t (1-\tau^\star)\widetilde{w}_t}\right] + \beta\left(\frac{z_t y_{t+1}\widetilde{w}_{t+1}}{z_{t+1} y_t \widetilde{w}_t}\right)(\Pi_{t+1} - 1)\Pi_{t+1} \tag{10}$$

1.5 Government

The monetary authority sets the interest rate on nominal government debt. The wage bill subsidy is assumed to be equal to $\tau^\star = \varepsilon^{-1}$ 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 at a rate τ^w. The government budget constraint is:

$$\frac{1}{R_t}\frac{B_{t+1}}{P_{t+1}} + T_t + (1-\vartheta)\sum_{s=-\infty}^{t}\vartheta^{t-s}\int_i [\tau^w \widetilde{w}_t \ell_t^s(i)di + \tau_t^a a_t^s(i)]\,di = \tau^\star w_t \int_0^1 n_t(j)dj + \frac{B_t}{P_t} \tag{11}$$

and we assume in what follows that $B_t = 0$.[6]

1.6 Market clearing

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

$$y_t = c_t \equiv (1-\vartheta)\sum_{s=-\infty}^{t}\vartheta^{t-s}\int_i c_t^s(i)di \tag{12}$$

$$\int_0^1 n_t(j)dj = (1-\vartheta)\sum_{s=-\infty}^{t}\vartheta^{t-s}\int_i \ell_t^s(i)di \tag{13}$$

$$0 = \frac{B_{t+1}}{\vartheta P_{t+1}} = a_{t+1} = (1-\vartheta)\sum_{s=-\infty}^{t}\vartheta^{t-s}\int_i a_{t+1}^s(i)di \tag{14}$$

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.

1.8 Equilibrium

Given initial conditions, and a sequence of shocks $\{z_t, \varepsilon_t\}_{t=0}^\infty$ and a monetary-fiscal policy $\{i_t, \tau_t^a, T_t\}_{t=0}^\infty$ and τ^w, τ^*, an equilibrium is a sequence $\{\{c_t^s(i), \ell_t^s(i), a_{t+1}^s(i), D_t^s(i)\}_{i,s}, \{P_t(j), y_t(j), n_t(j)\}_j, y_t, c_t, B_{t+1}, a_{t+1}, P_t, \widetilde{w}_t, \Pi_t, q_t, D_t\}_{t=0}^\infty$ that satisfies: (i) households choose $\{c_t^s(i), \ell_t^s(i), a_t^s(i)\}_{i,s,t}$ maximize (1) subject to (2) and (3); (ii) financial intermediaries choose $\{B_{t+1}, a_{t+1}\}_t$ to maximize (6); (iii) intermediate goods producers choose $\{P_t(j), y_t(j), n_t(j)\}_{j,t}$ to maximize (9) subject to (7), (8) and satisfy $P_t(j) = P_t$, $y_t(j) = y_t$ for all $j \in [0,1]$; (iv) dividends satisfy $D_t^s(i) = D_t = y_t - (1-\tau^\star)\widetilde{w}_t n_t$, inflation is defined as $\Pi_t = P_t/P_{t-1}$ and (v) market clearing conditions (12)-(14) are satisfied.

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

[6] Previous versions showed that the path of government debt is irrelevant for real allocations and optimal monetary policy provided that the new-born households receive a transfer equal to average household wealth.

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.[7]

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

$$c_t^s(i) = \mathcal{C}_t + \mu_t x_t^s(i) \tag{15}$$

$$\ell_t^s(i) = \rho \ln w_t - \gamma\rho c_t^s(i) + \xi_t^s(i) \tag{16}$$

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

$$\mathcal{C}_t = -\frac{1}{\gamma}\ln \beta R_t + \mathcal{C}_{t+1} - \frac{\gamma\mu_{t+1}^2 w_{t+1}^2 \sigma_{t+1}^2}{2} \tag{17}$$

$$\mu_t^{-1} = 1 + \gamma\rho w_t + \frac{\vartheta}{R_t}\mu_{t+1}^{-1} \tag{18}$$

Proof. See Appendix A. □

To understand how market incompleteness affects consumption and labor supply, it is useful to compare (15) and (16) 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 $\partial c_t^s(i)/\partial \xi_t^s(i) = 0$ and $\partial \ell_t^s(i)/\partial \xi_t^s(i) = 1$: a household with a temporarily higher disutility from working ($\xi_t^s(i) < \overline{\xi}$) can reduce hours without a fall in consumption. Instead, when markets are incomplete (15) and (16) imply that

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

A household with $\xi_t^s(i) < \overline{\xi}$ 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 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; (18) describes its evolution. Iterating this forwards yields

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

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,

[7] Since 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.

making it easier for a household with $\xi_t^s(i) < \overline{\xi}$ 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 $\mathcal{C}_t$ depends on interest rates relative to impatience and on households' precautionary motive, as shown in (17). Absent idiosyncratic risk, $\sigma_t = 0$ in (17) and we revert to the RANK Euler equation; higher real interest rates relative to household impatience raise consumption growth. The last term in (17) reflects precautionary savings. Given (15), the conditional variance of date $t+1$ consumption of household i is $\mathbb{V}_t\left(c_{t+1}^s(i)\right) = \mu_{t+1}^2 w_{t+1}^2 \sigma_{t+1}^2$. 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 $\mathbb{V}_t\left(x_{t+1}^s(i)\right) = w_{t+1}^2 \sigma_{t+1}^2$, 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 (8) across firms, we have $y_t = z_t n_t - \frac{\Psi}{2}\left(\Pi_t - 1\right)^2 y_t$. Similarly, aggregating (16) and using (12) and (13):

$$n_t = \rho \ln w_t - \gamma \rho y_t + \overline{\xi} \tag{19}$$

Combining these two equations, we have:

$$y_t = z_t \frac{\rho \ln w_t + \overline{\xi}}{1 + \gamma \rho z_t + \frac{\Psi}{2}\left(\Pi_t - 1\right)^2} \tag{20}$$

where $\frac{\Psi}{2}\left(\Pi_t - 1\right)^2$ denotes the resource cost of inflation.

Deriving the aggregate IS equation Imposing goods market clearing in (17) yields the aggregate IS equation which describes the relation between output today and tomorrow:

$$y_t = y_{t+1} - \frac{1}{\gamma} \ln \beta \left(\frac{1 + i_t}{\Pi_{t+1}}\right) - \frac{\gamma}{2} \mu_{t+1}^2 w_{t+1}^2 \sigma_{t+1}^2 \tag{21}$$

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 cyclicality of wages, i.e. it would be *procyclical*. In contrast, the empirical literature (Storesletten, Telmer and Yaron, 2004; Nakajima and Smirnyagin, 2019) generally finds that earnings risk is countercyclical. We assume $\sigma_t^2 w_t^2 = \sigma^2 w^2 \exp\left\{2\varphi(y_t - y)\right\}$ where y denotes steady state output and $\varphi = \frac{\partial \ln \mathbb{V}(x)}{\partial y}$ is the semi-elasticity of the variance of cash-on-hand $\mathbb{V}_t(x)$ w.r.t output. This allows $\mathbb{V}_t(x)$ to be 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 $\varphi = 0$.[8] Importantly, what we mean by cyclicality 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

[8] More 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.

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 $\tau^\star = \varepsilon^{-1}$, imposing $\Pi_t = \Pi_{t+1} = 1$ in (10) requires that $\widetilde{w} = 1$ and $w = 1 - \tau^w$; steady state output is $y = \frac{\rho \ln w + \bar{\xi}}{1+\gamma\rho}$. Imposing steady state in (18) and (21) yields

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

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 $\widetilde{\beta} = \vartheta/R$ 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

$$F(x) = (1-\vartheta) \sum_{s=0}^{\infty} \vartheta^s \Phi\left(\frac{x}{w\sigma\sqrt{s+1}}\right), \tag{22}$$

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 ($\vartheta \to 1$), the sum in (22) 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 (21), the MPC recursion (18), the definition of GDP (20) and the Phillips curve (10). 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 (20) to substitute out for wages, we have:

$$\widehat{y}_t = \Theta \widehat{y}_{t+1} - \frac{1}{\gamma y}\left(\widehat{i}_t - \pi_{t+1}\right) - \frac{\Lambda}{\gamma y}\widehat{\mu}_{t+1} \tag{23}$$

$$\widehat{\mu}_t = -\gamma\mu w y\Big[(1+\gamma\rho)\,\widehat{y}_t - \widehat{z}_t\Big] + \widetilde{\beta}(\widehat{\mu}_{t+1} + \widehat{i}_t - \pi_{t+1}) \tag{24}$$

$$\pi_t = \beta\pi_{t+1} + \kappa\,(\widehat{y}_t - \widehat{y}^n_t) \tag{25}$$

where $\widehat{i}_t = \ln(1+i_t) - \ln R$, $\Theta = 1 - \frac{\Lambda\varphi}{\gamma}$, $\kappa = \frac{\varepsilon}{\Psi}\frac{1+\gamma\rho}{\rho/y}$, $\widehat{y}^n_t = \frac{(1+\rho/y)\widehat{z}_t - (\rho/y)\widehat{\varepsilon}_t}{1+\gamma\rho}$ is the log deviation from steady state of the "natural" level of output i.e. which would prevail under flexible prices ($\Psi = 0$), and $\widehat{\varepsilon}_t = \ln\left(\frac{\varepsilon_t}{\varepsilon_t - 1}\right) - \ln\left(\frac{\varepsilon}{\varepsilon-1}\right)$.[9] In RANK, there is no idiosyncratic risk, i.e. $\sigma^2 = 0$ which implies $\Theta = 1$ and $\Lambda = 0$, so that (23) 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

[9] Note 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.

$\Theta > 1$, 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 (25) 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 cyclicality, we calibrate the equilibrium of the HANK economy in which the labor income tax is absent ($\tau^w = 0$). We choose $\overline{\xi}$ 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, Ozkan and Song (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 cyclicality of income risk $\varphi = -5.76$ which is broadly consistent with Storesletten, Telmer and Yaron (2004).[10] 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 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 $\varrho_z = 0.95^4$ and $\varrho_\varepsilon = 0.9^4$ (Bayer, Born and Luetticke, 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, Born and Luetticke (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

[10] Storesletten, Telmer and Yaron (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 $y\frac{d\sigma_y}{dy} = \frac{0.12-0.21}{0.03} = -3$. Using $\sigma_{y,t} = (1 - \gamma\rho\mu_t w_t)\, w\sigma e^{\varphi(y_t - y)}$, the equilibrium relationship between μ_t, w_t and y_t, and because we are calibrating cyclicality of income risk in the economy with $\tau^w = 0$, we have:

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

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

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:[11]

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

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

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

$$\mathbb{U}_t=u\Big(c_t,n_t;\overline{\xi}\Big)\times\Sigma_t\qquad\text{where}\qquad\Sigma_t=(1-\vartheta)\sum_{s=-\infty}^{t}\vartheta^{t-s}e^{\frac{1}{2}\gamma^2\sigma_c^2(s,t)}\tag{26}$$

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, $u(c_t,n_t;\overline{\xi})$ 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 $\overline{\xi}$. 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 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

$$\ln\Sigma_t=\frac{\gamma^2}{2}\mu_t^2w_t^2\sigma_t^2+\ln\left[1-\vartheta+\vartheta\Sigma_{t-1}\right]\tag{27}$$

$$\text{with}\qquad\ln\Sigma_0=\frac{\gamma^2}{2}\mu_0^2w_0^2\sigma_0^2+\underbrace{\ln\left[\frac{1-\vartheta}{1-\vartheta e^{-\frac{\Lambda}{2}\left(1-\tau_0^a\right)^2\left(\frac{\mu_0}{\mu}\right)^2}}\right]}_{\text{welfare cost of pre-existing wealth inequality}}\tag{28}$$

The evolution of consumption inequality is an increasing function of consumption risk $\mu_t^2w_t^2\sigma_t^2$, 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 (27).[12] 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.

[11] Note 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.

[12] Within-cohort consumption dispersion $\sigma_c^2(t,s)$ rises without bounds as the cohort ages (i.e., as $t-s\to\infty$) 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., $\sigma_c^2(t,t)=\mu_t^2w_t^2\sigma_t^2$), Σ_t does not necessarily diverge. In fact, provided that the survival rate $\vartheta<e^{-\Lambda/2}$, Σ_t is stationary.

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 $\sum_{t=0}^{\infty} \beta^t u\left(y_t, n_t; \overline{\xi}\right) \Sigma_t$ absent aggregate shocks, and given an initial wealth distribution. The constraints faced by the planner are the aggregate Euler equation (21), aggregate labor supply (19), the evolution of μ_t (18), the Phillips curve (10), the evolution of Σ_t (27) and the relationship between GDP and wages (20). 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 $\sum_{t=0}^{\infty} \beta^t u\left(y_t, n_t; \overline{\xi}\right) \Sigma_t$ subject to the constraints (21), (19), (18), (10), (27) and (20). 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.[13] In the RANK version of our economy, $\sigma = 0$ and (27) 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.

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 (28) to (27) 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 (28):

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

[13] Since 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.

While the first two terms on the RHS above are the same as that in (27), 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 (28). To be clear, anticipated cuts in rates *do* reduce inequality as discussed earlier: lower μ_t reduces Σ_t in equation (27). 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.[14] 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 (28) 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 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 $t > 0$.[15] 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 $\overline{\xi}/(1+\gamma\rho)$ – 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.

[14] A 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.

[15] This 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 $t > 0$.

Appendix D.1 shows that this tradeoff can be summarized by the following optimality condition:

$$\underbrace{\Omega}_{\substack{\text{benefit from reduction}\\ \text{of consumption risk}}} \equiv \underbrace{\frac{\Lambda}{1-\Lambda}}_{\substack{\text{reduction of consumption risk}\\ \text{due to reduced passthrough}}} + \underbrace{\frac{\Theta-1}{1-\Lambda}}_{\substack{\text{reduction of consumption risk}\\ \text{due to reduced income risk}}} = \underbrace{\frac{1-\widetilde{\beta}}{1+\gamma\rho w}(w-1)}_{\substack{\text{cost of deviating from}\\ \text{productive efficiency}}}, \tag{29}$$

which implies that the optimal income tax is $\tau^w = 1 - \frac{1-\widetilde{\beta}+\Omega}{1-\widetilde{\beta}-\gamma\rho\Omega}$. Ω denotes the sensitivity of *consumption* risk to aggregate output (we will also refer to this as the *cyclicality* of consumption risk) , and summarizes the benefit from a reduction in consumption inequality due to higher output. In RANK ($\Lambda = 0, \Theta = 1$), there is no risk and thus no benefit from reducing inequality ($\Omega = 0$), so that $w = 1$ or $\tau^w = 0$ is optimal. In the presence of income risk, higher output (implemented via lower τ^w) affects consumption risk through both a *self-insurance* channel and an *income-risk channel*. (29) states that at an optimum, the marginal benefit of lower consumption risk 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 (which is equal to 1 in steady state).

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 (29) 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 income risk, we have $\Omega = \Omega^c \equiv \frac{\Lambda}{1-\Lambda} > 0$ and consumption risk is *countercyclical*. Consequently, 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 risk even for a fixed μ. In addition, higher output reduces μ, further reducing consumption risk. Thus, the benefit from higher output is even larger than if $\Theta = 1$ – both LHS components in (29) are positive and consumption risk is more countercyclical ($\Omega > \Omega^c$). 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 risk 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.

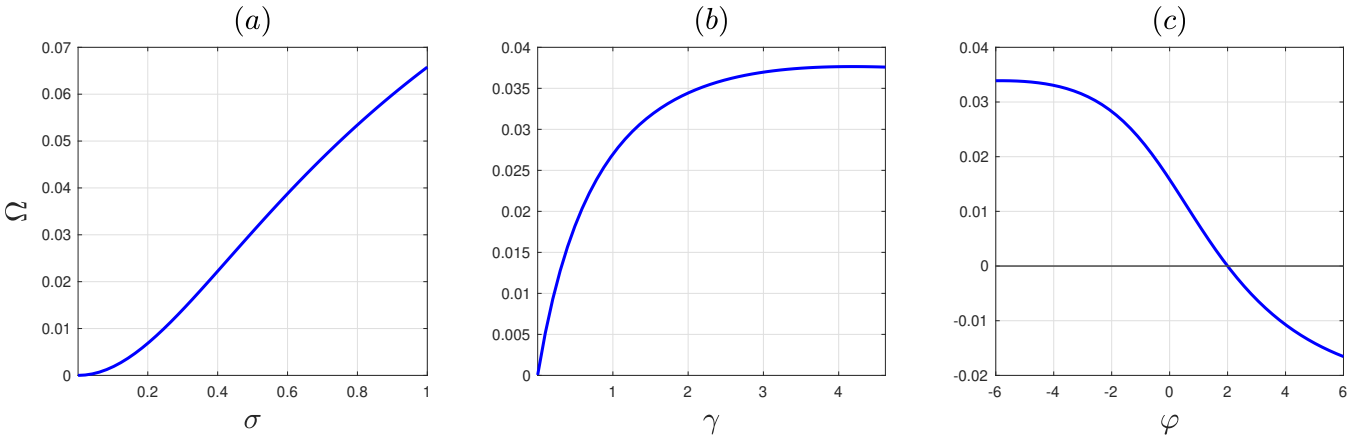

Figure 1: **Comparative Statics of Ω.** The curves plot the values of Ω for different values of σ, φ, γ, in each case holding all parameters other than that on the x-axis fixed at their levels in our baseline calibration with countercyclical risk ($\varphi < 0$).

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 (25), setting $\widehat{y}_t = \widehat{y}_t^n$ 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 $\widehat{y}_t - \widehat{y}_t^e$, where $\widehat{y}_t^e$ does not respond to inefficient cost-push shocks:

$$\widehat{y}_t^e = \frac{1+\rho/y}{1+\gamma\rho}\widehat{z}_t \qquad \text{so that} \qquad \widehat{y}_t^n = \widehat{y}_t^e - \frac{\varepsilon}{\kappa\Psi}\widehat{\varepsilon}_t$$

This implies that the Phillips curve (25) can equivalently be written as

$$\pi_t \quad = \beta\pi_{t+1} + \kappa(\widehat{y}_t - \widehat{y}_t^e) + \tfrac{\varepsilon}{\Psi}\widehat{\varepsilon}_t \tag{30}$$

To 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 $\frac{1+\gamma\rho}{\rho/y}(\widehat{y}_t - \widehat{y}_t^e)$, i.e., it is proportional to the output gap. In what follows, with some abuse of terminology, we refer to $\widehat{y}_t^e$ as the productively efficient level of output.[16]

[16]To be clear, a zero output gap $\widehat{y}_t - \widehat{y}_t^e$ 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.

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 (27) 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 (27) and using our assumptions about $w_t\sigma_t$, we have

$$\widehat{\Sigma}_t = \underbrace{\Lambda\widehat{\mu}_t - \gamma y\left(\Theta-1\right)\widehat{y}_t}_{\text{consumption risk}} + \beta^{-1}\widetilde{\beta}\widehat{\Sigma}_{t-1} \tag{31}$$

(31) 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 $\widehat{\mu}_t$. 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 $-\gamma y(\Theta-1)\widehat{y}_t$. For example, with countercyclical income risk $\Theta > 1$, raising output $\widehat{y}_t$ reduces income and hence, consumption risk.

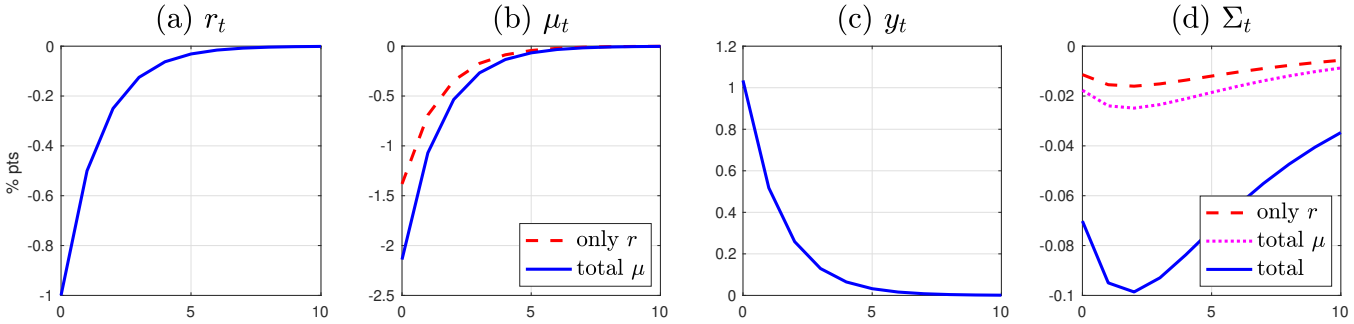

Figure 2: **The effect of monetary policy on consumption inequality** Σ_t. blue curves in panels b,c and d depict the reaction of μ_t, y_t and Σ_t respectively to a mean reverting cut in real interest rates depicted in panel a. The red-dashed line in panel b and d depict the effect of self-insurance 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 $\times 100$.

But the planner cannot vary $\widehat{y}_t$ and $\widehat{\mu}_t$ 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 $\widehat{r}_t = (0.5)^t\widehat{r}_0$ 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 $\widehat{\mu}_t$ 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 $\widehat{\mu}_t$ further: the blue line in panel c shows the total effect on passthrough via both real interest rates and wages. This 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 $\widehat{\Sigma}_t$ 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, Lin and Mano, 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, the effect of monetary policy via the self-insurance channel could potentially be 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 $\widehat{y}_t$, the overall effect of monetary policy on consumption inequality can be summarized by a relationship between $\widehat{\Sigma}_t$ and $\widehat{y}_t$, using the IS equation (23) and the μ recursion (24) to eliminate $\widehat{\mu}_t$ in (31). The coefficient on $\widehat{y}_t$ 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,* $\widehat{\Sigma}_t$ *evolves according to*

$$\widehat{\Sigma}_t = \underbrace{-\gamma y \Omega \left(\widehat{y}_t - \varkappa(\Omega)\, \widehat{y}_t^e \right)}_{\text{consumption risk}} + \beta^{-1} \widetilde{\beta} \widehat{\Sigma}_{t-1} \tag{32}$$

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

Proof. See Appendix E.1. □

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 $-\gamma y \Omega$, the *cyclicality of consumption risk* (multiplied by the coefficient of relative risk aversion γy). In other words, changes in interest rates affect output through the IS equation (23), and output in turn affects consumption risk through (32). 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: $\partial \widehat{\Sigma}_t / \partial \widehat{y}_t = -\gamma y \Omega < 0$. Even when income risk is acyclical ($\Theta = 1$), expansionary policy still reduces passthrough, i.e., consumption risk is still countercyclical, $-\gamma y \Omega < 0$. When income risk is strongly procyclical ($\Theta << 1 \Rightarrow \Omega < 0$), consumption risk is also procyclical: higher output increases inequality as lower passthrough is outweighed by higher income risk, $\partial \widehat{\Sigma}_t / \partial \widehat{y}_t > 0$. 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 $\widehat{y}_t = \varkappa(\Omega) \widehat{y}_t^e$, where $\varkappa(\Omega) < 1$, i.e., by moving output less than one-for-one with the productively efficient level $\widehat{y}_t^e$. Absent aggregate productivity

[17] If 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.

shocks ($\widehat{y}_t^e = 0$), consumption risk depends only on the level of output $\widehat{y}_t$, and is perfectly stabilized by setting $\widehat{y}_t = 0$. Stabilizing output perfectly keeps income risk constant; it also keeps real interest rates and wages constant, implying an unchanged passthrough $\widehat{\mu}_t = 0$. 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 ($\widehat{y}_t^e$), keeping output constant would require higher real wages to increase labor supply and compensate for the lower productivity. Higher wages would reduce passthrough $\widehat{\mu}_t < 0$, *reducing* consumption risk. However, letting output $\widehat{y}_t$ fall as much as its productively efficient level $\widehat{y}_t^e$ 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 $\widehat{y}_t^e$ – 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. The presence of consumption inequality means that the HANK economy is not at its first-best level in the zero inflation steady state. Consequently, a naive LQ approach (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, following Benigno and Woodford (2005), we eliminate these linear terms using a second order approximation of the constraints.[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 $\{\widehat{y}_t, \pi_t\}_{t=0}^{\infty}$ to minimize the loss function subject to the linearized Phillips curve (30).

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

$$\min_{\{\widehat{y}_t, \pi_t\}_{t=0}^{\infty}} \frac{1}{2} \sum_{t=0}^{\infty} \beta^t \left\{ \Upsilon(\Omega) \left(\widehat{y}_t - \delta(\Omega)\, \widehat{y}_t^e \right)^2 + \frac{\varepsilon}{\kappa} \pi_t^2 \right\} \tag{33}$$

$$s.t. \qquad \pi_t = \beta \pi_{t+1} + \kappa \left(\widehat{y}_t - \widehat{y}_t^e \right) + \frac{\varepsilon}{\Psi} \widehat{\varepsilon}_t,$$

where $\Upsilon(\Omega)$ *and* $\delta(\Omega)$ *are defined in Appendix E.2 and satisfy* $\Upsilon(0) = \delta(0) = 1$*. When income risk is*

[18] McKay and Wolf (2022) also use an LQ approach to characterize optimal policy in HANK, and address this issue in a different way. They consider the problem of a planner whose objective function aggregates the flow-utility of households using quasi-Pareto-weights which depend on households' histories of idiosyncratic shocks, i.e., the planner's objective is not a weighted sum of households' lifetime utilities. Specifically, these weights are chosen so that the planner regards allocations in the steady state with idiosyncratic but no aggregate risk as optimal. This assumption ensures that the quadratic approximation to the planner's objective function contains no first-order terms, and the naive LQ approach yields first-order accurate approximations to the allocations which solve their planners problem. However, these allocations are not generally Pareto-optimal since the planner's objective function does not respect households' individual preferences, i.e., their planner does not necessarily prefer feasible allocations which make all households better off (in particular by improving risk-sharing).

acyclical or countercyclical $\left(\Theta \geq 1 \Rightarrow \Omega \geq \Omega^c = \frac{\Lambda}{1-\Lambda} > 0\right)$, $\Upsilon(\Omega) > 1$ *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*

$$\min_{\{\widehat{y}_t, \pi_t\}_{t=0}^{\infty}} \frac{1}{2} \sum_{t=0}^{\infty} \beta^t \left\{ \left(\widehat{y}_t - \widehat{y}_t^e\right)^2 + \frac{\varepsilon}{\kappa} \pi_t^2 \right\}, \tag{34}$$

$$s.t. \qquad \pi_t = \beta \pi_{t+1} + \kappa \left(\widehat{y}_t - \widehat{y}_t^e\right) + \frac{\varepsilon}{\Psi} \widehat{\varepsilon}_t$$

The HANK and RANK planners in (33) and (34), 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 (34) is a special case of (33): 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 (34)) and price stability, which would be attained by setting $\pi_t = 0$ (last term in (34)).

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 (33) can be written as

$$\Upsilon(\Omega) \left(\widehat{y}_t - \delta(\Omega) \widehat{y}_t^e\right)^2 = \Upsilon(\Omega) \Big([1 - \varpi(\Omega)] \underbrace{\left(\widehat{y}_t - \varkappa(\Omega) \widehat{y}_t^e\right)}_{\substack{\text{consumption} \\ \text{risk}}} + \varpi(\Omega) \underbrace{\left(\widehat{y}_t - \widehat{y}_t^e\right)}_{\substack{\text{productive} \\ \text{efficiency}}} \Big)^2, \tag{35}$$

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 $\widehat{y}_t - \varkappa(\Omega)\widehat{y}_t^e$ where $\varkappa(\Omega) < 1$. In other words, consumption risk would be perfectly stabilized at its steady state level by setting $\widehat{y}_t = \varkappa(\Omega)\widehat{y}_t^e$, adjusting output less than one-for-one with changes in its productively efficient level $\widehat{y}_t^e$. The first term in the loss function 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 $\widehat{y}_t - \varkappa(\Omega)\widehat{y}_t^e$ and the output gap, $\widehat{y}_t - \widehat{y}_t^e$, with weights $1 - \varpi(\Omega)$ and $\varpi(\Omega)$ respectively. Since $\delta(\Omega) = \varpi(\Omega) + \left(1 - \varpi(\Omega)\right)\varkappa(\Omega)$ 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 $\widehat{y}_t$ less than one-for-one with $\widehat{y}_t^e$ – 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

[19] Again, 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.

output stabilization:

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

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 $\Upsilon(\Omega)\left(\widehat{y}_t - \delta(\Omega)\,\widehat{y}_t^e\right)^2$ 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.

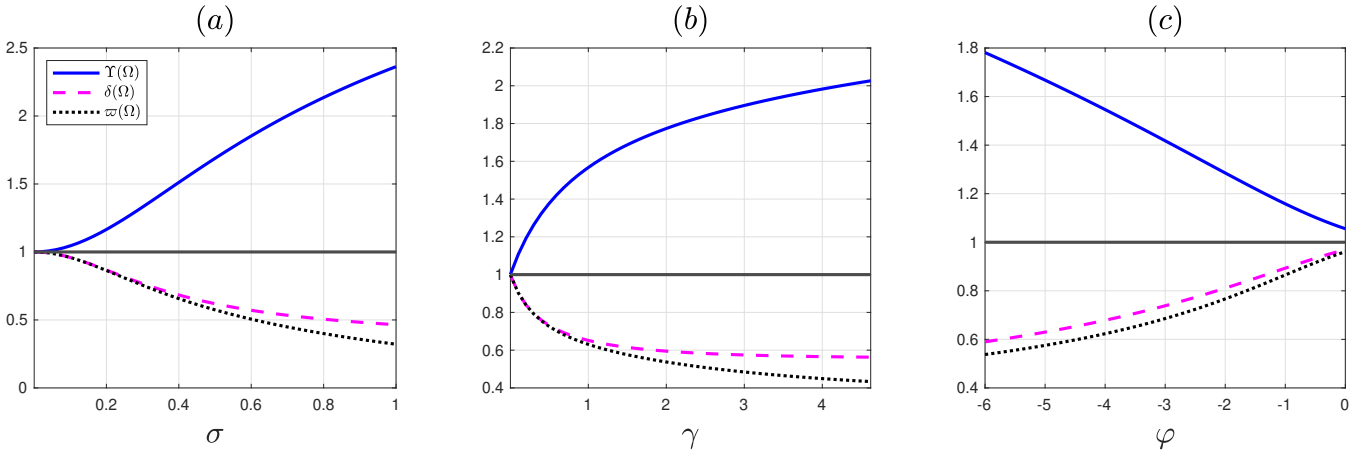

Figure 3: **Comparative Statics:** The blue curves denotes the values of Υ, the magenta-dashed curves denote δ and the black-dotted curves denote ϖ for different values of σ, φ, γ.

Figure 3 plots Υ (solid blue line) and δ (dashed-magenta line) as functions of income risk σ, risk aversion γ and cyclicality of 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 the cyclicality of consumption risk Ω, the sufficient statistic in equation (32). 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.[20] 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 (32)). 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.[21]

Lemma 2. *In HANK with $\Omega = 0$, the planner's objective function becomes (34) as in RANK. Consequently, optimal policy implements the same sequence $\{\widehat{y}_t, \pi_t\}$ 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 (33)), 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 $t \geq 0$:*

$$\Big(1 - \delta(\Omega)\Big)\widehat{y}_t + \delta(\Omega)\Big(\widehat{y}_t - \widehat{y}_t^e\Big) + \frac{\varepsilon}{\Upsilon(\Omega)}\widehat{p}_t = 0 \tag{36}$$

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

$$(\widehat{y}_t - \widehat{y}_t^e) + \varepsilon\widehat{p}_t = 0. \tag{37}$$

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 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.[22]

[20] As discussed in Section 3.5, Ω summarizes the effect of monetary policy on consumption risk via both the income risk and self-insurance channels. Even if income risk is acyclical, $\Omega = \Omega^c > 0$ and the planner puts some weight on stabilizing the level of output, because stabilizing output mitigates fluctuations in passthrough μ_t and hence consumption risk.

[21] Werning (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.

[22] Bilbiie (2008) derives a related optimal targeting rule in TANK, where optimal policy seeks to stabilize inequality arising

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 $\tau^\star$ 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 ($\widehat{p}_t = 0$) while closing the output gap ($\widehat{y}_t - \widehat{y}_t^e = 0$). Intuitively, maintaining output at its productively efficient level also keeps prices stable. This can be seen from the target criterion (37) along with the Phillips curve (30), which (given $\widehat{\varepsilon}_t = 0$) imply $\widehat{p}_t = \widehat{y}_t - \widehat{y}_t^e = 0$. 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 $\widehat{y}_t^e = \frac{1+\rho/y}{1+\gamma\rho}\widehat{z}_t < 0$ by tracking this level, $\widehat{y}_t = \widehat{y}_t^e < 0$, 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 $\widehat{y}_t = \widehat{y}_t^e$ 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 $\widehat{y}_t$ from falling as much as the flexible-price level of output $\widehat{y}_t^e$ initially. This entails positive inflation initially. In contrast, the planner commits to mildly negative output gaps ($\widehat{y}_t < \widehat{y}_t^e < 0$) 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 ($\widehat{z}_0 < 0$), at date 0, $\widehat{y}_0$ falls less than $\widehat{y}_0^e$ 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 $\widehat{y}_t < \widehat{y}_t^e$. Following an increase in productivity all these signs are reversed, i.e., π_t and $\widehat{y}_t - \widehat{y}_t^e$ 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 $\widehat{y}_t$ above $\widehat{y}_t^e = 0$ absent aggregate shocks. With acyclical or countercyclical income risk ($\Omega \geq \Omega_c > 0$), increasing $\widehat{y}_t$ 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 $\widehat{y}_t = \widehat{y}_t^e < 0$ $\forall t \geq 0$ (also implying $\pi_t = 0$ $\forall t \geq 0$). The fall in $\widehat{y}_t$ 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 $\widehat{y}_t^e$ to curtail the rise in inequality. Meanwhile, at $\widehat{y}_t = \widehat{y}_t^e$ the cost of marginally increasing output above $\widehat{y}_t^e$, measured by the output gap, $\widehat{y}_t - \widehat{y}_t^e$, remains unchanged. Since the benefit of increasing $\widehat{y}_t$ above $\widehat{y}_t^e$ increases while the cost of doing so remains unchanged, the planner sets $0 > \widehat{y}_t > \widehat{y}_t^e$. Output still falls on impact, but by less than the flexible-price level of output $\widehat{y}_t^e$, implying a positive output gap (blue curve

from unequal access to asset markets, rather than idiosyncratic risk. Importantly, while the TANK planner puts a lower weight on the price level relative to economic activity, the planner only tries to stabilize the output gap and not the level of output.

in panel (a)). This tradeoff is also reflected in the target criterion (36), which can be rewritten as:

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

Intuitively, rather than tracking $\widehat{y}_t^e$ one-for-one, which would stabilize the output gap, the planner seeks to minimize the gap between $\widehat{y}_t$ and $\delta(\Omega)\widehat{y}_t^e$ (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.

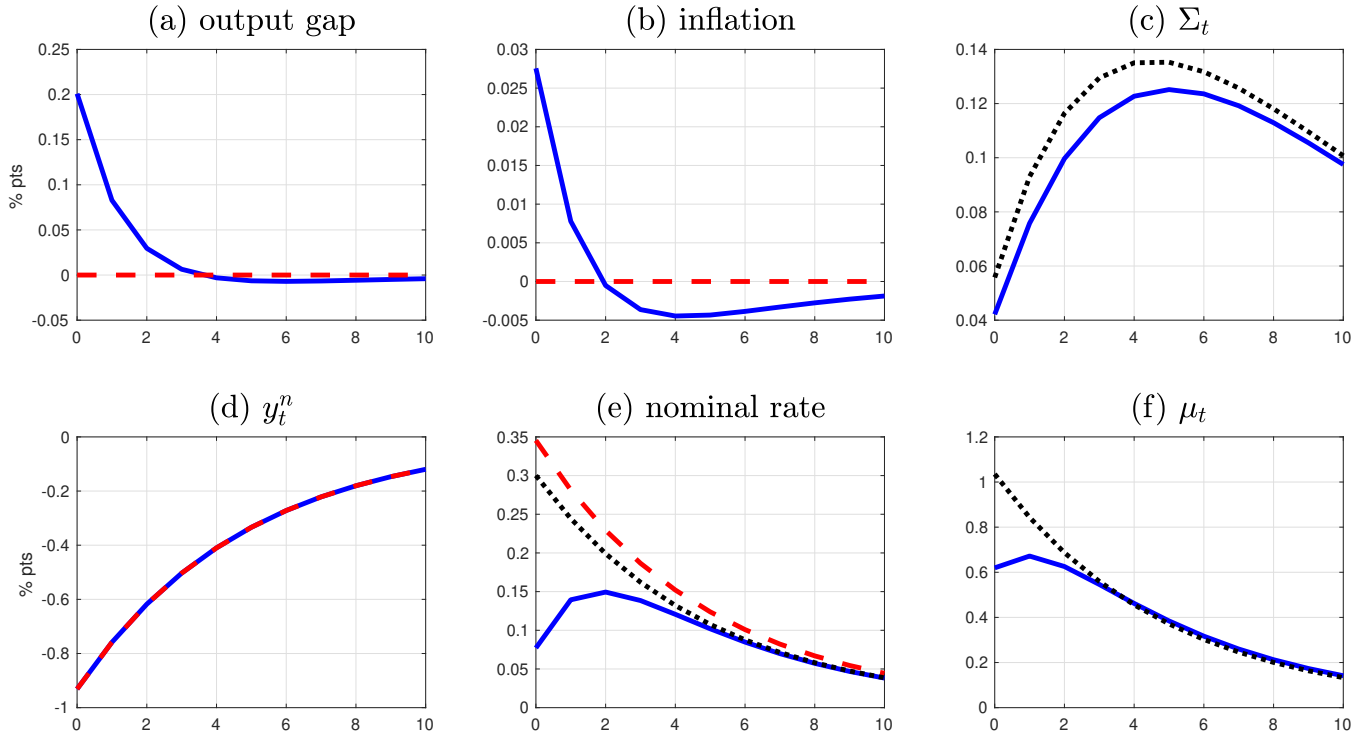

Figure 4: **Optimal policy in response to productivity shocks** in HANK (solid blue curves) and RANK (dashed red curves). Black-dotted curves denote outcomes in HANK under the non-optimal policy which sets $\widehat{y}_t = \widehat{y}_t^n, \pi_t = 0$ for all $t \geq 0$. All panels plot log-deviations 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 $\widehat{y}_t = \widehat{y}_t^e$ and $\pi_t = 0$ (black-dashed curve in panel (f)). Given the higher path of $\widehat{y}_t$ and lower path of $\widehat{\mu}_t$, while inequality still increases (blue curve in panel (c)), it is lower than it would have been, had the planner implemented $\widehat{y}_t = \widehat{y}_t^e$ and $\pi_t = 0$ for all $t \geq 0$ (black-dotted curve in panel (c)).

Implementing $\widehat{y}_t > \widehat{y}_t^e$ 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 (36) (since $\Upsilon(\Omega) > 1$). Nonetheless, to mitigate this rise in inflation, the planner commits to set $\widehat{y}_t$ slightly below $\widehat{y}_t^e$ 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 (36) is simply $\widehat{y}_t - \widehat{y}_t^e = \widehat{y}_t$. 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 $\widehat{y}_t^e$ (which remains unchanged) and the level of output consistent with zero inflation i.e. $\widehat{y}_t^n = -\frac{\varepsilon}{\kappa\Psi}\widehat{\varepsilon}_t$. Keeping $\pi_t = 0$ by setting $\widehat{y}_t = \widehat{y}_t^n < 0$ is not optimal as this would entail too large a fall in output relative to its efficient level $\widehat{y}_t^e = 0$. Conversely, keeping output at its efficient level $\widehat{y}_t = \widehat{y}_t^e = 0$ 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 $\widehat{y}_t$ below $\widehat{y}_t^n$ 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 ($\widehat{y}_t = 0$) 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 (36) to the case with only markup shocks (implying $\widehat{y}_t^e = 0$):

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

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 (γy) and Frisch elasticity (ρ/y) in steady state. Under optimal policy in HANK, following an increase in firms' desired markup ($\widehat{\varepsilon}_0 > 0$), at date 0, $\widehat{y}_0$ falls (but less than $\widehat{y}_0^n$) 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 $\widehat{y}_t - \widehat{y}_t^n < 0$. 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 an 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

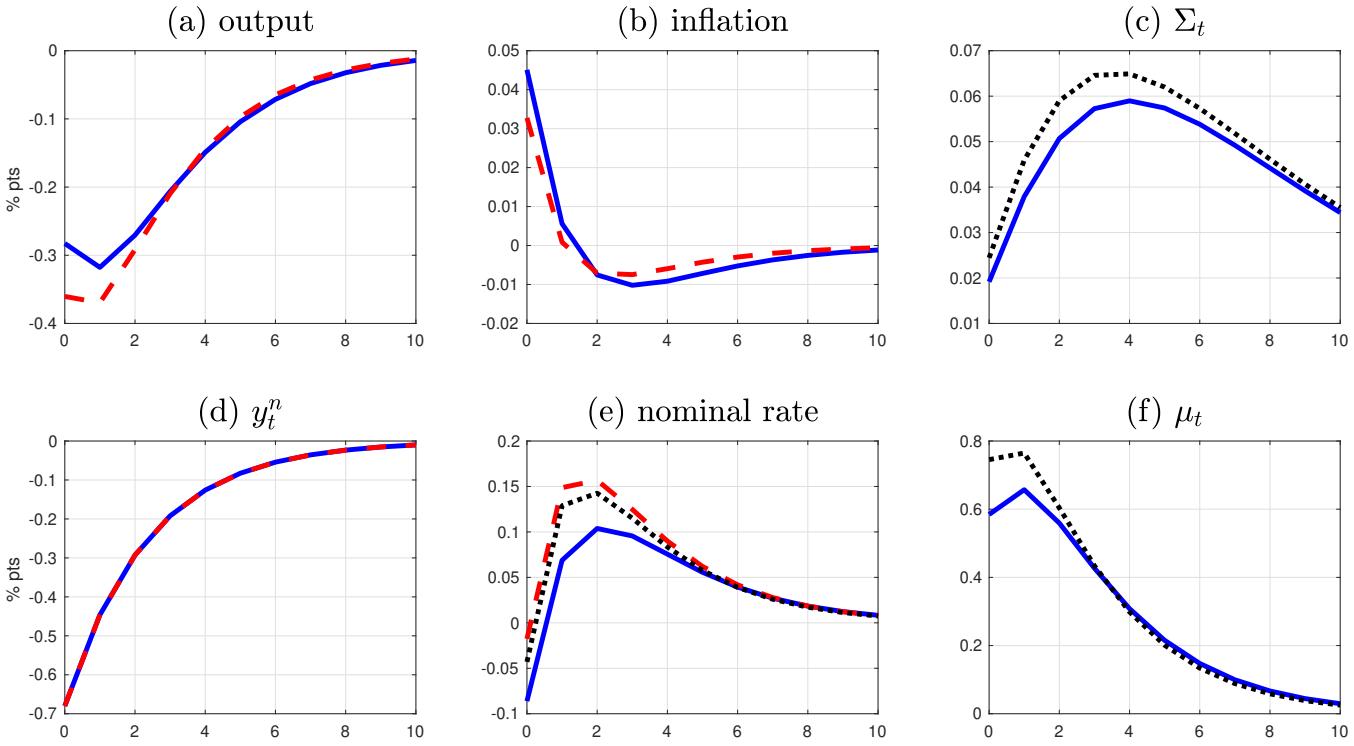

Figure 5: **Optimal policy in response to markup shocks** in HANK (solid blue lines) and RANK (dashed red lines). Black-dotted lines denote outcomes in HANK under non-optimal policy which implements the same $\{\widehat{y}_t, \pi_t\}_{t=0}^{\infty}$ as in RANK. All panels plot log-deviations from steady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

rates (panel (e)), which also translates into a smaller increase in passthrough $\widehat{\mu}_t$ (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 (36) 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 allocations and inflation. One such interest rate rule is

$$i_t = i_t^\star + \phi \pi_t + \phi_{\text{gap}} \left(\Delta y_t - \Delta \widehat{y}_t^e \right) + \phi_y \Delta y_t \tag{38}$$

where $\phi_{\text{gap}} = \phi \frac{\Upsilon(\Omega)}{\varepsilon} \delta(\Omega)$ is the weight on the change in output gap, $\phi_y = \phi \frac{\Upsilon(\Omega)}{\varepsilon} \left(1 - \delta(\Omega)\right)$ is the weight on output growth and $i_t^\star$, 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 (38) 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

$$\mathcal{C}_t^d = y_t + \left(\frac{1-\eta^d}{\eta^d}\right)\mu_t \mathcal{V}_t, \qquad \mathcal{C}_t^{nd} = y_t - \mu_t \mathcal{V}_t, \qquad \text{and} \qquad \mathcal{V}_t = (D_t - D) + \frac{\vartheta}{R_t}\mathcal{V}_{t+1}. \tag{39}$$

$\mathcal{V}_t$ denotes the present discounted value of dividends relative to their steady-state value D. Appendix G shows that linearizing (39) yields the valuation equation

$$\widetilde{\mathcal{V}}_t = \mathcal{D}_y \widehat{y}_t + \mathcal{D}_z \widehat{z}_t + \widetilde{\beta}\widetilde{\mathcal{V}}_{t+1} \qquad \text{where} \qquad \mathcal{D}_y = \frac{1}{\varepsilon} - \frac{\varepsilon - 1}{\varepsilon}\left(\frac{1+\gamma\rho}{\rho/y}\right), \mathcal{D}_z = \frac{1+\rho/y}{(\rho/y)}\frac{\varepsilon-1}{\varepsilon}, \tag{40}$$

where $\widetilde{\mathcal{V}}_t = \widehat{\mathcal{V}}_t/y$ 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):

$$\ln \Sigma_t \quad = \quad \frac{\gamma^2}{2}\mu_t^2 w_t^2 \sigma_t^2 + \ln\left[(1-\vartheta)\,\mathbb{B}_t + \vartheta \Sigma_{t-1}\right] \tag{41}$$

where $\mathbb{B}_t = \eta e^{-\gamma\left(\mathcal{C}_t^d - y_t\right)} + (1-\eta)\, e^{-\gamma\left(\mathcal{C}_t^{nd} - y_t\right)}$ 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 (42) in the Proposition below.

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

$$\min_{\{\widehat{y}_t, \pi_t, \widetilde{\mathcal{V}}_t\}_{t=0}^{\infty}} \frac{1}{2}\sum_{t=0}^{\infty}\beta^t\left\{\Upsilon\left(\Omega\right)\left(\widehat{y}_0 - \delta\left(\Omega\right)\widehat{y}_t^e\right)^2 + \frac{\varepsilon}{\kappa}\pi_0^2\right\} + \frac{\mathbb{K}\left(\eta^d\right)}{2}\left\{\widetilde{\mathcal{V}}_0^2 + \left(1-\beta^{-1}\widetilde{\beta}\right)\sum_{t=1}^{\infty}\beta^t\widetilde{\mathcal{V}}_t^2\right\} \tag{42}$$

subject to the Phillips curve (30) and the valuation equation (40). $\mathbb{K}(\eta^d) \geq 0$ 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 $t = 0$:

$$\Upsilon\left(\Omega\right)x_0 + \varepsilon\widehat{p}_0 + \mathbb{K}(\eta^d)\mathcal{D}_y\widehat{\mathcal{V}}_0 \quad = \quad 0 \tag{43}$$

and for $t > 0$:

$$\Upsilon\left(\Omega\right)\left(x_t - \beta^{-1}\widetilde{\beta}x_{t-1}\right) + \varepsilon\left(\widehat{p}_t - \beta^{-1}\widetilde{\beta}\widehat{p}_{t-1}\right) + \mathbb{K}(\eta^d)\mathcal{D}_y\left(1-\beta^{-1}\widetilde{\beta}\right)\widehat{\mathcal{V}}_t = 0 \tag{44}$$

where $x_t = \widehat{y}_t - \delta(\Omega)\widehat{y}_t^e$ 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 (42) vanishes and the loss function and target criteria are identical to (33) and (36) 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 $\widehat{\mathcal{V}}_t$, in addition to output, the output gap and the price level. This is because fluctuations in $\widehat{\mathcal{V}}_t$ generate between-group consumption inequality: higher $\widehat{\mathcal{V}}_t$ widens the average consumption gap between stockholders and non-stockholders. Stabilizing $\widehat{\mathcal{V}}_t$ 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 $\widehat{\mathcal{V}}_t$. Higher $\widehat{\mathcal{V}}_t$ 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 $\widehat{\mathcal{V}}_0$ 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 $\widehat{\mathcal{C}}_t^d$ above that of non-stockholders $\widehat{\mathcal{C}}_t^{nd}$ (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 0.1% pts. in response to a positive markup shock, in line with the numerical results of BEGS, whose HANK planner raises output by 0-0.1% pts.[23] Similarly, both

[23]Unequally 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.

HANK planners allow inflation to increase on impact (by around 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.

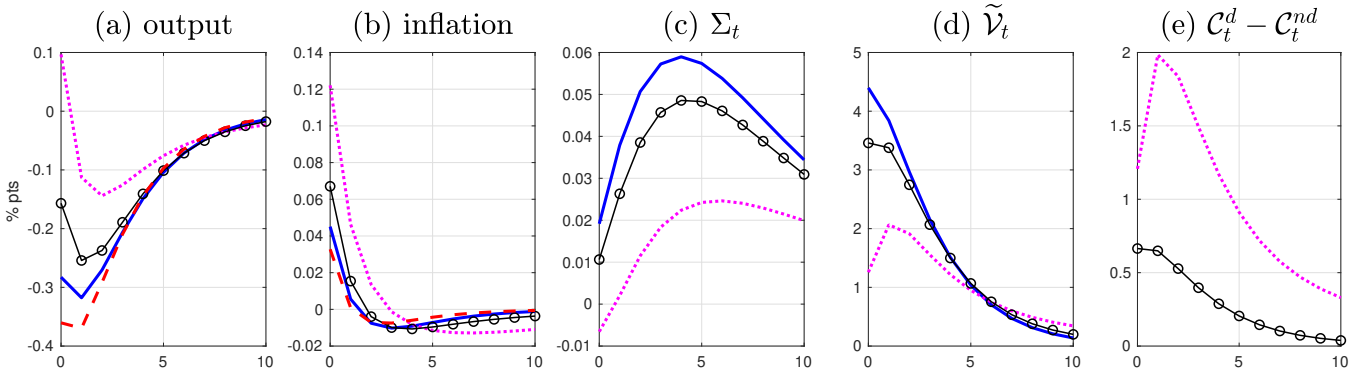

Figure 6: **Optimal policy in response to markup shocks** The red-dashed lines depict dynamics in RANK. All other curves depict dynamics in HANK in response to an increase in firms' desired markups. The blue curves depict the case with equal distribution of profits, the black lines with circle markers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

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 $e^{\gamma \alpha a_0^s(i)}$ to households with wealth $a_0^s(i)$ at date 0.[24] $\alpha \geq 0$ 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

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

[24] Since 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 11, this weight is β^s.

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

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

Unlike in the baseline, Σ_t is not unambiguously increasing in consumption inequality: the planner does not regard all consumption inequality arising from differences in pre-existing wealth inequality at date 0 as undesirable. However, they still regard all inequality resulting from idiosyncratic shocks from date 0 onwards as undesirable. This is reflected in the fact that while (27) is unchanged, (28) is now given by

$$\ln\Sigma_0 = \frac{1}{2}\gamma^2\mu_0^2 w_0^2\sigma_0^2 + \underbrace{\ln\left[\frac{1-\vartheta}{1-\vartheta e^{\frac{\Lambda}{2}\left(\frac{\alpha-\left(1-\tau_0^a\right)\mu_0}{\mu}\right)^2}}\right]}_{\text{welfare cost of date 0 wealth inequality}} \tag{45}$$

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 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 $\widehat{y}_t, \pi_t$ and $\widehat{\Sigma}_t$ 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 $1-\tau_0^{a\star} = \alpha/\mu$. Absent shocks, this tax sets the second term on the RHS of (45) 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.[25] 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*

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

subject to the Phillips curve (30)*.* $\Upsilon_0(\Omega)\geq\Upsilon(\Omega)$ *is increasing in* α *and* $\Upsilon_0(\Omega)=\Upsilon(\Omega)$ *when* $\alpha=0$*. The solution to this problem is characterized by the following target criterion for* $t=0$

$$\Big(1-\delta_0(\Omega)\Big)\widehat{y}_0+\delta_0(\Omega)\Big(\widehat{y}_0-\widehat{y}_0^e\Big)+\frac{\varepsilon}{\Upsilon_0(\Omega)}\widehat{p}_0=0, \tag{46}$$

while for $t>0$*, the target criterion is the same as that for the utilitarian planner,* (36) *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.

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 ($\tau_0^{a\star}=100\%$) and dotted-magenta lines depict the planner with $\alpha=\mu$ (who optimally sets $\tau_0^{a\star}=0\%$). Recall that the utilitarian planner

[25] Since 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 $(\mu_0-\mu)^2$.

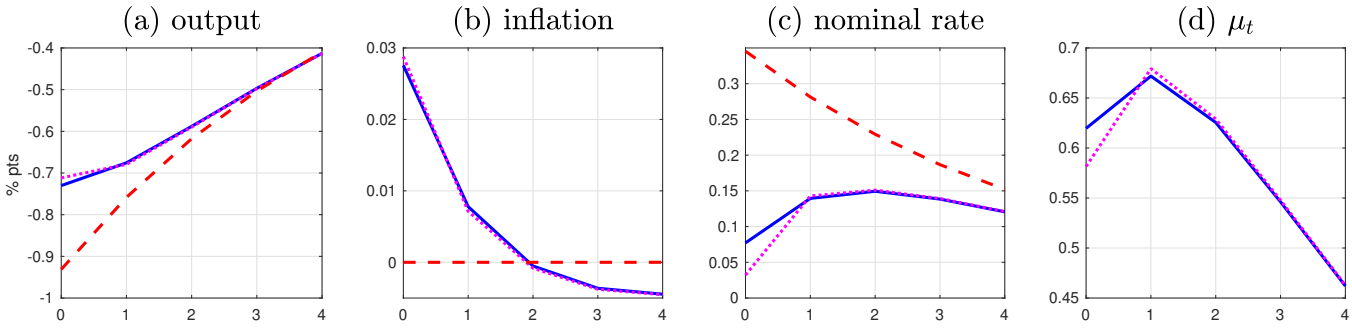

Figure 7: **Optimal policy in response to a negative productivity shock with the non-utilitarian planner** Blue lines depict the utilitarian baseline, dotted-magenta lines depict the planner with $\alpha = \mu$ and the red-dashed curves depict RANK. All panels plot log-deviations from steady state $\times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

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.[26] 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 introduces MPC heterogeneity by incorporating a fraction of hand-to-mouth 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 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 $\widehat{y}_t - \widehat{y}_t^n = \pi_t = 0$. 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.

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

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 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 cyclicality of income risk would be different. In our model with CARA utility, the relevant measure is the cyclicality of *level* income risk. In the CARA-Normal framework, it is the cyclicality of level 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 cyclicality of *log* income that affects differences between the positive properties of HANK and RANK models.[27] Thus, when using our results to assess what optimal policy would be in a quantitative HANK model with CRRA utility, one should check the cyclicality of log income risk.

Finally, a practical implication of our analysis is that monetary policymakers who are concerned with inequality may not need to explicitly incorporate some measure of inequality in their reaction function. Introducing the level of output in the target criterion – and accordingly reducing the relative weights on the output gap and prices – captures the planner's concern for consumption inequality, at least partly.

[27] In Bilbiie (2021), for example, the cyclicality of log 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.

References

Acharya, Sushant, and Keshav Dogra. 2020. "Understanding HANK: Insights From a PRANK." *Econometrica*, 88(3): 1113–1158.

Angeletos, George-Marios, and Laurent Emmanuel Calvet. 2006. "Idiosyncratic Production Risk, Growth and the Business Cycle." *Journal of Monetary Economics*, 53: 1095 – 1115.

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

Auclert, Adrien, Matt Rognlie, and Ludwig Straub. 2018. "The Intertemporal Keynesian Cross." Working Paper 25020.

Bayer, Christian, Benjamin Born, and Ralph Luetticke. 2020. "Shocks, Frictions, and Inequality in US Business Cycles." CEPR Discussion Papers 14364.

Benigno, Pierpaolo, and Michael Woodford. 2005. "Inflation Stabilization and Welfare: The Case of a Distorted Steady State." *Journal of the European Economic Association*, 3(6): 1185–1236.

Bhandari, Anmol, David Evans, Mikhail Golosov, and Thomas J Sargent. 2021. "Inequality, Business Cycles and Monetary-Fiscal- Policy." *Econometrica*, 89(6): 2559–2599.

Bilbiie, Florin O. 2008. "Limited Asset Markets Participation, Monetary Policy and (Inverted) Aggregate Demand Logic." *Journal of Economic Theory*, 140(1): 162–196.

Bilbiie, Florin O. 2021. "Monetary Policy and Heterogeneity: An Analytical Framework." Working Paper,.

Blanchard, Olivier Jean, and Jordi Galí. 2007. "Real Wage Rigidities and the New Keynesian Model." *Journal of Money, Credit and Banking*, 39: 35 – 65.

Caballero, Ricardo J. 1990. "Consumption Puzzles and Precautionary Savings." *Journal of Monetary Economics*, 25(1): 113 – 136.

Calvet, Laurent Emmanuel. 2001. "Incomplete Markets and Volatility." *Journal of Economic Theory*, 98: 295 – 338.

Calvo, Guillermo A, and Maurice Obstfeld. 1988. "Optimal Time-Consistent Fiscal Policy with Finite Lifetimes." *Econometrica*, 56(2): 411–432.

Challe, Edouard. 2020. "Uninsured Unemployment Risk and Optimal Monetary Policy in a Zero-Liquidity Economy." *American Economic Journal: Macroeconomics*, 12(2): 241–83.

Coibion, Olivier, Yuriy Gorodnichenko, Lorenz Kueng, and John Silvia. 2017. "Innocent Bystanders? Monetary policy and inequality." *Journal of Monetary Economics*, 88: 70–89.

Davila, Eduardo, and Andreas Schaab. 2022. "Optimal Monetary Policy with Heterogeneous Agents: A Timeless Ramsey Approach."

Debortoli, Davide, and Jordi Galí. 2018. "Monetary Policy with Heterogeneous Agents: Insights from TANK models." Working Paper.

Farhi, Emmanuel, and Iván Werning. 2019. "Monetary policy, bounded rationality, and incomplete markets." *American Economic Review*, 109(11): 3887–3928.

Galí, Jordi. 2015. *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications, Second Edition.* Princeton University Press.

Galí, Jordi, David Lopez-Salido, and Javier Valles. 2007. "Understanding the Effects of Government Spending on Consumption." *Journal of the European Economic Association*, 5(1): 227–270.

Guvenen, Fatih, Serdar Ozkan, and Jae Song. 2014. "The Nature of Countercyclical Income Risk." *Journal of Political Economy*, 122(3): 621–660.

Hansen, Niels-Jakob Harbo, Alessandro Lin, and Rui Mano. 2020. "Should Inequality Factor into Central Banks' Decisions?" IMF Working Papers.

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

Kekre, Rohan. 2022. "Unemployment Insurance in Macroeconomic Stabilization." Working Paper.

Le Grand, François, Alaïs Martin-Baillon, and Xavier Ragot. 2021. "Should monetary policy care about redistribution? Optimal fiscal and monetary policy with heterogeneous agents." Working Paper.

McKay, Alisdair, and Christian K Wolf. 2022. "Optimal Policy Rules in HANK." Working Paper.

McKay, Alisdair, and Ricardo Reis. 2016. "The Role of Automatic Stabilizers in the U.S. Business Cycle." *Econometrica*, 84: 141–194.

McKay, Alisdair, and Ricardo Reis. 2021. "Optimal Automatic Stabilizers." *The Review of Economic Studies.*

McKay, Alisdair, Emi Nakamura, and Jon Steinsson. 2016. "The Power of Forward Guidance Revisited." *American Economic Review*, 106(10): 3133–3158.

Mumtaz, Haroon, and Angeliki Theophilopoulou. 2017. "The impact of monetary policy on inequality in the UK. An empirical analysis." *European Economic Review*, 98(C): 410–423.

Nakajima, Makoto, and Vladimir Smirnyagin. 2019. "Cyclical Labor Income Risk." Working Paper.

Nisticò, Salvatore. 2016. "Optimal Monetary Policy And Financial Stability In A Non-Ricardian Economy." *Journal of the European Economic Association*, 14(5): 1225–1252.

Nuño, Galo, and Carlos Thomas. 2022. "Optimal Redistributive Inflation." *Annals of Economics and Statistics*, 146: 3–64.

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

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

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

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

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

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