A Behavioral Heterogeneous Agent New Keynesian Model

Oliver Pfäuti* Fabian Seyrich†

November 1, 2022
Link to most recent version

Abstract

Recent empirical evidence has led to a rethinking of how monetary policy is transmitted to the economy: (i) monetary policy influences household consumption to a large extent through indirect general equilibrium effects; (ii) households with higher marginal propensities to consume are on average more exposed to monetary policy; (iii) announcements of future monetary policy changes have relatively weak effects on current economic activity; and (iv) the effective lower bound on nominal interest rates does not lead to large instabilities. We develop a New Keynesian model with household heterogeneity and bounded rationality in the form of cognitive discounting that accounts for all these facts simultaneously. In contrast to existing models, we jointly account for these facts without having to rely on a specific monetary or fiscal policy. We use our model to revisit the policy implications of inflationary supply shocks and find a pronounced trade-off between price stability on the one hand, and fiscal and distributional consequences on the other hand. Monetary policy has to increase interest rates more strongly than under rational expectations to fully stabilize inflation. This leads to a stronger increase in consumption inequality and has larger fiscal implications. When monetary policy follows a standard Taylor rule, inequality declines after the negative supply shock. Inflation, however, increases substantially due to a novel amplification channel driven by cognitive discounting, household heterogeneity and the interaction of the two.

Keywords: Monetary Policy, Heterogeneous Households, Behavioral Macroeconomics, Forward Guidance, Lower Bound, Inflation, Macroeconomic Stabilization

JEL Codes: E21, E52, E62, E71

*University of Mannheim, pfaeuti.oliver@gmail.com
†Berlin School of Economics, DIW Berlin, and Freie Universität Berlin, fabian.seyrich@gmail.com.
We thank Klaus Adam, George-Marios Angeletos, Neele Balke, Nicholas Barberis, Christian Bayer, Florin Bilbiie, Eduardo Dávila, Stéphane Dupraz (discussant), Xavier Gabaix, Zhen Huo, Diego Känzig, Alexander Kriwoluzky, Max Jager, Timo Reinelt, Ricardo Reis, Hannah Seidl, Alp Simsek, Alejandro Van der Ghote (discussant), Maximilian Weiß (discussant) and seminar and conference participants at Yale University, the University of Chicago, the SED Annual Meeting 2022, the ECB Sintra Forum 2022, the PSE Macro Days 2022, the Bank of Finland, the SNDE Annual Symposium, QuickTalks, the RCEA Conference 2022, the RGS Doctoral Conference, the SMYE 2022, the MMF conference, the University of Warwick, the University of Mannheim, DIW Berlin and HU Berlin for helpful comments and suggestions. Oliver Pfäuti gratefully acknowledges financial support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 (Project C02) and the Stiftung Geld & Währung. Fabian Seyrich gratefully acknowledges financial support by the Leibniz Association through the project "Distributional effects of macroeconomic policies in Europe". First version online: January 2022.

1 Introduction

Recent empirical evidence has led to a rethinking of how monetary policy is transmitted to the economy: (i) monetary policy affects household consumption to a large extent through changing people's incomes rather than directly through changes in the real interest rate. These *indirect effects* tend to amplify the effects of conventional monetary policy on consumption as (ii) the incomes of households that exhibit higher marginal propensities to consume are found to be more exposed to aggregate income fluctuations induced by monetary policy; (iii) announcements of future monetary policy changes, in contrast, have relatively weak effects on current economic activity; and (iv) advanced economies have not experienced large instabilities in times the nominal interest rate has effectively been pegged to the lower bound.[1]

In this paper, we propose a new framework that accounts for these four facts *simultaneously*: the behavioral Heterogeneous Agent New Keynesian model—or *behavioral HANK model*, for short. The model features a standard New Keynesian core, but we allow for household heterogeneity and bounded rationality in the form of cognitive discounting. The presence of both—household heterogeneity and bounded rationality—is key to account for the four facts jointly. In contrast to existing models, our model accounts for the four facts without having to rely on a specific monetary or fiscal policy.

We first focus on a limited-heterogeneity setup which enables us to derive all results in closed form. The model is kept deliberately stylized and illustrates how cognitive discounting interacts with household heterogeneity. Households that exhibit higher marginal propensities to consume are more exposed to monetary policy which is crucial to account for the fact that monetary policy is amplified through indirect general equilibrium effects. Under cognitive discounting, households' expectations underreact to aggregate news which dampens the effects of announced future monetary policy changes and ensures that the model remains determinate under an interest-rate peg. We then replace the limited-heterogeneity setup with a standard quantitative incomplete markets setup with full heterogeneity and show that our results carry over to this more quantitative model.

Our model suggests that accounting for these four facts simultaneously has important implications for policy. We find a strong trade-off for monetary policy between price stability on the one side and fiscal and distributional consequences on the other side after an inflationary supply shock. To fully stabilize inflation, monetary policy needs to increase interest rates by more than under rational expectations, which pushes up the government debt level and inequality more strongly. In contrast, when monetary policy follows a standard Taylor rule it mitigates the fiscal consequences and decreases inequality. But now inflation increases more strongly than in the rational

[1] See, e.g., Ampudia et al. (2018), Slacalek et al. (2020) and Holm et al. (2021) for the empirical relevance of indirect channels in the transmission of monetary policy, Auclert (2019), Patterson (2019) and Slacalek et al. (2020) for evidence on households' income exposure and their marginal propensities to consume, and see, for example, Del Negro et al. (2015), D'Acunto et al. (2020), and Roth et al. (2021) for empirical evidence on the (in-)effectiveness of monetary policy announcements about its future actions, and Debortoli et al. (2020) and Cochrane (2018) on the stability at the lower bound.

model through a novel amplification channel induced by cognitive discounting, the underlying heterogeneity and the interaction of the two.

To arrive at the tractable version of our model, we consider a setup with two groups of households. One group is "unconstrained", in the sense that they participate in financial markets and are on their Euler equation. The other group consists of "hand-to-mouth" households who consume all their disposable income. They exhibit high marginal propensities to consume (MPCs) and are more exposed to monetary policy in line with the data. To introduce a precautionary-savings motive, households face an idiosyncratic risk of switching type.

Households anchor their expectations about future macroeconomic variables to the steady state and cognitively discount expected future deviations as in Gabaix (2020). As a result, expectations then underreact to aggregate news, as we show to be the case empirically across all income groups and consistent with findings in D'Acunto et al. (2020) or Roth et al. (2021).[2]

Like the textbook Representative Agent New Keynesian (RANK) model, our tractable model can be represented in just three equations. The key novelty arising from bounded rationality and household heterogeneity is a new aggregate IS equation. In contrast to the textbook model, our IS equation features a lower sensitivity of current output to changes in expected future output due to households' cognitive discounting and a stronger sensitivity of current output to changes in the real interest rate as households with higher MPCs are more exposed to monetary policy.

As a result of the lower sensitivity of current output to future expected output, announced policies that increase future output, such as announced future interest rate cuts, are less effective in stimulating current output. After such an announced future interest rate cut, unconstrained households want to consume more already today as they want to smooth their consumption intertemporally. Additionally, their precautionary savings motive decreases as they would be better off in case they become hand-to-mouth in the future because hand-to-mouth households benefit more from the future boom. Cognitive discounting weakens *both* of these channels and thus, explains the lower sensitivity of current output to future expected output. The farther away in the future the announced interest rate cut takes place, the smaller its effect on today's output. Hence, the model does not suffer from the *forward guidance puzzle,* which describes the paradoxical finding in many models that announced future interest-rate changes are at least as effective in stimulating current output than contemporaneous interest-rate changes (Del Negro et al. (2015), McKay et al. (2016)). In addition, our model remains determinate under an interest-rate peg and remains stable at the effective lower bound (ELB).

The second deviation from the textbook IS equation—the stronger sensitivity of current output

[2] We show how to microfound cognitive discounting as a noisy-signal extraction problem of otherwise rational agents. Angeletos and Lian (2022) show how other forms of bounded rationality or lack of common knowledge can be observationally equivalent. For further evidence on underreaction of expectations or general patterns of inattention, see, e.g., Coibion and Gorodnichenko (2015), Coibion et al. (2020) or Angeletos et al. (2021). Born et al. (2022) show that even when agents overreact to micro news, they underreact to macro news, which is what we focus on in this paper.

to changes in the real interest rate—arises because households with higher MPCs are more exposed to monetary policy. An expansionary monetary policy shock increases the income of the hand-to-mouth households more than one-for-one. As these households consume all their disposable income, this leads to a stronger response of aggregate consumption than if all households would be exposed equally to monetary policy. Thus, the model features amplification of conventional monetary policy shocks due to indirect general equilibrium effects. A decomposition into direct and indirect effects shows that indeed the major share of the monetary policy transmission works through indirect effects.

That our model simultaneously generates amplification of conventional monetary policy through indirect effects and rules out the forward-guidance puzzle is in stark contrast to rational models. Rational HANK models that generate amplification through indirect effects exacerbate the forward-guidance puzzle. Rational models that resolve the forward-guidance puzzle, on the other hand, cannot simultaneously generate amplification of monetary policy through indirect effects (see Werning (2015), Acharya and Dogra (2020), and Bilbiie (2021)).

We extend our tractable framework along several dimensions to show the model's compatibility with additional empirical patterns. We first analytically derive the intertemporal MPCs (iMPCs) and show they match empirical estimates—a key statistic in HANK models (Auclert et al. (2018), Kaplan and Violante (2020)). Second, we consider sticky wages and show how the model generates hump-shaped responses of macroeconomic variables in response to aggregate shocks, and expectations that initially underreact followed by a delayed overshooting (both consistent with the data, Auclert et al. (2020), Angeletos et al. (2021) and Adam et al. (2022)). Both findings are not present when abstracting from heterogeneity, bounded rationality or both.

We then replace the limited-heterogeneity assumptions and build on a standard incomplete-markets setup. We relax many of the strong assumptions we imposed in the tractable framework. In particular, we now consider ex-ante identical households that face uninsurable idiosyncratic productivity risk, incomplete markets and borrowing constraints that are endogenously binding. We further allow for heterogeneity in the degree of cognitive discounting.[3] We show numerically that the full model also accounts for facts (i)-(iv) simultaneously.

We use the model to revisit the monetary and fiscal policy implications of inflationary supply shocks. Many advanced economies have recently experienced a dramatic surge in inflation which is partly attributed to disruptions in production (see di Giovanni et al. (2022)). We analyze these supply disruptions by considering a negative productivity shock.

We first consider a monetary policy that fully stabilizes inflation and find that this policy closes the output gap independent of the presence of cognitive discounting. The required interest-rate response to fully stabilize inflation, however, needs to be much stronger when accounting for cognitive discounting. The reason is that households expect interest rates to remain elevated

[3] As in the tractable model, households cognitively discount expected deviations from the stationary equilibrium after an aggregate shock but are fully rational with respect to their idiosyncratic risk.

for some time due to the persistence of the shock and the higher expected interest rates help to stabilize current inflation. This expectation channel is dampened under cognitive discounting. Hence, the monetary authority needs to increase interest rates more forcefully.

These stronger interest-rate hikes create side effects. In particular, they have strong fiscal implications as they increase the cost of government debt, which leads to a larger increase in government debt. Furthermore, consumption inequality increases more strongly as wealthy households benefit more from higher interest rates than asset-poor householdss.

We then assume that monetary policy follows a standard Taylor rule and show that in this case, inflation and the output gap increase considerably due to a novel amplification channel. Both—the underlying heterogeneity and bounded rationality—amplify the inflationary pressure from the supply shock and the two mutually reinforce each other: the positive output gap especially benefits households with higher MPCs increasing the output gap further and, thus, calls for higher interest rates in each period. As households cognitively discount these higher (future) interest rates, this further increases the output gap amplifying the redistribution to high MPCs households and therefore the increase in the output gap until the economy ends up in an equilibrium with a higher output gap and higher inflation. Yet, consumption inequality now decreases as poorer households tend to benefit more from the higher output gap.

We also consider cost-push shocks as an alternative explanation for high inflationary pressure and find similar implications for monetary and fiscal policy. In sum, our model predicts that it is more difficult for monetary policy to stabilize the economy after an inflationary supply side shock and that monetary policy faces a severe trade-off between aggregate stabilization and price stability on the one hand, and fiscal and distributional consequences on the other hand.

Related literature. The literature treats the facts (i)-(iv) mostly independent from each other. The heterogeneous-household literature has highlighted the transmission of monetary policy through indirect, general equilibrium effects (Kaplan et al. (2018), Auclert (2019), Auclert et al. (2020), Bilbiie (2020), Luetticke (2021)), and proposed potential resolutions of the forward guidance puzzle (McKay et al. (2016, 2017), Hagedorn et al. (2019), McKay and Wieland (2022)). Werning (2015) and Bilbiie (2021) combine the themes of policy amplification and forward guidance puzzle in HANK and establish a trade-off inherent in models with household heterogeneity: if HANK models amplify contemporaneous monetary (and fiscal) policy through redistribution towards high MPC households, they dampen precautionary savings desires after a forward guidance shock which aggravates the forward guidance puzzle.

Few resolutions of this trade-off—what Bilbiie (2021) calls the *Catch-22*—have been put forward. In contrast to our model, they all rely on a specific design for either monetary or fiscal policy. Bilbiie (2021) shows that if monetary policy follows a Wicksellian price level targeting rule or fiscal policy follows a nominal bond rule, his tractable HANK model can simultaneously account for facts (i)-(iv).[4] Hagedorn et al. (2019) shows how introducing nominal government bonds and coupling it with a particular nominal bond supply rule can resolve the forward guidance puzzle in a quantitative HANK model (following the theoretical arguments in Hagedorn et al. (2016) and Hagedorn (2018)). In contrast, we account for the four facts even in the case in which monetary policy follows a standard Taylor rule and absent any nominal bonds or specific fiscal rules.

Farhi and Werning (2019) also combine household heterogeneity with some form of bounded rationality, but focus entirely on resolving the forward-guidance puzzle. Our model generates a number of additional desirable features, such as amplification of monetary policy through indirect effects in a setting with unequal exposure of households to monetary policy. We also consider a different form of bounded rationality, cognitive discounting, while Farhi and Werning (2019) focus on level-k thinking. Our setup is consistent with the empirical findings in Roth et al. (2021) who show that households adjust their interest-rate expectations only by about half of what the Fed announces, even when being told the Fed's intended interest-rate path.[5] Furthermore, we are the first ones to consider supply shocks and show that the interaction of household heterogeneity and bounded rationality has very different implications for supply shocks than for forward guidance shocks.

Few other papers share the combination of nominal rigidities, household heterogeneity and some deviation from full information rational expectations (FIRE). Laibson et al. (2021) introduces *present bias* in a model of household heterogeneity but the model is set in partial equilibrium and they do not consider how the power of forward guidance or the stability at the lower bound are affected by the presence of the two frictions. Auclert et al. (2020) incorporate sticky information into a HANK model to generate hump-shaped responses of macroeconomic variables to aggregate shocks while simultaneously matching iMPCs. We obtain similar results in our extension of the tractable model with sticky wages. Their paper, however, does not discuss the implications of the deviation from FIRE and heterogeneity for forward guidance or stability at the lower bound.[6]

Outline. The rest of the paper is structured as follows. We present our tractable behavioral HANK model in Section 2 and our analytical results in Section 3. In Section 4, we develop the quantitative behavioral HANK model, show how it can account for the facts simultaneously and discuss the role of heterogeneity in the behavioral bias. We use the quantitative model to study the policy implications of an inflationary supply-side shock in Section 5. We discuss three extensions of the tractable model in Section 6 and Section 7 concludes.

[4] Bilbiie (2021) proposes an additional resolution: a pure risk channel which can, in theory, break the co-movement of income risk and inequality. However, it requires a calibration which is at odds with the data.

[5] In an extension, we consider the case in which some households (financial markets, for example) fully incorporate the announced interest-rate paths into their expectations (see Section 4.2 where we discuss heterogeneous degrees of cognitive discounting) and show that our results remain robust in that scenario.

[6] Wiederholt (2015), Angeletos and Lian (2018), Andrade et al. (2019), Gabaix (2020) consider deviations from FIRE and Michaillat and Saez (2021) introduce wealth in the utility function (all in non-HANK setups) and show how to resolve the forward guidance puzzle.

2 A Tractable Behavioral HANK Model

In this section, we present our tractable New Keynesian model featuring household heterogeneity and bounded rationality before we then turn to the full-blown incomplete-markets setup later on. To ensure closed-form solutions, we make a number of assumptions that are typical in the analytical HANK literature (e.g., McKay et al. (2017), Bilbiie (2021)).

2.1 Structure of the Model

Households. Time is discrete and denoted by $t = 0, 1, 2, ...$ The economy is populated by a unit mass of households, indexed by $i \in [0, 1]$. Households obtain utility from (non-durable) consumption, C_t^i, and dis-utility from working N_t^i. Households discount future utility at rate $\beta \in [0, 1]$. We assume a standard CRRA utility function

$$\mathcal{U}(C_t^i, N_t^i) \equiv \begin{cases} \frac{(C_t^i)^{1-\gamma}}{1-\gamma} - \frac{(N_t^i)^{1+\varphi}}{1+\varphi}, & \text{if } \gamma \neq 1, \\ \log(C_t^i) - \frac{(N_t^i)^{1+\varphi}}{1+\varphi}, & \text{if } \gamma = 1, \end{cases} \tag{1}$$

where φ denotes the inverse Frisch elasticity and γ the relative risk aversion.

Households can save in government bonds B_{t+1}^i, paying nominal interest i_t, and face an exogenous borrowing constraint which we set to zero. Households participate in financial markets infrequently. When they do participate, they can freely trade bonds. Otherwise, they simply receive the payoff from their previously acquired bonds. For now, asset-market participation is exogenous and can be interpreted, for example, as a shock to the household's taste or patience. We denote households participating in financial markets by U as, in equilibrium, they will be Unconstrained in the sense that they are on their Euler equation. We denote the non-participants by H as they neither save nor borrow and are thus, $H$and-to-mouth. An unconstrained household remains unconstrained with probability s and becomes hand-to-mouth with probability $1 - s$. Hand-to-mouth households remain hand-to-mouth with probability h and switch to being unconstrained with probability $1 - h$. In what follows, we focus on stationary equilibria where $\lambda \equiv \frac{1-s}{2-s-h}$ denotes the constant share of hand-to-mouth households. Unconstrained households receive a share $1 - \mu^D$ of the intermediate firm profits, D_t, and hand-to-mouth households the remaining share μ^D.

Households belong to a family whose intertemporal welfare is maximized by its utilitarian family head. The head can only provide insurance within types but not across types, i.e., the head pools all the resources within types. Thus, in equilibrium every U household will consume and work the same amount and every H household will consume and work the same amount but the H households' consumption and labor supply is not necessarily the same as that of U households. When households switch from being unconstrained to being hand-to-mouth, they keep their government bonds.

We allow for the possibility that the family head is boundedly rational in the way we describe in the following.[7] The program of the family head is

$$V\left(B_t^U\right) = \max_{\{C_t^U, C_t^H, B_{t+1}^U, N_t^U, N_t^H\}} \left[(1-\lambda)\mathcal{U}\left(C_t^U, N_t^U\right) + \lambda\mathcal{U}\left(C_t^H, N_t^H\right)\right] + \beta\mathbb{E}_t^{BR} V\left(B_{t+1}^U\right)$$

subject to the flow budget constraints of unconstrained households

$$C_t^U + B_{t+1}^U = W_t N_t^U + \frac{1-\mu^D}{1-\lambda} D_t + s\frac{1+i_{t-1}}{1+\pi_t} B_t^U t, \tag{2}$$

and the hand-to-mouth households

$$C_t^H = W_t N_t^H + \frac{\mu^D}{\lambda} D_t + (1-s)\frac{1+i_{t-1}}{1+\pi_t}\frac{1-\lambda}{\lambda} B_t^U, \tag{3}$$

as well as the borrowing constraint $B_{t+1}^U \geq 0$, where W_t is the real wage. The budget constraints reflect our assumptions that only U households can save in government bonds, but that households keep their acquired government bonds when switching their type as well as the assumption of full-insurance within type, as the bonds are equally shared within types.

The optimality conditions are given by the Euler equation of unconstrained households

$$\left(C_t^U\right)^{-\gamma} \geq \beta\mathbb{E}_t^{BR}\left[R_t\left(s\left(C_{t+1}^U\right)^{-\gamma} + (1-s)\left(C_{t+1}^H\right)^{-\gamma}\right)\right], \tag{4}$$

where $R_t \equiv \frac{1+i_t}{1+\pi_{t+1}}$ denotes the real interest rate, and the respective labor-leisure equations of both types are given by:

$$\left(N_t^i\right)^{\varphi} = W_t\left(C_t^i\right)^{-\gamma}.$$

Importantly, the Euler equation of the unconstrained households features a self-insurance motive as unconstrained households demand bonds to self-insure their idiosyncratic risk of becoming hand-to-mouth.

We focus on the zero liquidity equilibrium (i.e., bond supply B_t^G is equal to zero for all t) to keep our model tractable (as in Krusell et al. (2011), McKay et al. (2017), Ravn and Sterk (2017), and Bilbiie (2021)).

Bounded rationality. We follow Gabaix (2020) and model bounded rationality in the form of cognitive discounting.[8] Let X_t be a random variable (or vector of variables) and let us define X_t^d as some default value the agent may have in mind and let $\tilde{X}_{t+1} \equiv X_{t+1} - X_t^d$ denote the deviation from this default value.[9] The behavioral agent's expectation about X_{t+1} is then defined as

$$\mathbb{E}_t^{BR}\left[X_{t+1}\right] = \mathbb{E}_t^{BR}\left[X_t^d + \tilde{X}_{t+1}\right] \equiv X_t^d + \bar{m}\mathbb{E}_t\left[\tilde{X}_{t+1}\right], \tag{5}$$

[7]We show in Appendix A.9 how the family head's expectation can be understood as an average expectation over all households' expectations within the family where each household receives a noisy signal about the future state.

[8]While Gabaix (2020) embeds bounded rationality in a NK model the basic idea of behavioral inattention (or sparsity) has been proposed by Gabaix earlier already (see Gabaix (2014, 2016)) and a handbook treatment of behavioral inattention is given in Gabaix (2019). Benchimol and Bounader (2019) and Bonciani and Oh (2021) study optimal monetary policy in a RANK and TANK model, respectively, with this kind of behavioral frictions.

[9]Gabaix (2020) focuses on the case in which X_t denotes the state of the economy. He shows (Lemma 1 in Gabaix (2020)) that this form of cognitive discounting also applies to all other variables. Appendix A.8 derives our results following the approach in Gabaix (2020). The results remain exactly the same.

where $\mathbb{E}_t[\cdot]$ is the rational expectations operator and $\bar{m} \in [0,1]$ is the behavioral parameter capturing the degree of rationality. A higher $\bar{m}$ denotes a smaller deviation from rational expectations and rational expectations are captured by $\bar{m} = 1$. Intuitively, the behavioral agent anchors her expectations to the default value and cognitively discounts expected future deviations from this default value. We focus on the steady state as the default value but we discuss an alternative assumption in Section 6.3. Note, that absent aggregate shocks the agents are fully rational as they know their idiosyncratic risk.

While we present a way how to microfound cognitive discounting as a noisy-signal extraction problem in Appendix A.9, note, that the exact microfoundation or underlying behavioral friction is not crucial for the rest of our analysis. Angeletos and Lian (2022) show how other forms of bounded rationality or lack of common knowledge can lead to observationally-equivalent expectations.

Log-linearizing equation (5) around the steady state yields

$$\mathbb{E}_t^{BR}[\widehat{x}_{t+1}] = (1-\bar{m})\widehat{x}_t^d + \bar{m}\mathbb{E}_t[\widehat{x}_{t+1}] \tag{6}$$

and when X_t^d is the steady state value, we obtain $\mathbb{E}_t^{BR}[\widehat{x}_{t+1}] = \bar{m}\mathbb{E}_t[\widehat{x}_{t+1}]$. Given $\bar{m} < 1$, expectations underreact to aggregate news about the future, consistent with findings in D'Acunto et al. (2020) or Roth et al. (2021). In Appendix B, we estimate $\bar{m}$ for different household groups based on their income and we find that households of all income groups underreact to macroeconomic news. We also discuss other empirical evidence on $\bar{m}$ and how we can map recent evidence in Coibion and Gorodnichenko (2015) and Angeletos et al. (2021) to $\bar{m}$. As a benchmark, we follow Gabaix (2020) and set $\bar{m}$ to 0.85, which is a rather conservative deviation from rational expectations, given that the empirical evidence points towards a $\bar{m}$ between 0.6 and 0.85.

Firms. We assume a standard New Keynesian firm side with sticky prices. All households consume the same aggregate basket of individual goods, $j \in [0,1]$, $C_t = \left(\int_0^1 C_t(j)^{\frac{\epsilon-1}{\epsilon}} dj\right)^{\frac{\epsilon}{\epsilon-1}}$, where $\epsilon > 1$ is the elasticity of substitution between the individual goods. Each firm faces demand $C_t(j) = \left(\frac{P_t(j)}{P_t}\right)^{-\epsilon} C_t$, where $P_t(j)/P_t$ denotes the individual price relative to the aggregate price index, $P_t^{1-\epsilon} = \int_0^1 P_t(j)^{1-\epsilon} dj$, and produces with the linear technology $Y_t(j) = N_t(j)$. The real marginal cost is given by W_t. We assume that the government pays a constant subsidy τ^S on revenues to induce marginal cost pricing in the steady state. The subsidy is financed by a lump-sum tax on firms T_t^F. Hence, the profit function is $D_t(j) = (1+\tau^S)[P_t(j)/P_t]Y_t(j) - W_t N_t(j) - T_t^F$. Total profits are then $D_t = Y_t - W_t N_t$ and are zero in steady state. Zero steady-state profits imply full insurance in steady state, as households only differ in their profit income, i.e., in the steady state we have $C^H = C^U = C$. In the log-linear dynamics around this steady state, profits vary inversely with the real wage, $\widehat{d}_t = -\widehat{w}_t$, where variables with a hat on top denote log-deviations from steady state. We allow for steady state inequality in Appendix D and show that our results are not driven by this assumption and are in fact barely affected even by substantial inequality in the steady state.

Government. Fiscal policy induces the optimal steady state subsidy financed by lump-sum taxation of firms. We introduce taxes and government bonds in our quantitative model in Section 4.

In most of the analysis, we assume that monetary policy follows a standard (log-linearized) Taylor rule

$$\widehat{i}_t = \phi \pi_t + \epsilon_t^{MP}, \tag{7}$$

with ϵ_t^{MP} being a monetary policy shock (Appendix A discusses more general Taylor rules). For now, monetary policy shocks are the only source of aggregate uncertainty.

Market clearing. Market clearing requires that the goods market clears $Y_t = C_t = \lambda C_t^H + (1-\lambda)C_t^U$ and the labor market clears $N_t = \lambda N_t^H + (1-\lambda)N_t^U$. Bond market clearing implies $B_{t+1}^U = 0$ at all t.

2.2 Log-Linearized Dynamics

We now focus on the log-linearized dynamics around the full-insurance, zero-liquidity steady state. First, we can write consumption of the hand-to-mouth households as

$$\widehat{c}_t^H = \chi \widehat{y}_t, \tag{8}$$

with

$$\chi \equiv 1 + \varphi\left(1 - \frac{\mu^D}{\lambda}\right) \tag{9}$$

measuring the cyclicality of the H household's consumption (see appendix A.1). Patterson (2019) documents that households with higher MPCs tend to be more exposed to aggregate income fluctuations induced by monetary policy or other demand shocks—fact (ii) in the introduction. We can account for fact (ii) by setting $\chi > 1$. Similarly, Auclert (2019) finds that poorer households tend to exhibit higher MPCs. Together with the finding in Coibion et al. (2017) that poorer households' income is on average more exposed to monetary policy, this also implies $\chi > 1$. For given φ, this requires $\mu^D < \lambda$.

Why does $\mu^D < \lambda$ imply that the consumption of hand-to-mouth households moves more than one-to-one with aggregate output after a monetary policy shock? Consider an expansionary monetary policy shock, i.e., an unexpected decrease in the nominal interest rate. Unconstrained households want to consume more and save less, leading to an increase in demand. Firms then increase their labor demand, leading to an increase in wages. Due to the assumption of sticky prices and flexible wages, profits in the New Keynesian model decrease. In the representative agent model, the representative agent both incurs the increase in wages and the decrease in profits coming from firms. With household heterogeneity, this is not necessarily the case. If the hand-to-mouth households receive less of the profits than their share in the population ($\mu^D < \lambda$) the increase in the real wage is fully transmitted to their income whereas the decrease in profits is not. Thus, the income of H households increases more than aggregate output increases. The unconstrained households whose profit share is disproportionally large, on the other hand, work more to make up for the income loss due to lower profit income. It is thus mainly the unconstrained households who produce the additional output.

Combining equation (8) with the goods market clearing condition yields

$$\widehat{c}_t^U = \frac{1-\lambda\chi}{1-\lambda}\widehat{y}_t, \tag{10}$$

which implies that consumption inequality is given by:[10]

$$\widehat{c}_t^U - \widehat{c}_t^H = \frac{1-\chi}{1-\lambda}\widehat{y}_t. \tag{11}$$

Thus, if $\chi > 1$, inequality is countercyclical as it varies negatively with total output, i.e., inequality increases in recessions and decreases in booms. In line with the empirical evidence on the covariance between MPCs and income exposure the data also points towards $\chi > 1$ when looking at the cyclicality of inequality, conditional on monetary policy: Coibion et al. (2017), Mumtaz and Theophilopoulou (2017), Ampudia et al. (2018) and Samarina and Nguyen (2019) all provide evidence of countercyclical inequality conditional on monetary policy shocks.

The second key equilibrium equation is the log-linearized bond Euler equation of U households:

$$\widehat{c}_t^U = s\mathbb{E}_t^{BR}\left[\widehat{c}_{t+1}^U\right] + (1-s)\mathbb{E}_t^{BR}\left[\widehat{c}_{t+1}^H\right] - \frac{1}{\gamma}\left(\widehat{i}_t - \mathbb{E}_t^{BR}\pi_{t+1}\right). \tag{12}$$

For the case without type-switching, i.e., for $s = 1$, equation (12) boils down to a standard Euler equation. For $s \in [0, 1)$, however, the agent takes into account that she might switch her type and self-insures against becoming hand-to-mouth next period. How strongly this precautionary-saving motive affects the household's consumption away from the stationary equilibrium will depend on the household's degree of bounded rationality. We will, following the assumption in Gabaix (2020), often focus on the case in which households are rational with respect to the real rate, i.e., we replace $\mathbb{E}_t^{BR}\pi_{t+1}$ with $\mathbb{E}_t\pi_{t+1}$ in equation (12). We show in Appendix D that our results go through with boundedly-rational real-rate expectations.

Supply side. For simplicity and to get a clear understanding of the mechanisms driving our results, we focus on a static Phillips curve for now:

$$\pi_t = \kappa\widehat{y}_t, \tag{13}$$

where $\kappa \geq 0$ captures the slope of the Phillips curve. Such a static Phillips curve arises if we assume that firms are either completely myopic or if they face Rotemberg-style price adjustment costs relative to yesterday's market average price index, instead of their own price (see Bilbiie (2021)). In Appendix D we show that a forward-looking Phillips Curve (rational or behavioral) does not qualitatively affect our results.

[10] We denote the case in which unconstrained households consume relatively more than hand-to-mouth households as higher inequality, even though they consume the same amount in steady state. As we move away from the tractable model in Sections 4 and 5, households' consumption levels will differ in the stationary equilibrium.

Discussion of assumptions. Throughout this section, we have imposed several assumptions that allow us in the following section to analytically characterize our results as well as to generate analytical insights into how household heterogeneity and bounded rationality interact. In particular, we assume full insurance within types, exogenous type switching, a zero-liquidity equilibrium, no inequality in the steady state, a static Phillips Curve and homogeneous degrees of bounded rationality. We relax all these assumptions in Section 4 and show that our results presented in the following do not depend on these assumptions.

3 Results

In this section, we derive the three-equation representation of the tractable behavioral HANK model and show that the model is consistent with facts (i)-(iv). We also show that the model nests a wide spectrum of existing models—none of which can account for all the empirical facts simultaneously.

3.1 The Three-Equation Representation

The behavioral HANK model can be summarized by three equations: a Phillips curve, representing the aggregate supply side captured by equation (13), a rule for monetary policy (equation (7)), which together with the aggregate IS equation determines aggregate demand. To obtain the aggregate IS equation, we combine the hand-to-mouth households' consumption (8) with the consumption of unconstrained households (10) and their consumption Euler equation (12) (see appendix A for all the derivations).

Proposition 1. *The aggregate IS equation is given by*

$$\widehat{y}_t = \psi_f \mathbb{E}_t \widehat{y}_{t+1} - \psi_c \frac{1}{\gamma} \left(\widehat{i}_t - \mathbb{E}_t \pi_{t+1} \right), \tag{14}$$

where

$$\psi_f \equiv \bar{m}\delta = \bar{m} \left[1 + (\chi - 1) \frac{1 - s}{1 - \lambda\chi} \right] \quad \textit{and} \quad \psi_c \equiv \frac{1 - \lambda}{1 - \lambda\chi}.$$

Compared to RANK, two new coefficients show up: ψ_c and ψ_f. ψ_c governs the sensitivity of today's output with respect to the contemporaneous real interest rate. ψ_c is shaped by household heterogeneity, in particular by the share of H households λ and their income exposure χ. As the H households' incomes are more exposed to the aggregate income ($\chi > 1$), $\psi_c > 1$ which makes current output more sensitive to changes in the contemporaneous real interest rate due to general equilibrium forces, as we show later.

The second new coefficient in the behavioral HANK IS equation (14), ψ_f, captures the sensitivity of today's output with respect to changes in expected future output. ψ_f is shaped by household heterogeneity and the behavioral friction as it depends on the precautionary-savings motive, captured by δ, *and* the degree of bounded rationality of households as well as the interaction of these two. Given that $\chi > 1$, unconstrained households take into account that they will be more exposed to aggregate income fluctuations in case they become hand-to-mouth. Thus, income risk is countercyclical, which manifests itself in $\delta > 1$. Countercyclical risk induces compounding in the Euler equation and, thus, competes with the empirically observed underreaction of aggregate expectations ($\bar{m} < 1$) which induces discounting in the Euler equation. We see in the following sections that even for a small degree of bounded rationality—much smaller than the empirics suggest—the discounting through bounded rationality dominates the compounding through countercyclical income risk. Hence, in the behavioral HANK model it holds that $\psi_f < 1$ which makes the economy less sensitive to expectations and news about the future which is key to resolve the forward guidance puzzle as well as to obtain a determinate, locally unique equilibrium.

Equation (14) nests a wide range of existing IS equations: the IS equation in the standard rational-expectations RANK model by setting $\psi_f = \psi_c = 1$, RANK models deviating from FIRE by $\delta = \psi_c = 1$, Two-Agent NK (TANK) models by setting $\bar{m} = \psi_f = 1$, and rational HANK models by $\bar{m} = 1$.

Calibration. We set the share of H agents to one third, $\lambda = 0.33$, and μ^D such that $\psi_c = 1.2$, consistent with the empirical findings in Patterson (2019). This implies $\chi = 1.35$ and thus, implies that high-MPC households' incomes are relatively more sensitive to aggregate fluctuations induced by monetary policy, in line with the findings in Coibion et al. (2017) and Auclert (2019). We set the probability of a U household to become hand-to-mouth next period to 5.4%, i.e., $s = 0.946$ (this corresponds to $s = 0.8$ in annual terms). We focus on log utility, $\gamma = 1$, set $\beta = 0.99$, and the slope of the Phillips Curve to $\kappa = 0.02$, as in Bilbiie et al. (2021). The cognitive discounting parameter, $\bar{m}$ is set to 0.85, as explained in Section 2. Details on the calibration and a discussion of the robustness of our findings for different calibrations are presented in Appendix B. Note, that even when we vary certain parameters, we keep $\lambda < \chi^{-1}$.

3.2 Monetary Policy

We now show how the behavioral HANK model generates amplification of contemporaneous monetary policy through indirect effects while resolving the forward guidance puzzle at the same time. Additionally, we discuss determinacy conditions and show that the model remains stable at the effective lower bound.

To derive these results, it is sometimes convenient to combine the IS equation (14) with the static Phillips Curve (13) and the Taylor rule (7) so that we can represent the model in a single first-order difference equation:

$$\widehat{y}_t = \frac{\psi_f + \psi_c \frac{\kappa}{\gamma}}{1 + \psi_c \phi \frac{\kappa}{\gamma}} \mathbb{E}_t \widehat{y}_{t+1} - \frac{\psi_c \frac{1}{\gamma}}{1 + \psi_c \phi \frac{\kappa}{\gamma}} \varepsilon_t^{MP}. \tag{15}$$

General equilibrium amplification and forward guidance. We start by showing how the behavioral HANK model generates amplification of current monetary policy through indirect general equilibrium effects while simultaneously ruling out the forward guidance puzzle. The forward guidance puzzle states that announcements about future changes in the interest rate affect output today as strong (or even stronger) than contemporaneous changes in the interest rate.[11] Such strong effects of future interest rate changes, however, seem puzzling and are not supported by the data (Del Negro et al. (2015), Roth et al. (2021)).

Let us consider two different monetary policy experiments: (i) a contemporaneous monetary policy shock, i.e., a surprise decrease in the interest rate today, and (ii) a forward guidance shock, i.e., a news shock today about a decrease in the interest rate k periods in the future. In both cases, we focus on *i.i.d.* shocks and $\phi = 0$.[12]

Proposition 2. *In the behavioral HANK model, there is amplification of contemporaneous monetary policy relative to RANK if and only if*

$$\psi_c > 1 \Leftrightarrow \chi > 1, \tag{16}$$

and the forward guidance puzzle is ruled out if

$$\psi_f + \frac{\kappa}{\gamma}\psi_c < 1. \tag{17}$$

The behavioral HANK model generates amplification of contemporaneous monetary policy with respect to RANK whenever $\chi > 1$, that is, when high-MPC households' consumption is relatively more sensitive to aggregate income fluctuations. As discussed in Section 2.2, this is the case empirically.

After a decrease in the interest rate, wages increase and profits decline. As H agents receive a relatively smaller share of profits but fully benefit from the increase in wages, their income increases more than one-to-one with aggregate income. As they consume their income immediately, the initial effect on total output increases. The unconstrained households, on the other hand, experience a smaller increase in their income due to the fall in their profit income. To make up for this, they supply more labor and hence, produce the extra output. As a result, $\psi_c > 1$ and the increase in output is amplified through general equilibrium effects. To see the importance of GE or indirect effects, the following Lemma disentangles the direct and indirect effects.

Lemma 1. *The consumption function in the behavioral HANK model is given by*

$$\widehat{c}_t = [1 - \beta(1 - \lambda\chi)]\,\widehat{y}_t - \frac{(1-\lambda)\beta}{\gamma}\widehat{r}_t + \beta\bar{m}\delta(1 - \lambda\chi)\mathbb{E}_t\widehat{c}_{t+1}. \tag{18}$$

[11] Detailed analyses of the forward guidance puzzle in RANK are provided by McKay et al. (2016) and Del Negro et al. (2015).

[12] If we instead impose $\phi > 0$, contemporaneous amplification in the following proposition is not affected but the condition to rule out the forward guidance puzzle is further relaxed. Similarly, assuming completely fixed prices ($\kappa = 0$), as for example in Farhi and Werning (2019), or modelling forward guidance as changes in the *real* interest rate, as for example in McKay et al. (2016), would also leave the amplification condition unaltered but relaxes the condition to rule out the forward guidance puzzle.

Let ρ denote the exogenous persistence and define the indirect effects as the change in total consumption due to the change in total income but for fixed real rates. The share of indirect effects, Ξ^{GE}, out of the total effect is then given by

$$\Xi^{GE} = \frac{1-\beta(1-\lambda\chi)}{1-\beta\bar{m}\delta\rho(1-\lambda\chi)}.$$

Given our calibration and assuming an AR(1) monetary policy shock with a persistence of 0.6, indirect effects account for about 63%, consistent with larger quantitative models as for example in Kaplan et al. (2018) and thus, the model accounts for fact (i). Holm et al. (2021) state that the overall importance of indirect effects they find in the data is comparable to those in Kaplan et al. (2018), with the difference that these effects unfold after some time, whereas direct effects are more important on impact. Because in our stylized model the response to a monetary policy shock peaks on impact indirect effects are important right away. Slacalek et al. (2020) provide further evidence that indirect effects are strong drivers of aggregate consumption in response to monetary policy shocks.

For comparison, the representative agent model generates an indirect share of

$$\Xi^{GE} = \frac{1-\beta}{1-\beta\bar{m}\rho},$$

which, given our calibration, amounts to about 2%.

Note, that in the case of an i.i.d. shock the behavioral friction leaves the relative importance of direct vs. indirect effects—i.e., amplification of contemporaneous monetary policy—unaltered, as amplification of a contemporaneous i.i.d. shock is solely determined by the static redistribution towards the high MPC households. It is through these indirect general equilibrium effects that monetary policy gets amplified as the H households do not directly respond to interest rate changes because they do not participate in asset markets.

Turning to forward guidance, note, that the forward guidance puzzle is ruled out if the term $\frac{\psi_f+\psi_c\frac{\kappa}{\gamma}}{1+\psi_c\phi\frac{\kappa}{\gamma}}$ in front of $\mathbb{E}_t\widehat{y}_{t+1}$ in the first-order difference equation (15) is smaller than 1. Given that we consider $\phi = 0$, this boils down to the condition stated in Proposition 2.

What determines whether condition (17) holds or not? First, note that as in the discussion of contemporaneous monetary policy, with $\chi > 1$ the income of H agents moves more than one for one with aggregate income. In this case, unconstrained households who self-insure against becoming hand-to-mouth in the future want less insurance when they expect a decrease in the interest rate since if they become hand-to-mouth they would benefit more from the increase in aggregate income. Hence, after a forward guidance shock, unconstrained households decrease their precautionary savings which compounds the increase in output today ($\delta > 1$). Yet, as households are boundedly rational, they cognitively discount these effects taking place in the future. Importantly, unconstrained households cognitively discount both the usual consumption-smoothing response due to the future increase in consumption as well as the general equilibrium implications for their precautionary savings, thereby decreasing the effects of the forward guidance shock on today's consumption. Thus, the model not only accounts for facts (i) and (ii) but simultaneously accounts for fact (iii).

This last part clearly illustrates the main interaction of bounded rationality and household heterogeneity that enables the behavioral HANK model to resolve the forward guidance puzzle while simultaneously generating amplification through indirect effects. Households fully understand their idiosyncratic risk of switching their type as well as the implications of switching type in case there are no aggregate shocks, i.e., in the steady state. If the monetary authority makes an unexpected announcement about its future policy, however, behavioral households do not fully incorporate the effects of this policy on their own income risk and thus, their precautionary savings. Already a small underreaction of the behavioral households is enough to resolve the forward guidance puzzle. Given our calibration there is no forward guidance puzzle in the behavioral HANK model as long as $\bar{m} < 0.94$ which is above the upper bounds for empirical estimates (see Section 2).[13]

We now compare the behavioral HANK model to its rational counterpart to show how the behavioral HANK model overcomes a shortcoming inherent in the rational HANK model – the *Catch-22* (Bilbiie (2021); see also Werning (2015)). The Catch-22 describes the tension that the rational HANK model can either generate amplification of contemporaneous monetary policy *or* solve the forward guidance puzzle. To see this, note that with $\bar{m} = 1$ the forward guidance puzzle is resolved when

$$\delta + \frac{1}{\gamma} \frac{1-\lambda}{1-\lambda\chi} \kappa < 1$$

which requires $\chi < 1$, as otherwise $\delta > 1$. Assuming $\chi < 1$, however, leads to *dampening* of contemporaneous monetary policy instead of amplification. We graphically illustrate the Catch-22 of the rational model and its resolution in the behavioral HANK model in Appendix C. Note that also rational TANK models (thus, turning off type switching) or the behavioral RANK model would not deliver amplification and resolve the forward guidance puzzle simultaneously. TANK models would face the same issues as the rational RANK model in the sense that they cannot solve the forward guidance puzzle while bounded rationality in a RANK model does not deliver initial amplification.

3.2.1 Stability at the Effective Lower Bound

In this section, we revisit the determinacy conditions in the behavioral HANK model and discuss the implications for the stability at the effective lower bound constraint on nominal interest rates.

According to the Taylor principle, monetary policy needs to respond sufficiently strongly to inflation in order to guarantee a determinate equilibrium. In the rational RANK model the Taylor

[13] A related paradox in the rational model is that as the persistence of the shock increases, the effects become unboundedly large and as the persistence approaches unity, an exogenous increase in the nominal interest rate becomes expansionary. The behavioral HANK model, on the other hand, does not suffer from this. We elaborate these points in more detail in Appendix D.3.

principle is given by $\phi > 1$, where ϕ is the inflation-response coefficient in the Taylor rule (7). We now derive a similar determinacy condition in the behavioral HANK model and show that both household heterogeneity and bounded rationality affect this condition. The following proposition provides the behavioral HANK Taylor principle.[14]

Proposition 3. *The behavioral HANK model has a determinate, locally unique equilibrium if and only if:*

$$\phi > \phi^* = 1 + \frac{\bar{m}\delta - 1}{\frac{\kappa}{\gamma}\frac{1-\lambda}{1-\lambda\chi}}. \tag{19}$$

We obtain Proposition 3 directly from the difference equation (15). For determinacy, we need that the coefficient in front of $\mathbb{E}_t\widehat{y}_{t+1}$ is smaller than 1 (the eigenvalues associated with any exogenous variables are assumed to be $\rho < 1$, and are thus stable). Solving this condition for ϕ yields Propositon 3. Appendix A.4 outlines the details and extends the result to more general Taylor rules.

To understand the condition in Proposition 3, consider first $\bar{m} = 1$ and, thus, focus solely on the role of household heterogeneity. With $\chi > 1$, it follows that $\phi^* > 1$ and, hence, the threshold is higher than the RANK Taylor principle states. This insufficiency of the Taylor principle in the rational HANK model has been shown by Bilbiie (2021) and in a similar way by Ravn and Sterk (2021) and Acharya and Dogra (2020). As a future aggregate sunspot increases the income of households in state H disproportionately, unconstrained households cut back on precautionary savings today which further increases output today. This calls for a stronger response of the central bank to not let the sunspot become self-fulfilling.

On the other hand, bounded rationality $\bar{m} < 1$ relaxes the condition as unconstrained households now cognitively discount both the future aggregate sunspot as well as its implications for their idiosyncratic risk. A smaller response of the central bank is needed in order to prevent the sunspot to become self-fulfilling. Given our calibration the cutoff value for $\bar{m}$ to restore the RANK Taylor principle in the behavioral HANK model is 0.966. What is more, given our baseline choice of $\bar{m} = 0.85$, we obtain $\phi^* = -4$. Thus, in the behavioral HANK model it is not necessary that monetary policy responds to inflation at all as the economy features a stable unique equilibrium even under an interest rate peg.

Stability at the effective lower bound. Related to the indeterminacy issues under a peg the traditional New Keynesian model struggles to explain how the economy can remain stable when the effective lower bound (ELB) on nominal interest rates is binding for an extended period of time, as observed in many advanced economies over recent decades (see, e.g., Debortoli et al. (2020) and Cochrane (2018)). If the ELB binds for a sufficiently long time, RANK predicts unreasonably large recessions and, in the limit case in which the ELB binds forever, even indeterminacy.[15] Similar to

[14] We focus on local determinacy and bounded equilibria.

[15] A forever binding ELB basically implies that the Taylor coefficient is equal to zero and, thus, the nominal rate is pegged at the lower bound, thereby violating the Taylor principle. Note, that this statement also extends

the forward guidance puzzle, this is even more severe in rational HANK models.

We now show that the behavioral HANK model resolves these issues and thus accounts for fact (iv). To this end, let us add a *natural rate shock* (i.e., a demand shock) $\widehat{r}_t^n$ to the IS equation:

$$\widehat{y}_t = \psi_f \mathbb{E}_t \widehat{y}_{t+1} - \psi_c \left(\widehat{i}_t - \mathbb{E}_t \pi_{t+1} - \widehat{r}_t^n \right).$$

We assume that in period t the natural rate decreases to a value $\widetilde{r}^n$ that is sufficiently negative such that the natural rate in levels is below the ELB. The natural rate stays at $\widetilde{r}^n$ for $k \geq 0$ periods and after k periods the economy returns immediately back to steady state. Agents correctly anticipate the length of the binding ELB. Iterating the IS equation forward, it follows that output in period t is given by

$$\widehat{y}_t = -\frac{1}{\gamma} \psi_c \underbrace{\left(\widehat{i}_{ELB} - \widetilde{r}^n \right)}_{>0} \sum_{j=0}^{k} \left(\psi_f + \frac{\kappa}{\gamma} \psi_c \right)^j , \tag{20}$$

where the term $\left(\widehat{i}_{ELB} - \widetilde{r}^n \right) > 0$ captures the shortfall of the policy response due to the binding ELB. Under rational expectations, we have that $\psi_f > 1$, meaning that output implodes as $k \to \infty$. The same is true in the rational RANK model which is captured by $\psi_f = \psi_c = 1$. In the behavioral HANK model, however, this is not the case. As long as $\psi_f + \frac{\kappa}{\gamma} \psi_c < 1$ the output response in t is bounded even as $k \to \infty$. It follows that $\bar{m} < 0.94$ is enough to rule out unboundedly-severe recessions at the ELB even if the ELB is expected to persist forever. We graphically illustrate in Appendix C that the behavioral HANK model remains stable also for long spells of the ELB in which output in the rational models collapses.

Fiscal multipliers. In Appendix D, we derive analytically the fiscal multiplier in the behavioral HANK model and show that it generates consumption responses that are consistent with empirical evidence (Dupor et al. (2021)). As for monetary policy, the exposure of high-MPC households to aggregate income changes leads to positive fiscal multipliers in the benchmark case of constant real rates and cognitive discounting ensures that the multiplier remains finite, even for highly-persistent government spending shocks.

Nesting existing models. The behavioral HANK model nests three classes of models in the literature: the representative-agent rational expectations (RANK) model for $\lambda = 0$ and $\bar{m} = 1$ (see Galí (2015), Woodford (2003)), representative agent models without FIRE for $\lambda = 0$ and $\bar{m} \in (0, 1)$ as, for example, in Gabaix (2019), Angeletos and Lian (2018), and Woodford (2019); and TANK and tractable HANK models as e.g. in Bilbiie (2008), Bilbiie (2021), McKay et al. (2017), or Debortoli and Galí (2018) for $\bar{m} = 1$. In contrast to these classes of models, the behavioral HANK model combines the indirect general equilibrium amplification of monetary policy with a resolution of the forward guidance puzzle and stability at the ELB. In representative agent models monetary

to models featuring more elaborate monetary policy rules including Taylor rules responding to output or also the Wicksellian price-level targeting rule, as they all collapse to a constant nominal rate in a world of an ever-binding ELB.

policy mainly works through direct intertemporal substitution channels and do not feature $\psi_c \neq 1$, rational HANK models on the other hand do not feature $\psi_f < 1$ and $\psi_c > 1$ simultaneously as discussed in Section 3.2. Hence, if we abstract from either bounded rationality or household heterogeneity (or both), the models fails in accounting for the four facts about the transmission mechanisms and effectiveness of monetary policy.

4 The Full Behavioral HANK Model

In this section, we move from the limited-heterogeneity setup described in Section 2 to an incomplete markets setup as in Bewley (1986), Huggett (1993), and Aiyagari (1994) which is standard in the quantitative HANK literature. There is a continuum of ex-ante identical households all subject to idiosyncratic productivity risk, incomplete markets, and exogenous borrowing constraints. Households self-insure against their idiosyncratic risk by accumulating government bonds. Bonds are now in positive net supply as the fiscal authority issues real government debt, B_t^G. To finance its interest payments, the fiscal authority collects tax payments from households. Given these assumptions, households differ ex-post in their productivity level, e, and their wealth B. The households' utility function is the same as in the tractable model (equation (1)).

Household i faces the budget constraint

$$C_{i,t} + \frac{B_{i,t+1}}{R_t} = B_{i,t} + W_t e_{i,t} N_{i,t} + D_t d(e_{it}) - \tau_t(e_{it})$$

and the borrowing constraint $B_{i,t+1} \geq \underline{B}$, where $\underline{B}$ denotes an exogenous borrowing limit. Households receive a share of the dividends, $D_t d(e_{it})$, conditional on their productivity, and pay taxes also conditional on their productivity, $\tau_t(e_{it})$. Note that this way, taxes are non-distortionary in the sense that they do not show up in the household's first-order conditions. In line with McKay et al. (2016), we assume that only the households with the highest productivity pay taxes.

We introduce bounded rationality in the same way as in our tractable model. Households are fully rational with respect to their idiosyncratic risk, but they cognitively discount the effects of aggregate shocks. Let $\bar{C}_{i,t} = C(e_{i,t}, B_{i,t}, \bar{Z})$ denote consumption of household i in period t with idiosyncratic productivity $e_{i,t}$ and asset holdings $B_{i,t}$ when all aggregate variables are in steady state, indicated by $\bar{Z}$. Here, Z potentially denotes a whole matrix of aggregate variables, including, for example, news shocks (i.e., forward guidance shocks). In other words, $\bar{C}_{i,t}$ denotes consumption of household i with productivity $e_{i,t}$ and asset holdings $B_{i,t}$ in the stationary equilibrium. In case of an aggregate shock, $Z_t \neq \bar{Z}$, consumption is denoted by $C_{i,t} = C(e_{i,t}, B_{i,t}, Z_t)$. Assuming as in the tractable model that the household anchors her expectations to the stationary equilibrium implies the following Euler equation

$$\begin{aligned} C_{i,t}^{-\gamma} &\geq \beta R_t \mathbb{E}_t^{BR}\left[C_{i,t+1}^{-\gamma}\right] \\ &= \beta R_t \mathbb{E}_t^{BR}\left[\bar{C}_{i,t+1}^{-\gamma} + \left(C_{i,t+1}^{-\gamma} - \bar{C}_{i,t+1}^{-\gamma}\right)\right] \\ &= \beta R_t \mathbb{E}_t\left[\bar{C}_{i,t+1}^{-\gamma} + \bar{m}\left(C_{i,t+1}^{-\gamma} - \bar{C}_{i,t+1}^{-\gamma}\right)\right], \end{aligned} \tag{21}$$

where the rational expectations operator $\mathbb{E}_t[\cdot]$ denotes the expectations that a fully rational household would have in the behavioral economy. Equation (21) illustrates that when households form expectations about their marginal utility in the next period, their expectations about the marginal utilities associated with each possible individual state are anchored to the marginal utilities associated with these states in stationary equilibrium.[16] As usual, the Euler equation holds with equality for non-constrained households while it holds with strict inequality for households whose borrowing constraint binds. The labor-leisure condition is identical to the one in the tractable model and holds for every household. With rational expectations ($\bar{m} = 1$), the model collapses to a standard rational one-asset HANK model, similar to McKay et al. (2016) or Debortoli and Galí (2018).

Calibration. Most of the calibration is standard in the literature. As in the tractable model, we set the cognitive discounting parameter to $\bar{m} = 0.85$. We set the dividend shares in order to match the fact that households with higher MPCs (which is highly correlated with lower productivity states) tend to be more exposed to aggregate income changes induced by monetary policy (Patterson (2019)). Consistent with the data and the tractable model, this results in a calibration in which households with higher productivity receive a larger share of the dividends than households with a lower productivity. Our calibration implies a correlation between the individual MPC and the change in gross incomes after a monetary policy shock of 0.36. We set the amount of government debt to match the aggregate MPC of 0.16 out of an income windfall of 500$, as in Kaplan et al. (2018). This results in a government debt-to-annual-GDP level of 50%. Further details and the rest of the calibration are relegated to Appendix E.

4.1 Monetary Policy

We now consider two monetary policy experiments. First, a one-time conventional expansionary monetary policy shock and second, a forward guidance shock that is announced today to take place k periods in the future. In particular, we assume that the monetary authority announces in period 0 to decrease the nominal interest rate by 10 basis points in period k and keeps the nominal rate at its steady state value in all other periods. Following Farhi and Werning (2019), we focus on the case with fully rigid prices such that the change in the nominal rate translates one for one to changes in the real rate and is thus also consistent with the exercise in McKay et al. (2016). In addition, we also follow Farhi and Werning (2019) and McKay et al. (2016) in assuming that the government debt level remains constant, $B_t^G = \bar{B}^G$.

Figure 1 shows on the vertical axis the response of output in period 0, dY_0, to an announced real rate change implemented in period k (horizontal axis). The white horizontal line represents the response in the rational RANK model (normalized to 1). The constant response in RANK is

[16] Note, that this assumption about the anchor value has the desirable property that to first order the individual endogenous state variable, $B_{i,t}$, is not cognitively discounted, reflecting the fact that the household knows her asset holdings.

Figure 1: Monetary Policy and Forward Guidance

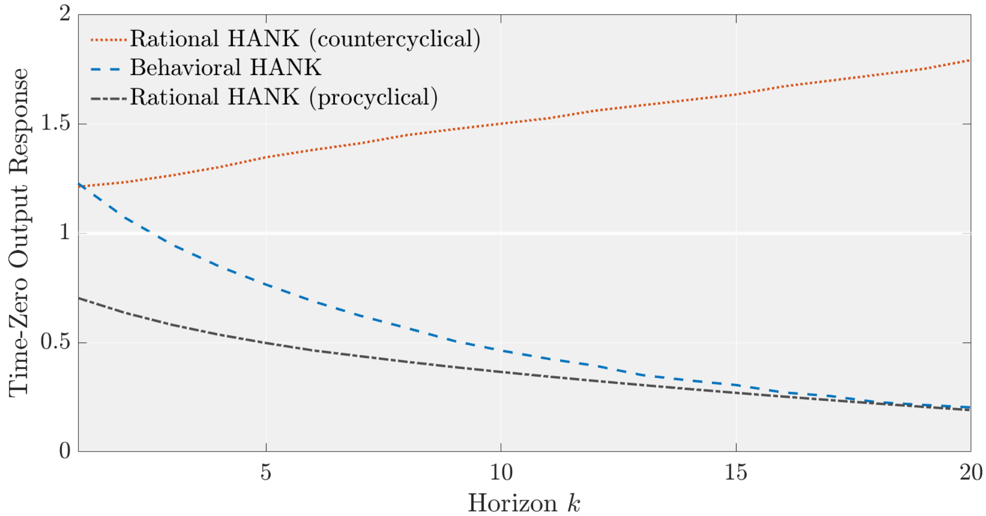

Note: This figure shows the response of total output in period 0 to anticipated one-time monetary policy shocks occurring at different horizons k, relative to the response in the representative agent model under rational expectations (normalized to 1). The blue-dashed line shows the results for the behavioral HANK model, the orange-dotted line for the rational HANK model with countercyclical inequality and the black-dashed-dotted line for the rational HANK model with procyclical inequality.

a consequence of the assumption that forward guidance is implemented through changes in the real rate.

The blue-dashed line shows the results for the behavioral HANK model. We see that contemporaneous monetary policy has stronger effects than in RANK. As estimated by Patterson (2019), we obtain an amplification of 20% compared to the case in which all households are equally exposed. The intuition is the same as in the tractable model: as households with higher MPCs tend to be more exposed to aggregate income changes, monetary policy is amplified through indirect general equilibrium effects. Turning again to an AR(1)-process with a persistence of 0.6, we find that indirect effects account for 53% of the total effect in the quantitative behavioral HANK consistent with our tractable model. At the same time, the behavioral HANK model does not suffer from the forward guidance puzzle, as shown by the decline in the blue-dashed line. Interest rate changes announced to take place in the future have relatively weaker effects on contemporaneous output and the effects decrease with the horizon.[17]

In contrast, the orange-dotted and the black-dashed-dotted lines highlight the tension in rational HANK models. When households with high MPCs tend to be more exposed to aggregate income fluctuations—which corresponds to $\chi > 1$ in the tractable model and which we refer to as the *countercyclical* HANK model—contemporaneous monetary policy is as strong as in the behavioral model. But with rational expectations the amplification through indirect effects extends intertemporally and results in an aggravation of the forward guidance puzzle. Indeed, we see

[17]We find that for our baseline calibration the behavioral HANK model resolves the forward guidance puzzle as long as $\bar{m} < 0.96$.

from the orange-dotted line that the farther away the announced interest rate change takes place, the stronger the response of output today. A change that is announced to take place in twenty quarters leads to a response of today's output that is about 50% stronger than a contemporaneous monetary policy shock.

When, in contrast to the data, households with higher MPCs tend to be less exposed to aggregate income fluctuations—$\chi < 1$ in the tractable model and which we refer to as the *procyclical* HANK model—the rational HANK model resolves the forward guidance puzzle (see McKay et al. (2016)). But the procyclical HANK model is unable to generate amplification of contemporaneous monetary policy (see black-dashed-dotted line) and will have quite different policy implications, as we will see in Section 5.

Stability at the effective lower bound. To test the stability of the model at the effective lower bound—fact (iv)—we consider a shock to the discount factor that pushes the economy to the ELB for twelve periods, in the behavioral and the rational model. After that the shock jumps back to its steady state value. Consistent with the tractable model, the recession in the rational model is substantially more severe. While output drops on impact by 5% in the behavioral model, it drops by 10% in the rational model (see Appendix E for details).

Fiscal multipliers. We furthermore show that the quantitative behavioral HANK model generates positive consumption responses to a fiscal spending shock under a constant real rate. To this end, we consider a temporary increase in government consumption financed by lump-sum transfers. To such a fiscal policy shock private consumption increases independently of the persistence of the fiscal shock (see Appendix E for details).

Overall, we conclude that our main insights of the tractable behavioral HANK model carry over to the quantitative behavioral HANK model and that also the quantitative model simultaneously accounts for the facts (i)-(iv) outlined in the introduction.

4.2 Heterogeneous Cognitive Discounting

So far, we have assumed that all households exhibit the same degree of rationality. In reality, however, there might be heterogeneity with respect to the degree of cognitive discounting. Indeed, as we show in Appendix B.1, while underreaction is found across all income groups, the data suggests that higher income households deviate somewhat less from rational expectations. To model this, we assume that a household's rationality is a function of her productivity level e: $\bar{m}(e = e_L) = 0.8$, $\bar{m}(e = e_M) = 0.85$ and $\bar{m}(e = e_H) = 0.9$.

This parameterization serves three purposes: first, the lowest-productivity households exhibit the largest deviation from rational expectations and the degree of rationality increases monotonically with productivity. Second, the average degree of bounded rationality remains 0.85 such that we can isolate the effect of heterogeneity in bounded rationality from its overall level. And third, this is a rather conservative parameterization—both in terms of the degree of heterogeneity and in the level of rationality—compared to the results in the data which points more towards lower levels of rationality across all households and less dispersion. We discuss alternative calibrations—including one in which a subgroup of households is fully rational—in Appendix E.

Figure 2: Heterogeneous $\bar{m}$ and Monetary Policy

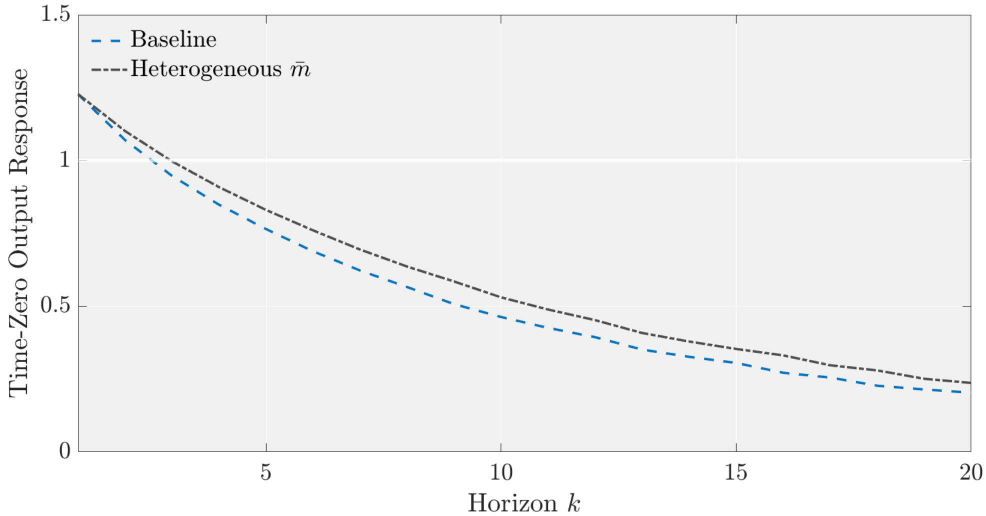

Note: This figure shows the response of total output in period 0 to anticipated i.i.d. monetary policy shocks occurring at different horizons k for the baseline calibration with $\bar{m} = 0.85$ for all households (blue-dashed line) and for the model in which households differ in their levels of cognitive discounting (black-dashed-dotted line).

Figure 2 compares the model with heterogeneous degrees of bounded rationality (black-dashed-dotted line) to our baseline quantitative behavioral HANK model (blue-dashed line) for the same monetary policy experiments as above. The effect of a contemporaneous monetary policy shock is practically identical across the two scenarios consistent with the insight that amplification of a contemporaneous monetary policy shock is barely affected by the degree of rationality. At longer horizons, however, monetary policy is more effective in the economy in which households differ in their degrees of rationality.

There are two competing effects: first, high productivity households are now more rational such that they react stronger to announced future changes in the interest rate compared to the baseline which increases the effectiveness of forward guidance. Second, low productivity households are less rational which tends to dampen the effectiveness of forward guidance. Yet, a large share of low productivity households are at their borrowing constraint and, thus, do not directly react to future changes in the interest rate anyway while most of the high productivity households are unconstrained. Hence, the first effect dominates and forward guidance is more effective compared to the baseline model. Overall, however, the differences across the two calibrations are rather small. As we show in Appendix B.1, even when the highest productivity households are fully rational the forward guidance puzzle is resolved and the effects of forward guidance vanish quite quickly with the horizon.

5 Policy Implications: Inflationary Supply Shocks

Having established that the behavioral HANK model is consistent with recent facts about the transmission and effectiveness of monetary policy, we now use it to revisit the policy implications of inflationary supply shocks. Many advanced economies have recently experienced a dramatic surge in inflation and at least part of this is attributed to disruptions in production, such as supply-chain "bottlenecks" (see, e.g., di Giovanni et al. (2022)). We model these disruptions as a negative total factor productivity (TFP) shock and analyze how monetary policy has to be implemented after such a shock in order to stabilize inflation.

Production of intermediate-goods firm j is now given by $Y_t(j) = A_t N_t(j)$, where A_t is total factor productivity following an AR(1)-process, $A_t = (1 - \rho_A)\bar{A} + \rho_A A_{t-1} + \varepsilon_t^A$, and ε_t^A is an i.i.d. shock, $\bar{A}$ the steady-state level of TFP and ρ_A the persistence of A_t which we set to $\rho_A = 0.9$. Each firm can adjust its price with probability 0.15 in a given quarter and we assume that firms have rational expectations to fully focus on the role of bounded rationality on the household side.

Government debt is time-varying and total tax payments, T_t, follow a standard debt feedback rule, $T_t - \bar{T} = \vartheta \frac{B_{t+1} - \bar{B}}{\bar{Y}}$, where we set $\vartheta = 0.05$. We consider two different monetary policy regimes: in the first one, monetary policy follows a strict inflation-targeting rule and implements a zero inflation rate in all periods. In the second one, monetary policy follows a standard Taylor rule.

The size of the shock is such that output in the model with fully-flexible prices, complete markets and rational expectations—what we from now on call *potential output*—decreases by 1% in terms of deviations from its steady state. We normalize the leisure parameter in the complete markets, flexible price model such that it has the same steady state output as our behavioral HANK model. The *output gap* is then defined as the difference between actual output and potential output divided by steady state output.

Figure 3 shows the impulse-response functions of output, the output gap, inflation, nominal interest rates, government debt and the consumption Gini index as a measure of inequality after the negative supply shock when monetary policy fully stabilizes inflation. The blue-dashed lines show the responses in the behavioral HANK model and the orange-dotted lines in the rational HANK model. The output responses are almost indistinguishable across the two models and practically identical to the fall in potential output such that the output gap is essentially zero.

Yet, the reaction of monetary policy differs significantly across the two models. The nominal interest rate in the behavioral HANK model increases twice as much on impact as in the rational HANK model. The reason is that behavioral households cognitively discount the future higher interest rates that they expect due to the persistence of the shock. Hence, these expected higher future rates are less effective in stabilizing inflation today. Thus, to induce zero inflation in every period, monetary policy needs to increase interest rates by more than in the rational HANK model, in which the expected future interest rate hikes are very powerful. As this line of reasoning applies in each period, the interest rate in the behavioral HANK model remains above the interest rate

Figure 3: Inflationary supply shock: strict inflation-targeting

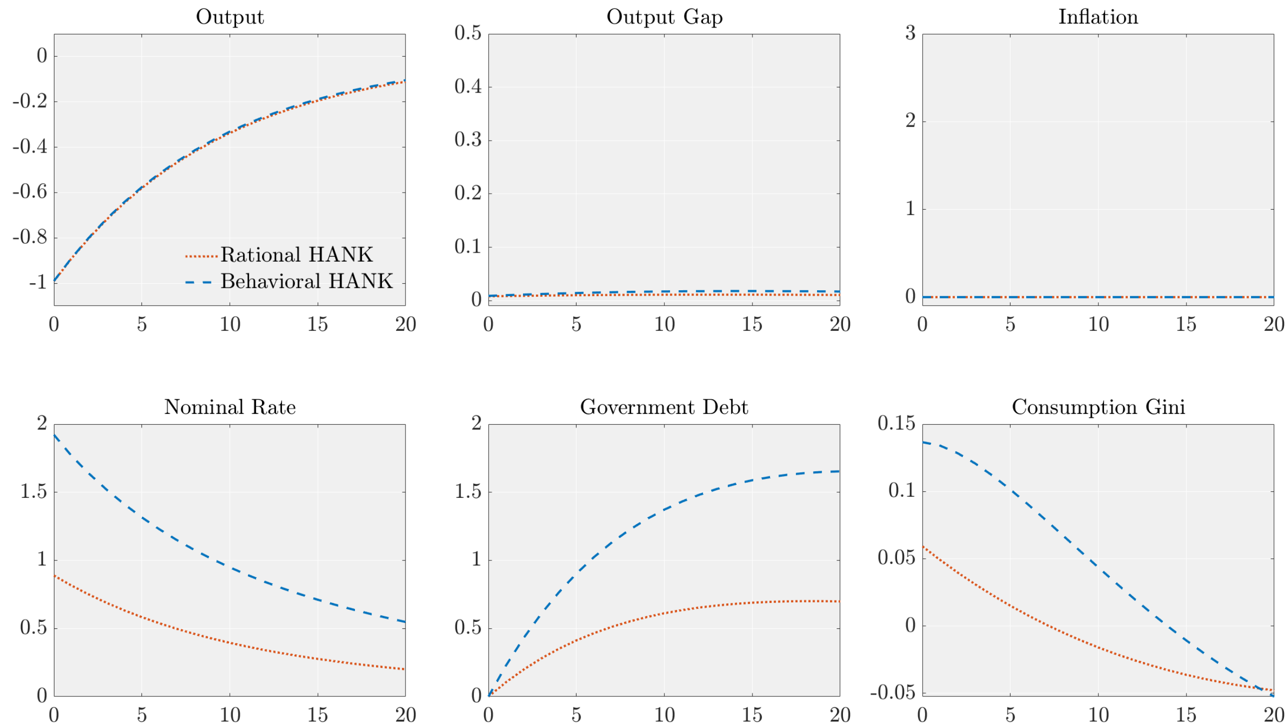

Note: This figure shows the impulse responses after a TFP shock that decreases potential output by 1% in the inflation-stabilizing monetary policy regime. Output and the output gap are shown as percentage deviations from steady state output, the nominal interest rate and inflation as annualized percentage points and the government debt level as percentage point deviations of the debt-per-annual GDP level. The lower-right figure shows the change in the consumption Gini index as a percentage deviation from the stationary equilibrium.

in the rational model.

Raising interest rates increases the cost of debt for the government which it finances in the short run by issuing additional debt. The bottom-middle panel in Figure 3 shows that government debt in the behavioral model increases by more than twice as much as in the rational model. Thus, the fiscal implications of monetary policy are larger due to the stronger response of monetary policy.

On top of the stronger increase in government debt and interest rates, consumption inequality increases more strongly in the behavioral model compared to the rational model. The reason is that along the wealth distribution, increases in the real interest rate redistribute to wealthier households and, hence, to households who already have a higher consumption level. As the increases in the real interest rate are higher in the behavioral HANK model, these redistribution effects are more pronounced. Because monetary policy fully stabilizes inflation and the output gap, dividends and wages fall by the same relative amount after the productivity shock, such that each household's labor and dividend income falls by the same amount. Hence, the redistribution channels present in Sections 3 and 4 after policy shocks are muted here.

In Appendix F.6, we show that both the increase in consumption inequality and the implications for fiscal policy are even more evident when initial debt levels are high, especially in the behavioral HANK model.

Figure 4: Inflationary supply shock: Taylor rule

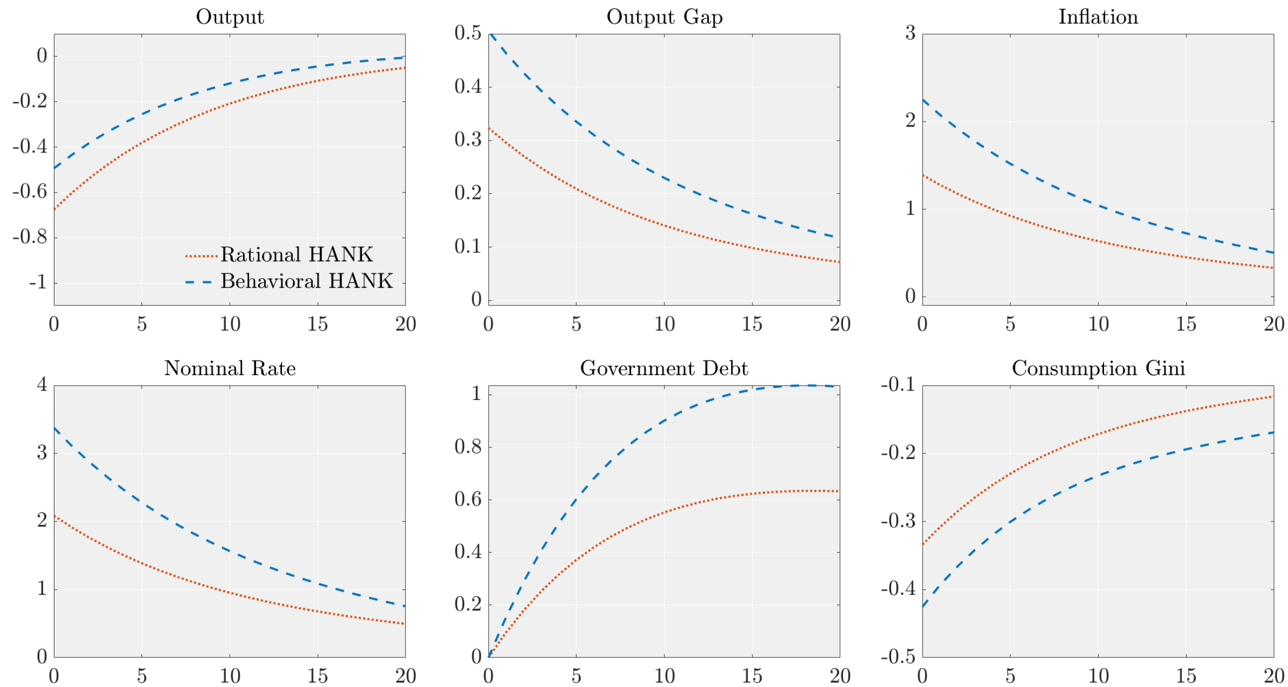

Note: This figure shows the impulse responses after a productivity shock for the case that monetary policy follows a Taylor rule. Output and the output gap are shown as percentage deviations from steady state output, the nominal interest rate and inflation as annualized percentage points and the government debt level as percentage point deviations of the debt-per annual-GDP level. The lower-right figure shows the change in the consumption Gini index as a percentage deviation from the stationary equilibrium.

How do these results change when monetary policy follows a standard Taylor rule? Figure 4 shows the case when monetary policy follows a simple Taylor rule (7) with a feedback coefficient of 1.5. Inflation and the output gap increase by substantially more in the behavioral HANK than in its rational counterpart even though the (nominal and real) interest rate increases by more as well. The reason is a novel amplification channel arising because of household heterogeneity, cognitive discounting and the interaction of the two: the positive output gap increases wages and decreases profits relative to the inflation stabilizing regime in the same way as expansionary policy shocks in Sections 3 and 4 do. This redistributes on average towards lower income and higher MPC households which further increases the output gap and inflation. In addition, the higher expected real interest rates in response to the negative supply shock lead to a negative deviation of expected consumption from stationary equilibrium. In the behavioral HANK model, households cognitively discount these negative deviations of expected consumption from stationary equilibrium and hence, they decrease today's consumption by less compared to fully rational households. This further increases the output gap which amplifies the redistribution to high MPCs households which again amplifies the increase in the output gap until the economy ends up in an equilibrium with a higher output gap and higher inflation.[18]

[18]Note that the fact that both the underlying heterogeneity and bounded rationality amplify persistent supply shocks is in stark contrast to a persistent demand shock like a monetary policy shock in which case heterogeneity

Thus, while inflation and the output gap increase substantially when monetary policy follows a Taylor rule, consumption inequality is now decreasing instead of increasing both in the rational as well as in the behavioral HANK model and it decreases even more in the behavioral model (see lower-right panel in Figure 4). While higher interest rates still redistribute to relative consumption-rich households, this effect on consumption inequality is now dominated by the increase in the output gap which redistributes to relatively consumption-poor households. Finally, the government debt level also increases more than in the rational HANK model, but both increase by less than when monetary policy fully stabilizes inflation.

Decomposition of the amplification channel. In the behavioral HANK model inflation increases substantially more than in RANK (inflation in RANK increases by 1 percentage point on impact, see Appendix F.3). How much of the additional inflation increase in the behavioral HANK model is due to the underlying heterogeneity, cognitive discounting and the interaction of the two? Figure 5 decomposes the amplification channel into these three components. It shows the additional inflation increase in the behavioral HANK model compared to the inflation increase in the RANK benchmark, that is $\Delta\pi_t^{BHANK} \equiv \pi_t^{BHANK} - \pi_t^{RANK}$, depicted by the black-solid line. The figure further shows the additional inflation increase in the rational HANK model $\Delta\pi_t^{HANK} \equiv \pi_t^{HANK} - \pi_t^{RANK}$ (orange-dotted line) and the sum of the additional inflation increases in the rational HANK model and the behavioral RANK model, $\Delta\pi_t^{HANK} + \Delta\pi_t^{BRANK}$, with $\Delta\pi_t^{BRANK} \equiv \pi_t^{BRANK} - \pi_t^{RANK}$, which is depicted by the blue-dashed line.

$\Delta\pi_t^{HANK}$ shows how much of the amplification compared to RANK is caused solely by household heterogeneity (indicated by the orange area). $\Delta\pi_t^{BRANK}$ (which corresponds to the blue area) can be interpreted as the amplification coming from cognitive discounting alone. Correspondingly, $\Delta\pi_t^{HANK} + \Delta\pi_t^{BRANK}$ captures the amplification that would arise if cognitive discounting and household heterogeneity are both present, but would not interact with each other. Thus, the difference between between $\Delta\pi_t^{BHANK}$ and $\Delta\pi_t^{HANK} + \Delta\pi_t^{BRANK}$ (the gray area) depicts the extra amplification coming from the complementarity of heterogeneity and cognitive discounting.

Under our baseline calibration, this complementarity amounts to about 10% on impact of the inflation response in RANK (the inflation response in RANK is 1 percentage point on impact). As the additional increase in the behavioral HANK, $\Delta\pi_t^{BHANK}$, is about 1.2 percentage points, the complementarity explains about 8.5% of the *additional increase*. Under a more extreme calibration, i.e., a cognitive discounting parameter of 0.6 instead of 0.85 and an unequal exposure of households that implies an amplification of 30% instead of 20% of conventional monetary policy, implies that the complementarity amounts to more than 70% of the impact inflation response in RANK (or about 25% of the *additional increase* in that case). As we show in Figure 16 in Appendix F.1, the inflation increase that is due to the interaction becomes even larger than the one due to household

would amplify and cognitive discounting dampen the effect. For example, a monetary shock of 1 percentage point with persistence of 0.6 increases inflation on impact by 1.44 annualized percentage points in the rational HANK model, by 1.24 in the behavioral HANK and by 1.03 in RANK.

heterogeneity under this alternative calibration.

Figure 5: Decomposition of the Additional Inflation Increase

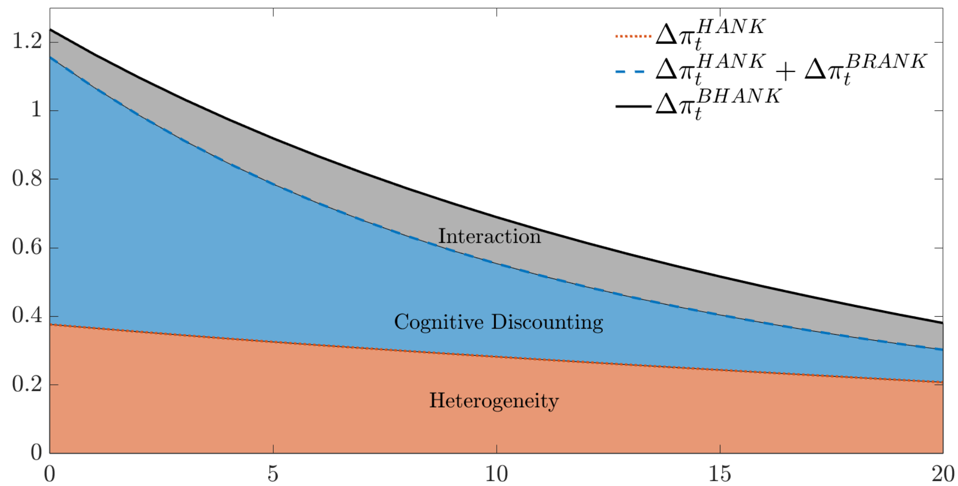

Note: This figure shows the additional inflation increase in the rational HANK model compared to the rational RANK model (orange-dotted line), the sum of the additional increase in the rational HANK and the behavioral RANK compared to the rational RANK model (blue-dashed line) and the additional increase in behavioral HANK model compared to the rational RANK model (black-solid line).

Given the larger sensitivity of inflation to supply shocks due to this novel amplification channel, our model may hence offer a (partial) explanation for why many advanced economies have seen large inflation increases following the Covid-19 pandemic.

Comparison to the procyclical HANK model. One of the reasons why the behavioral HANK model amplifies supply shocks is that it is less responsive to expected future interest rates. A natural question is then: how do its policy implications compare to those derived in rational HANK models that are calibrated to resolve the forward guidance puzzle? As shown in Section 4, when all households receive an equal share of the dividends, the rational model can resolve the forward guidance puzzle (McKay et al. (2016)). This implies that households with high MPCs benefit less from income increases induced by monetary policy, thereby violating fact (ii).

Figure 18 in the appendix shows that this "procyclical" rational HANK model predicts a much weaker response of inflation to the same supply shock in the case of a standard Taylor rule. The reason is that now the positive output gap redistributes on average to high-income and low MPC households which further dampens aggregate demand. In other words, this model features a dampening channel compared to RANK after supply shocks instead of an amplification channel as in the behavioral HANK model.[19]

The two models also differ in terms of their cross-sectional implications: in the procyclical

[19] Another take-away is that for a given persistent demand shock, the behavioral HANK model and a recalibrated version of the procyclical HANK model could be observationally equivalent in terms of the output and inflation response. Yet, these two models differ drastically after supply shock.

HANK model, consumption inequality increases strongly whereas it decreases in the behavioral HANK model.

If monetary policy fully stabilizes inflation, the procyclical HANK model becomes observationally equivalent with the countercyclical HANK model (our baseline rational HANK model that accounts for fact (ii))—both in terms of aggregates and in the cross-section—but not with the behavioral HANK model. The reason is that now the output gap is closed which switches off the channels that weaken the effects of expected future interest rate changes in the forward guidance experiment. Households still fully incorporate the whole path of future interest rates, inconsistent with the large empirical evidence on households' inattention and underreaction to monetary policy (Coibion and Gorodnichenko (2015), Coibion et al. (2020), D'Acunto et al. (2020) or Roth et al. (2021)). In the behavioral HANK model, however, the resolution of the forward guidance puzzle is due to a weakening of the expectations channel more generally, even when the output gap is kept at zero. Thus, the underlying reason for the resolution of the forward guidance puzzle matters when it comes to the monetary implications of supply shocks.

The role of the tax system. As the monetary authority raises interest rates more strongly in the behavioral model after inflationary supply shocks, the spillovers to fiscal policy become larger. Therefore, the tax system or the design of fiscal policy becomes more important. We highlight this by considering a different tax system, in particular, a less-progressive one. While in the baseline case only the high productivity households pay taxes, now all households pay taxes proportional to their productivity (see Appendix F.5 for details).

We find that the increase in inequality in the full-inflation-stabilization case is much more pronounced in the behavioral model when taxes are less progressive. Furthermore, the increase in inequality is much more persistent. The reason is that more productive households tend to be less borrowing constrained and thus, adjust their consumption on impact in expectation of higher future taxes. Households at the borrowing constraint, however, reduce their consumption once taxes actually increase. As these households tend to consume relatively little, consumption inequality increases over time as taxes increase. Appendix F.5 shows these results graphically.

Heterogeneous $\bar{m}$. When we incorporate heterogeneous $\bar{m}$ along the lines described in Section 4.2, we find that the trade-off is slightly weaker compared to our baseline while still significantly stronger than in the rational HANK model (see Appendix F.4). The reason is that high-productivity households tend to be more likely to directly respond to monetary policy and as these households are now closer to rational expectations, they respond more strongly to expected future higher interest rates. Therefore, monetary policy has to react slightly less than in the case with homogeneous degrees of cognitive discounting. In the case of a Taylor rule, the amplification channel is also slightly less strong. Quantitatively, however, the differences are tiny and the trade-off that arises due to households' cognitive discounting remains substantial.

Furthermore, heterogeneity in cognitive discounting affects the role of the fiscal regime. In particular, as more productive households are less behavioral, the future expected increase in taxes is almost fully accounted for by them. Less productive households, on the other hand, cognitively discount these future expected tax increases more strongly and thus, respond less to them on impact. Under a more progressive tax system, the first channel is more important and hence, the initial effect on inequality is smaller in the case when monetary policy fully stabilizes inflation and taxes are progressive. The reason is that more productive households—who tend to consume more—decrease their consumption more strongly on impact as they expect future higher taxes.

Cost-push shocks. So far, we have focused on the inflationary pressure coming from negative TFP shocks. We show in Appendix F.7 that if the inflationary pressure comes from a cost-push shock instead, the monetary and fiscal implications are very similar: the central bank needs to raise interest rates much more strongly in the behavioral HANK model than in the rational HANK model to fully stabilize inflation. This pushes up the government debt level, especially in the behavioral HANK model. If monetary policy instead follows a Taylor rule, again inflation is much more sensitive to the supply shock in the behavioral HANK model than in the rational HANK model.

6 Model Extensions

We now extend the tractable model along three dimensions to show how the interaction of household heterogeneity and bounded rationality helps to match further empirical facts. First, we show that the behavioral HANK model matches the empirical estimates of intertemporal marginal propensities to consume (iMPCs) and how they depend on bounded rationality, heterogeneity and the interaction of the two. Second, we allow for sticky wages and show how the interplay of sticky wages, household heterogeneity and bounded rationality leads to hump-shaped responses of macroeconomic variables in response to aggregate shocks as well as forecast-error dynamics consistent with recent findings from survey data. Third, we derive an equivalence result between HANK models with bounded rationality and HANK models with incomplete information and learning.

6.1 Intertemporal MPCs

The HANK literature shows that intertemporal marginal propensities to consume are a key statistic for conducting policy analysis (see, e.g., Auclert et al. (2018), Auclert et al. (2020), and Kaplan and Violante (2020)).[20] We follow the tractable HANK literature and define the aggregate iMPCs as the partial derivative of aggregate consumption at time k, $\widehat{c}_k$, with respect to aggregate disposable income, $\tilde{y}_0$, keeping everything else fixed (see Bilbiie (2021), Cantore and Freund (2021)). The

[20]See, e.g., Lian (2021) or Boutros (2022) for MPC analyses in models deviating from FIRE.

Figure 6: Intertemporal MPCs, Bounded Rationality and Household Heterogeneity

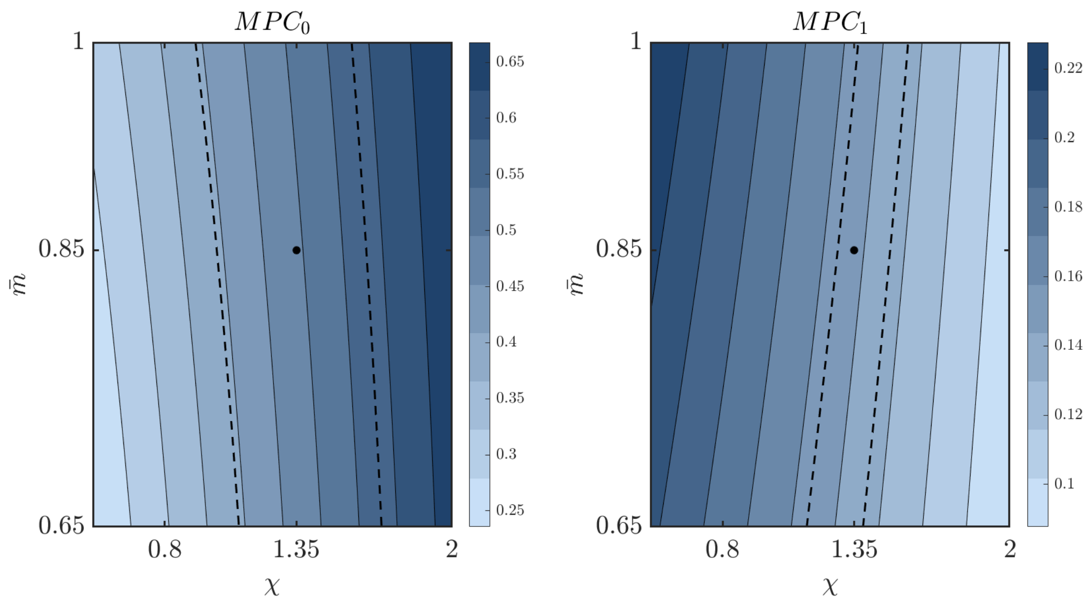

Note: This figure shows the aggregate intertemporal MPCs in our tractable model, i.e., the aggregate consumption response in year 0 (left) and year 1 (right) to a change in aggregate disposable income in year 0 for different χ (horizontal axis) and $\bar{m}$ (vertical axis). The dashed lines show the range of empirically-estimated iMPCs and the black dot shows the model estimate given our baseline calibration. Darker colors represent higher MPCs (see the colorbars on the right side of the figures).

following Proposition characterizes the iMPCs in the behavioral HANK model (see Appendix G for the derivation).

Proposition 4. *The intertemporal MPCs in the behavioral HANK model, i.e., the aggregate consumption response in period k to a one-time change in aggregate disposable income in period 0, are given by*

$$MPC_0 \equiv \frac{d\widehat{c}_0}{d\widetilde{y}_0} = 1 - \frac{1-\lambda\chi}{s\bar{m}}\mu_2^{-1}$$

$$MPC_k \equiv \frac{d\widehat{c}_k}{d\widetilde{y}_0} = \frac{1-\lambda\chi}{s\bar{m}}\mu_2^{-1}\left(\beta^{-1}-\mu_1\right)\mu_1^{k-1}, \qquad \text{for } k>0,$$

where the parameters μ_1 and μ_2 depend on the underlying parameters, including $\bar{m}$ and χ and are explicitly spelled out in Appendix G.

We calibrate the model annually as most of the empirical evidence on iMPCs is annual (see Fagereng et al. (2021) and Auclert et al. (2018)). We set $s = 0.8$ and $\beta = 0.95$, and keep the rest of the calibration as in Section 3. Figure 6 graphically depicts how the interplay of bounded rationality $\bar{m}$ and household heterogeneity χ determines the size of the aggregate iMPCs. The left panel depicts the aggregate MPCs within the first year (in period 0) and the right panel the aggregate MPCs within the second year (in period 1). Darker colors represent higher MPCs. First, note that with our baseline calibration—$\chi = 1.35$ and $\bar{m} = 0.85$ as shown by the black dots—the behavioral HANK model generates iMPCs within the first year of 0.52 and within the second year of 0.15. These values lie within the estimated bounds for the iMPCs in the data (Auclert et al. (2018)) which are between $0.42 - 0.6$ for the first and $0.14 - 0.16$ for the second year (see dashed lines). Away from our baseline calibration, an increase in χ increases the MPCs in the first year but decreases them in the second year. In contrast, an increase in $\bar{m}$ increases the aggregate MPC in the first year and in the second year.

Let us first turn to the role of χ for the iMPCs. Recall, the higher χ the more sensitive is the income of the H households to a change in aggregate income. Thus, with higher χ, H households gain weight in relative terms for the aggregate iMPCs while unconstrained households loose weight in relative terms. This pushes up the aggregate MPC within the first year as the H households spend all of their income windfall, but pushes down the aggregate MPC within the second year as households that were hand-to-mouth in the period of the income windfall have a MPC of 0 in the second year.

Bounded rationality, captured by $\bar{m}$, affects only the MPCs of the initially-unconstrained households as these are the only households who intertemporarily optimize. Their Euler equation dictates that the decrease in today's marginal utility of consumption—due to the increase in consumption—is equalized by a decrease in tomorrow's expected marginal utility. For behavioral households, however, the decrease in tomorrow's marginal utility needs to be more substantial as they cognitively discount it. Hence, behavioral households save relatively more out of the income windfall. This pushes down the aggregate MPCs in $t = 0$. The same is true for the aggregate MPC in $t = 1$, in which there are now two opposing forces at work: on the one hand, unconstrained households again cognitively discount the expectations about the future decrease in their marginal utility which depresses their consumption. On the other hand, unconstrained households have accumulated more wealth from period $t = 0$ which tends to increase consumption. Given our calibration, the former dominates in $t = 1$. Figure 29 in Appendix G shows that, beginning in $k = 3$, the latter effect starts to dominate. For a higher idiosyncratic risk of becoming hand-to-mouth, i.e., an increase in the transition probability $1 - s$, the aggregate MPC is already higher in $t = 1$ for lower $\bar{m}$. The reason is that a smaller fraction of initially-unconstrained households remains unconstrained which pushes consumption upwards in $k = 1$ (see Figure 30 in Appendix G).

The effects of a change in $\bar{m}$ are more pronounced at lower levels of χ. This follows directly from our discussion about the role of χ and $\bar{m}$: the lower χ, the higher is the relative importance of unconstrained households for the aggregate iMPCs and, in turn, the stronger is the effect of $\bar{m}$ on the aggregate iMPCs. These interaction effects are quite substantial: at $\chi = 1.35$, a decrease of $\bar{m}$ from 1 to 0.65 decreases the MPC_0 by 8% and the MPC_1 by more than 11%.

6.2 Sticky Wages

Recent HANK models have relaxed the assumption of fully-flexible wages and rather assume wages to be sticky, bringing these models closer to the data (see, e.g., Auclert et al. (2020) or Broer et al. (2020)). To introduce sticky wages, we follow Colciago (2011) and assume a centralized labor market in which a labor union allocates the hours of households to firms and makes sure that U and H households work the same amount. The labor union faces the typical Calvo (1983) friction, such that it can re-optimize the wage within a given period only with a certain probability, giving rise to a wage Phillips Curve. We assume that the labor union sets wages based on rational expectations to focus on the effects of bounded rationality solely on the household side.

The wage Phillips curve is given by

$$\pi_t^w = \beta \mathbb{E}_t \pi_{t+1}^w + \kappa_w \widehat{\mu}_t^w,$$

where π_t^w denotes wage inflation, κ_w the slope of the wage Phillips curve and $\widehat{\mu}_t^w$ is a time-varying wage markup, given by $\widehat{\mu}_t^w = \gamma \widehat{c}_t + \varphi \widehat{n}_t - \widehat{w}_t$. We set $\kappa_w = 0.075$ as in Bilbiie et al. (2021).

We follow Auclert et al. (2020) and introduce interest-rate smoothing in the Taylor rule: $\widehat{i}_t = \rho_i \widehat{i}_{t-1} + (1-\rho_i)\phi\pi_t + \varepsilon_t^{MP}$ with $\rho_i = 0.89$ and $\phi = 1.5$ as estimated by Auclert et al. (2020) and assume the shocks ε_t^{MP} to be completely transitory. Similar to the wage setters, we assume price-setting firm managers to be fully rational, giving rise to the standard New Keynesian Phillips curve

$$\pi_t = \beta \mathbb{E}_t \pi_{t+1} + \kappa_\pi \widehat{mc}_t,$$

where $\widehat{mc}_t$ denotes the time-varying price markup. The rest of the model is as above. We relegate the details and the parameterization to Appendix H.

Hump-shaped responses to monetary policy shocks. Figure 7 shows the impulse-response functions of output, real wages and consumption of the two household types to a monetary policy shock for the behavioral HANK model with sticky wages (blue-dashed lines). Importantly, the figure shows that the output response to a monetary policy shock is hump-shaped in the behavioral HANK model but neither in its rational counterpart (orange-dotted lines) nor in its representative-agent counterpart (black-solid lines show the rational RANK results and the black-dashed-dotted lines the results for the behavioral RANK model).

Wage rigidity leads to a hump-shaped response in real wages, which is the case in all four models. Since wages determine the H households' income in the rational and the behavioral HANK, their consumption also follows a hump-shape (see lower right figure). Crucial for the overall response, however, is not only the response of H households but also the response of unconstrained households.

Under rational expectations, unconstrained households perfectly understand how the consumption of H agents will respond in the future and what this implies for their idiosyncratic risk induced by type switching. In particular, they understand already on impact that their self-insurance mo-

Figure 7: Monetary Policy Shock with Sticky Wages

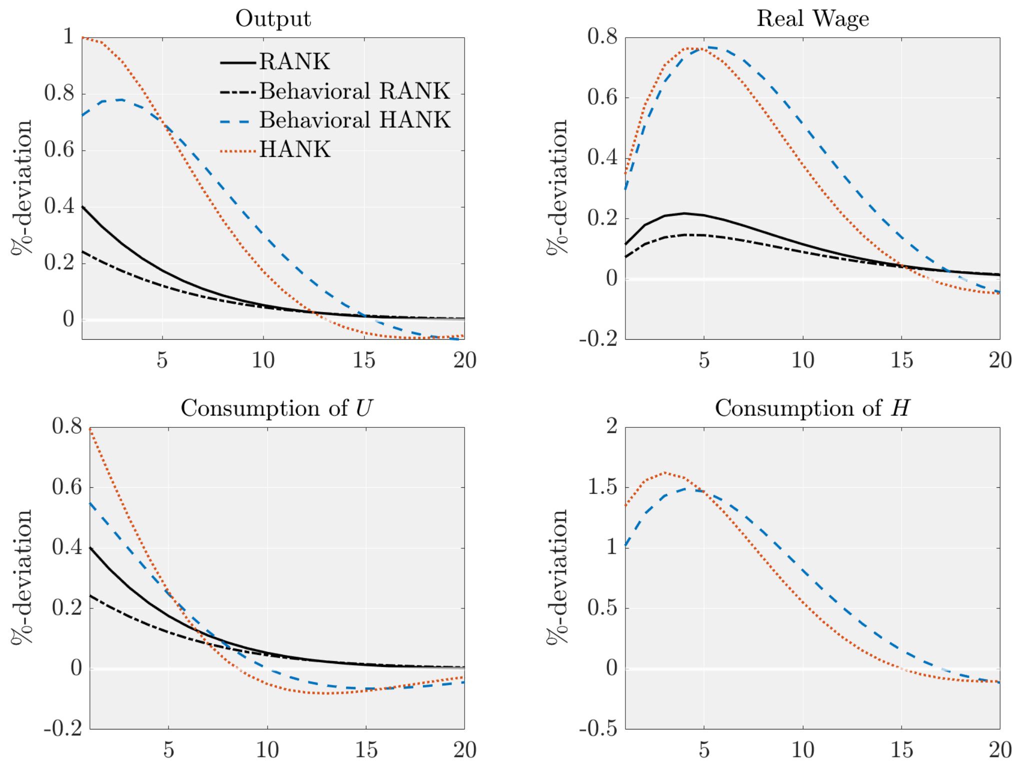

Note: This figure shows the impulse-response functions of output, real wages and consumption of the two household types to a monetary policy shock in the tractable behavioral HANK model, the rational and the behavioral RANK model and the rational HANK model with sticky wages. The shock size is normalized such that output in the rational HANK model increases by 1pp on impact.

tive will be relaxed for some periods. Thus, unconstrained households immediately cut back on precautionary savings and, thus, their consumption responds strongly on impact. Under bounded rationality, however, unconstrained households cognitively discount the future and underreact to the expected increase in wages and, thus, the relaxation of their idiosyncratic risk. Hence, on impact, they do not cut back on precautionary savings as strong as a rational household would. Going forward, they learn that their self-insurance motive is still (or even more) relaxed. As a consequence, their consumption decreases more slowly inducing a flatter consumption profile compared to a rational unconstrained household. It is the combination of the flatter consumption profile of unconstrained households and the hump-shaped consumption profile of the hand-to-mouth that generates the hump-shaped response of consumption in the aggregate.

The model with a representative (rational or behavioral) agent does not generate the hump-shaped response. The reason is that without hand-to-mouth agents, the wage profile does not translate into hump-shaped consumption of (a sub population of) households to begin with. It is thus indeed the *interaction* of household heterogeneity and bounded rationality that produces these hump-shaped responses.

Auclert et al. (2020) argue that many macroeconomic models fail to generate the *micro jumps and macro humps* that we observe in the data, i.e., iMPCs that respond strongly on impact and hump-shaped responses of macroeconomic variables to aggregate shocks. Our results on iMPCs in Section 6.1 as well as the results presented in Figure 7 show how the behavioral HANK model offers a tractable analogue to the full-blown HANK model presented in Auclert et al. (2020).

Forecast-error dynamics. We now show that the sticky-wage behavioral HANK model generates dynamic forecast errors as observed in survey data. In particular, households' expectations initially underreact followed by delayed overreaction as recently documented empirically in Angeletos et al. (2021) for unemployment and inflation and in Adam et al. (2022) for housing prices. Consistent with the empirical exercise in Angeletos et al. (2021), we focus on three-quarter ahead forecasts. For a variable $\widehat{x}$, the three-period ahead forecast error is defined as

$$FE^{\widehat{x}}_{t+h+3|t+h} \equiv \widehat{x}_{t+h+3} - \bar{m}^3 \mathbb{E}_{t+h}\left[\widehat{x}_{t+h+3}\right],$$

such that a positive forecast error means the forecast was lower than the actual outcome.

Figure 8 shows the forecast errors of output, the real wage and consumption of the two household types starting in the first period in which the expectations start to change which in this case corresponds to the fourth period after the shock. For completeness, the orange-dashed lines at zero show that under rational expectations, i.e., $\bar{m} = 1$, forecast errors are equal to 0. In the behavioral HANK model, however, this is not the case. In fact, forecast errors are positive in the first few quarters after the shock, illustrating the underreaction of the agents' expectations to the shock.

After about 10-15 quarters, however, forecast errors turn negative. Put differently, the behavioral agents' expectations show patterns of delayed overreaction. In contrast to Angeletos et al. (2021) or Adam et al. (2022), the behavioral HANK model with sticky wages generates these dynamic patterns of forecast errors even though the behavioral agents' expectations are purely forward looking.

Where does the delayed overshooting come from? As figure 7 shows, output falls below its steady-state level after some periods in the HANK models. The reason is that with sticky wages, wages increase very persistently. In HANK, this makes the consumption of the H households very persistent which, ceteris paribus, makes the increase in aggregate demand more persistent. Monetary policy reacts to this by increasing the nominal interest rate more strongly and more persistently. Due to inertia in the Taylor rule, however, the interest rate stays high even as aggregate demand returns to its steady state level, generating a mild recession after about 15 quarters (consistent with larger HANK models, see, for example, Auclert et al. (2020)). The behavioral agents then not only underestimate the boom after the monetary policy shock in the short-run, but also underestimate the mild recession in the medium-run, which causes the delayed overshooting in their expectations.

Note that the behavioral RANK model (black-solid lines) does not generate these delayed overreactions. Only when allowing for both—household heterogeneity and bounded rationality—the model is able to generate hump-shaped responses of macroeconomic aggregates and forecast

Figure 8: Forecast Errors with Sticky Wages in Response to a Monetary Policy Shock

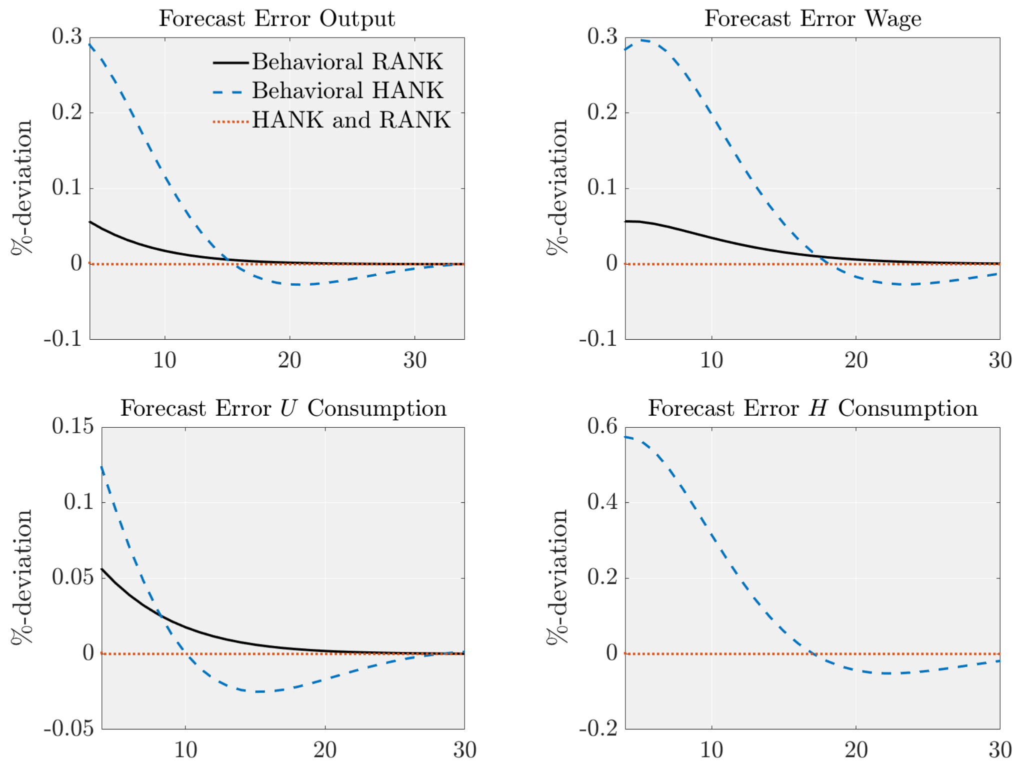

Note: This figure shows the forecast error dynamics of output, the real wage, consumption of unconstrained households and of hand-to-mouth households after an expansionary monetary policy shock in the tractable model.

error dynamics that are consistent with recent evidence from household survey expectations.

6.3 Bounded Rationality and Incomplete Information with Learning: An Equivalence Result

In this section, we derive an equivalence result of heterogeneous-household models featuring bounded rationality and those featuring incomplete information with learning. In particular, we show how a change in the default value in the behavioral setup leads to an observationally equivalent IS equation as in models with incomplete information and learning (see Angeletos and Huo (2021) and Gallegos (2021)). To this end, we now assume that behavioral agents anchor their expectations to their *last observation* instead of the steady state values which induces a backward-looking component in the expectations as well as in the IS equation:

Proposition 5. *Set the boundedly-rational agents' default value to the variable's past value $X_t^d = X_{t-1}$. In this case, the boundedly-rational agent's expectations of X_{t+1} becomes*

$$\mathbb{E}_t^{BR}\left[X_{t+1}\right] = (1-\bar{m})X_{t-1} + \bar{m}\mathbb{E}_t\left[X_{t+1}\right] \tag{22}$$

and the behavioral HANK IS equation is then given by

$$\widehat{y}_t = \psi_f \mathbb{E}_t \widehat{y}_{t+1} - \psi_c \frac{1}{\gamma}\left(\widehat{i}_t - \mathbb{E}_t \pi_{t+1}\right) + (1-\bar{m})\delta \widehat{y}_{t-1}. \tag{23}$$

Proposition 5 shows that the change in the agents' default value does not change the existing behavioral and heterogeneity coefficients ψ_f and ψ_c. Yet, anchoring to past realizations introduces an additional backward-looking term in the IS equation, similar to models relying on habit persistence. The IS equation thus features *myopia* and *anchoring* as in Angeletos and Huo (2021) and Gallegos (2021) who derive an IS equation with the same reduced form. Their setup, however, is based on incomplete-information and learning. We complement their findings by showing how we can generate the equivalent outcome based on a *behavioral* relaxation of FIRE.

7 Conclusion

In this paper, we develop a new framework for business-cycle and policy analysis: the behavioral HANK model. To arrive at this framework, we introduce bounded rationality in the form of cognitive discounting and household heterogeneity into a sticky price model. The model can account for recent empirical findings on the transmission mechanisms of monetary policy. In particular, households with higher marginal propensities to consume tend to be more exposed to changes in aggregate income that are induced by monetary policy, leading to an amplification of conventional monetary policy through indirect effects. Simultaneously, the model rules out the forward guidance puzzle and remains stable at the effective lower bound. The model thus overcomes a tension in existing models with household heterogeneity: when accounting for the underlying heterogeneity, these models tend to aggravate the forward guidance puzzle and the instability issues at the lower bound. Both, bounded rationality and household heterogeneity, are crucial to arrive at our results.

The behavioral HANK model predicts that central banks that want to stabilize inflation after an inflationary supply shock need to hike the nominal interest rate much more strongly than under rational expectations—even in rational models that do not suffer from the forward guidance puzzle. Hiking interest rates, however, leads to a more pronounced increase in public debt and inequality, especially when initial debt levels are already high. Finally, we uncover a novel amplification channel in response to negative supply shocks: when monetary policy follows a Taylor rule, household heterogeneity and cognitive discounting both generate a larger increase in inflation and the output gap and the two mutually reinforce each other.

Given its consistency with empirical facts about the transmission of monetary policy, the behavioral HANK model provides a natural laboratory for both business-cycle and policy analysis. Our framework can also easily be extended along many dimensions, some of which we have explored in the paper, whereas others are left for future work.

References

Acharya, S. and K. Dogra (2020): "Understanding HANK: Insights from a PRANK," *Econometrica*, 88, 1113–1158.

Adam, K., O. Pfaeuti, and T. Reinelt (2022): "Subjective Housing Price Expectations, Falling Natural Rates and the Optimal Inflation Target," .

Aiyagari, S. R. (1994): "Uninsured idiosyncratic risk and aggregate saving," *The Quarterly Journal of Economics*, 109, 659–684.

Ampudia, M., D. Georgarakos, J. Slacalek, O. Tristani, P. Vermeulen, and G. Violante (2018): "Monetary policy and household inequality," .

Andrade, P., G. Gaballo, E. Mengus, and B. Mojon (2019): "Forward guidance and heterogeneous beliefs," *American Economic Journal: Macroeconomics*, 11, 1–29.

Angeletos, G.-M., F. Collard, and H. Dellas (2020): "Business-cycle anatomy," *American Economic Review*, 110, 3030–70.

Angeletos, G.-M. and Z. Huo (2021): "Myopia and anchoring," *American Economic Review*, 111, 1166–1200.

Angeletos, G.-M., Z. Huo, and K. A. Sastry (2021): "Imperfect macroeconomic expectations: Evidence and theory," *NBER Macroeconomics Annual*, 35, 1–86.

Angeletos, G.-M. and C. Lian (2018): "Forward guidance without common knowledge," *American Economic Review*, 108, 2477–2512.

——— (2022): "Dampening General Equilibrium: Incomplete Information and Bounded Rationality," Tech. rep., National Bureau of Economic Research.

Auclert, A. (2019): "Monetary policy and the redistribution channel," *American Economic Review*, 109, 2333–67.

Auclert, A., M. Rognlie, and L. Straub (2018): "The intertemporal keynesian cross," Tech. rep., National Bureau of Economic Research.

——— (2020): "Micro jumps, macro humps: Monetary policy and business cycles in an estimated HANK model," Tech. rep., National Bureau of Economic Research.

Benchimol, J. and L. Bounader (2019): *Optimal monetary policy under bounded rationality*, International Monetary Fund.

Bewley, T. (1986): "Stationary monetary equilibrium with a continuum of independently fluctuating consumers," *Contributions to mathematical economics in honor of Gérard Debreu*, 79.

Bhandari, A., J. Borovička, and P. Ho (2019): "Survey data and subjective beliefs in business cycle models," *Available at SSRN 2763942*.

Bilbiie, F. O. (2008): "Limited asset markets participation, monetary policy and (inverted) aggregate demand logic," *Journal of economic theory*, 140, 162–196.

——— (2011): "Nonseparable preferences, frisch labor supply, and the consumption multiplier of government spending: One solution to a fiscal policy puzzle," *Journal of Money, Credit and*

Banking, 43, 221–251.

——— (2020): "The new Keynesian cross," *Journal of Monetary Economics*, 114, 90–108.

——— (2021): "Monetary policy and heterogeneity: An analytical framework," .

Bilbiie, F. O., D. Känzig, and P. Surico (2021): "Capital and income inequality: An aggregate-demand complementarity," .

Bonciani, D. and J. Oh (2021): "Optimal monetary policy mix at the zero lower bound," *Bank of England Working Paper.*

Born, B., Z. Enders, M. Menkhoff, G. J. Müller, and K. Niemann (2022): "Firm Expectations and News: Micro v Macro," .

Boutros, M. (2022): "Windfall Income Shocks with Finite Planning Horizons," .

Broer, T., N.-J. Harbo Hansen, P. Krusell, and E. Öberg (2020): "The New Keynesian transmission mechanism: A heterogeneous-agent perspective," *The Review of Economic Studies*, 87, 77–101.

Broer, T., A. Kohlhas, K. Mitman, and K. Schlafmann (2021a): "Information and Wealth Heterogeneity in the Macroeconomy," .

Broer, T., P. Krusell, and E. Öberg (2021b): "Fiscal Multipliers: A Heterogenous-Agent Perspective," Tech. rep., National Bureau of Economic Research.

Calvo, G. A. (1983): "Staggered prices in a utility-maximizing framework," *Journal of monetary Economics*, 12, 383–398.

Cantore, C. and L. B. Freund (2021): "Workers, capitalists, and the government: fiscal policy and income (re) distribution," *Journal of Monetary Economics*, 119, 58–74.

Carlson, J. A. and M. Parkin (1975): "Inflation expectations," *Economica*, 42, 123–138.

Chodorow-Reich, G. (2019): "Geographic cross-sectional fiscal spending multipliers: What have we learned?" *American Economic Journal: Economic Policy*, 11, 1–34.

Cochrane, J. H. (2018): "Michelson-Morley, Fisher, and Occam: The radical implications of stable quiet inflation at the zero bound," *NBER Macroeconomics Annual*, 32, 113–226.

Coibion, O. and Y. Gorodnichenko (2015): "Information rigidity and the expectations formation process: A simple framework and new facts," *American Economic Review*, 105, 2644–78.

Coibion, O., Y. Gorodnichenko, E. S. Knotek II, and R. Schoenle (2020): "Average inflation targeting and household expectations," Tech. rep., National Bureau of Economic Research.

Coibion, O., Y. Gorodnichenko, L. Kueng, and J. Silvia (2017): "Innocent Bystanders? Monetary policy and inequality," *Journal of Monetary Economics*, 88, 70–89.

Colciago, A. (2011): "Rule-of-thumb consumers meet sticky wages," *Journal of money, credit and banking*, 43, 325–353.

Debortoli, D. and J. Galí (2018): "Monetary policy with heterogeneous agents: Insights from TANK models," *Manuscript, September.*

Debortoli, D., J. Galí, and L. Gambetti (2020): "On the empirical (ir) relevance of the zero lower bound constraint," *NBER Macroeconomics Annual*, 34, 141–170.

Del Negro, M., M. Giannoni, and C. Patterson (2015): "The forward guidance puzzle," .

di Giovanni, J., Kalemli-Özcan, A. Silva, and M. A. Yildirim (2022): "Global Supply Chain Pressures, International Trade, and Inflation," Tech. rep., National Bureau of Economic Research.

Dupor, B., M. Karabarbounis, M. Kudlyak, and M. Mehkari (2021): "Regional consumption responses and the aggregate fiscal multiplier," .

D'Acunto, F., D. Hoang, and M. Weber (2020): "Managing households' expectations with unconventional policies," Tech. rep., National Bureau of Economic Research.

Erceg, C. J., D. W. Henderson, and A. T. Levin (2000): "Optimal monetary policy with staggered wage and price contracts," *Journal of monetary Economics*, 46, 281–313.

Fagereng, A., M. B. Holm, and G. J. Natvik (2021): "MPC heterogeneity and household balance sheets," *American Economic Journal: Macroeconomics*, 13, 1–54.

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

Ferriere, A. and G. Navarro (2018): "The Heterogeneous Effects of Government Spending: It's All About Taxes," *FRB International Finance Discussion Paper*.

Fuhrer, J. C. and G. D. Rudebusch (2004): "Estimating the Euler equation for output," *Journal of Monetary Economics*, 51, 1133–1153.

Gabaix, X. (2014): "A sparsity-based model of bounded rationality," *The Quarterly Journal of Economics*, 129, 1661–1710.

——— (2016): "Behavioral macroeconomics via sparse dynamic programming," Tech. rep., National Bureau of Economic Research.

——— (2019): "Behavioral inattention," in *Handbook of Behavioral Economics: Applications and Foundations 1*, Elsevier, vol. 2, 261–343.

——— (2020): "A behavioral New Keynesian model," *American Economic Review*, 110, 2271–2327.

Galí, J. (2015): *Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications*, Princeton University Press.

Galí, J., J. D. López-Salido, and J. Vallés (2007): "Understanding the effects of government spending on consumption," *Journal of the european economic association*, 5, 227–270.

Gallegos, J.-E. (2021): "HANK beyond FIRE," .

Hagedorn, M. (2018): "Prices and inflation when government bonds are net wealth," .

Hagedorn, M., J. Luo, I. Manovskii, and K. Mitman (2019): "Forward guidance," *Journal of Monetary Economics*, 102, 1–23.

Hagedorn, M. et al. (2016): *A demand theory of the price level*, Centre for Economic Policy Research.

Holm, M. B., P. Paul, and A. Tischbirek (2021): "The transmission of monetary policy under the microscope," *Journal of Political Economy*, 129, 2861–2904.

Huggett, M. (1993): "The risk-free rate in heterogeneous-agent incomplete-insurance economies," *Journal of economic Dynamics and Control*, 17, 953–969.

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

Kaplan, G. and G. L. Violante (2020): "The Marginal Propensity to Consume in Heterogeneous Agents Models," Tech. rep., Working Paper. Princeton University.

Krusell, P., T. Mukoyama, and A. A. Smith Jr (2011): "Asset prices in a Huggett economy," *Journal of Economic Theory*, 146, 812–844.

Kuhn, M., M. Schularick, and U. I. Steins (2020): "Income and wealth inequality in America, 1949–2016," *Journal of Political Economy*, 128, 3469–3519.

Laibson, D., P. Maxted, and B. Moll (2021): "Present bias amplifies the household balance-sheet channels of macroeconomic policy," Tech. rep., National Bureau of Economic Research.

Lian, C. (2021): "Mistakes in future consumption, high MPCs now," Tech. rep., National Bureau of Economic Research.

Luetticke, R. (2021): "Transmission of monetary policy with heterogeneity in household portfolios," *American Economic Journal: Macroeconomics*, 13, 1–25.

Mankiw, N. G. (2000): "The savers-spenders theory of fiscal policy," *American economic review*, 90, 120–125.

McKay, A., E. Nakamura, and J. Steinsson (2016): "The power of forward guidance revisited," *The American Economic Review*, 106, 3133–3158.

——— (2017): "The discounted euler equation: A note," *Economica*, 84, 820–831.

McKay, A. and J. F. Wieland (2022): "Forward guidance and durable goods demand," *American Economic Review: Insights*, 4, 106–22.

Michaillat, P. and E. Saez (2021): "Resolving New Keynesian anomalies with wealth in the utility function," *Review of Economics and Statistics*, 103, 197–215.

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

Nakamura, E. and J. Steinsson (2014): "Fiscal stimulus in a monetary union: Evidence from US regions," *American Economic Review*, 104, 753–92.

Patterson, C. (2019): "The matching multiplier and the amplification of recessions," *Unpublished Manuscript, Northwestern University.*

Ravn, M. O. and V. Sterk (2017): "Job uncertainty and deep recessions," *Journal of Monetary Economics*, 90, 125–141.

——— (2021): "Macroeconomic fluctuations with HANK & SAM: An analytical approach," *Journal of the European Economic Association*, 19, 1162–1202.

Roth, C., M. Wiederholt, and J. Wohlfart (2021): "The Effects of Forward Guidance: Theory with Measured Expectations," Tech. rep.

Samarina, A. and A. D. Nguyen (2019): "Does monetary policy affect income inequality in the euro area?" .

Slacalek, J., O. Tristani, and G. L. Violante (2020): "Household balance sheet channels of monetary policy: A back of the envelope calculation for the euro area," *Journal of Economic Dynamics and Control*, 115, 103879.

Werning, I. (2015): "Incomplete markets and aggregate demand," 2015 Meeting Papers 932, Society for Economic Dynamics.

Wiederholt, M. (2015): "Empirical properties of inflation expectations and the zero lower bound," *manuscript, Goethe University.*

Woodford, M. (2003): "Interest and prices," .

——— (2011): "Simple analytics of the government expenditure multiplier," *American Economic Journal: Macroeconomics*, 3, 1–35.

——— (2019): "Monetary policy analysis when planning horizons are finite," *NBER Macroeconomics Annual*, 33, 1–50.

Appendix

A Model Details and Derivations

A.1 Derivation of χ

In Section 2, we stated that

$$\widehat{c}_t^H = \chi \widehat{y}_t, \tag{24}$$

where $\chi \equiv 1 + \varphi \left(1 - \frac{\mu^D}{\lambda}\right)$ is *the* crucial statistic coming from the limited heterogeneity setup. We now show how we arrive at equation (24) from the H-households' budget constraint, optimality conditions and market clearing.

The labor-leisure condition of the H households is given by

$$(N_t^H)^{\varphi} = W_t (C_t^H)^{-\gamma}, \tag{25}$$

and similarly for the U households. As we focus on the steady state with no inequality, we have that in steady state $C = C^H = C^U$ and $N = N^U = N^H$ and market clearing and the production function imply $Y = C = N$, which we normalize to 1.

Log-linearizing the labor-leisure conditions yields

$$\varphi \widehat{n}_t^H = \widehat{w}_t - \gamma \widehat{c}_t^H$$
$$\varphi \widehat{n}_t^U = \widehat{w}_t - \gamma \widehat{c}_t^U.$$

Since both households work for the same wage, we obtain

$$\varphi \widehat{n}_t^H + \gamma \widehat{c}_t^H = \varphi \widehat{n}_t^U + \gamma \widehat{c}_t^U \tag{26}$$

Log-linearizing the market clearing conditions yields

$$\widehat{n}_t = \lambda \widehat{n}_t^H + (1-\lambda) \widehat{n}_t^U$$
$$\widehat{c}_t = \lambda \widehat{c}_t^H + (1-\lambda) \widehat{c}_t^U,$$

which can be re-arranged as (using $\widehat{y}_t = \widehat{c}_t = \widehat{n}_t$)

$$\widehat{n}_t^U = \frac{1}{1-\lambda}\left(\widehat{y}_t - \lambda \widehat{n}_t^H\right)$$
$$\widehat{c}_t^U = \frac{1}{1-\lambda}\left(\widehat{y}_t - \lambda \widehat{c}_t^H\right).$$

Replacing $\widehat{n}_t^U$ and $\widehat{c}_t^U$ in equation (26) then gives

$$\varphi \widehat{n}_t^H + \gamma \widehat{c}_t^H = (\varphi + \gamma) \widehat{y}_t. \tag{27}$$

The budget constraint of H households (accounting for the fact that bond holdings are zero in equilibrium) is given by

$$C_t^H = W_t N_t^H + \frac{\mu^D}{\lambda} D_t. \tag{28}$$

In log-linearized terms, we get

$$\widehat{c}_t^H = \widehat{w}_t + \widehat{n}_t^H + \frac{\mu^D}{\lambda}\widehat{d}_t, \tag{29}$$

and using that $\widehat{w}_t = -\widehat{d}_t = \varphi\widehat{n}_t^H + \gamma\widehat{c}_t^H$, we get

$$\widehat{c}_t^H = \left(\varphi\widehat{n}_t^H + \gamma\widehat{c}_t^H\right)\left(1 - \frac{\mu^D}{\lambda}\right) + \widehat{n}_t^H. \tag{30}$$

Using (27) to solve for $\widehat{n}_t^H$ and plugging it into (30) yields

$$\widehat{c}_t^H = \widehat{c}_t^H\gamma\left(1 - \frac{\mu^D}{\lambda}\right) + \chi\left(\frac{\varphi+\gamma}{\varphi}\widehat{y}_t - \frac{\gamma}{\varphi}\widehat{c}_t^H\right).$$

Grouping terms, we obtain

$$\widehat{c}_t^H = \chi\widehat{y}_t,$$

with $\chi \equiv 1 + \varphi\left(1 - \frac{\mu^D}{\lambda}\right)$, as stated above.

A.2 Derivation of Proposition 1.

Combining equations (8) and (10) with the bounded-rationality setup in equation (6) for $\widehat{x}_t^d = 0$ as X_t^d is given by the steady state, we have

$$\mathbb{E}_t^{BR}\left[\widehat{c}_{t+1}^H\right] = \bar{m}\mathbb{E}_t\left[\widehat{c}_{t+1}^H\right] = \bar{m}\chi\mathbb{E}_t\left[\widehat{y}_{t+1}\right]$$

$$\mathbb{E}_t^{BR}\left[\widehat{c}_{t+1}^U\right] = \bar{m}\mathbb{E}_t\left[\widehat{c}_{t+1}^U\right] = \bar{m}\frac{1-\lambda\chi}{1-\lambda}\mathbb{E}_t\left[\widehat{y}_{t+1}\right].$$

Plugging these two equations as well as equation (10) into the Euler equation of unconstrained households (12) yields

$$\frac{1-\lambda\chi}{1-\lambda}\widehat{y}_t = s\bar{m}\frac{1-\lambda\chi}{1-\lambda}\mathbb{E}_t\left[\widehat{y}_{t+1}\right] + (1-s)\bar{m}\chi\mathbb{E}_t\left[\widehat{y}_{t+1}\right] - \frac{1}{\gamma}\left(\widehat{i}_t - \mathbb{E}_t\pi_{t+1}\right).$$

Combining the $\mathbb{E}_t\left[\widehat{y}_{t+1}\right]$ terms and dividing by $\frac{1-\lambda\chi}{1-\lambda}$ yields the following coefficient in front of $\mathbb{E}_t\left[\widehat{y}_{t+1}\right]$:

$$\begin{aligned}
\psi_f &\equiv \bar{m}\left[s + (1-s)\chi\frac{1-\lambda}{1-\lambda\chi}\right] \\
&= \bar{m}\left[1 - 1 + s + (1-s)\chi\frac{1-\lambda}{1-\lambda\chi}\right] \\
&= \bar{m}\left[1 - \frac{1-\lambda\chi}{1-\lambda\chi} + s + (1-s)\chi\frac{1-\lambda}{1-\lambda\chi}\right] \\
&= \bar{m}\left[1 - \frac{1-\lambda\chi}{1-\lambda\chi} + \frac{(1-\lambda\chi)s}{1-\lambda\chi} + (1-s)\chi\frac{1-\lambda}{1-\lambda\chi}\right] \\
&= \bar{m}\left[1 + (\chi - 1)\frac{1-s}{1-\lambda\chi}\right].
\end{aligned}$$

Defining $\psi_c \equiv \frac{1-\lambda}{1-\lambda\chi}$ yields the behavioral HANK IS equation in Proposition 1:

$$\widehat{y}_t = \psi_f\mathbb{E}_t\widehat{y}_{t+1} - \psi_c\frac{1}{\gamma}\left(\widehat{i}_t - \mathbb{E}_t\pi_{t+1}\right).$$

A.3 Derivation of Proposition 2.

The first part comes from the fact that amplification is obtained when

$$\psi_c = \frac{1-\lambda}{1-\lambda\chi} > 1,$$

which requires $\chi > 1$, given that we assume throughout $\chi\lambda < 1$.

For the second part, recall how we define the forward guidance experiment (following Bilbiie (2021)). We assume a Taylor coefficient of 0, i.e., $\phi = 0$, such that the nominal interest rate is given by $\widehat{i}_t = \varepsilon_t^{MP}$. Replacing inflation using the Phillips curve (13), i.e., $\pi_t = \kappa\widehat{y}_t$, we can re-write the behavioral HANK IS equation from Proposition 1 as

$$\begin{aligned}\widehat{y}_t &= \psi_f \mathbb{E}_t \widehat{y}_{t+1} - \psi_c \frac{1}{\gamma}\left(\varepsilon_t^{MP} - \kappa \mathbb{E}_t \widehat{y}_{t+1}\right) \\ &= \left(\psi_f + \psi_c \frac{1}{\gamma}\kappa\right)\mathbb{E}_t \widehat{y}_{t+1} - \psi_c \frac{1}{\gamma}\varepsilon_t^{MP}\end{aligned}$$

The forward guidance puzzle is ruled out if and only if

$$\left(\psi_f + \psi_c \frac{1}{\gamma}\kappa\right) < 1,$$

which is the same as the condition stated in Proposition 2:

$$\bar{m}\delta + \frac{1}{\gamma}\frac{1-\lambda}{1-\lambda\chi}\kappa < 1.$$

Solving this for $\bar{m}$ yields

$$\bar{m} < \frac{1 - \frac{1-\lambda}{\gamma(1-\lambda\chi)}\kappa}{\delta},$$

which completes Proposition 2.

A.4 Derivation of Proposition 3.

Replacing $\widehat{i}_t$ by $\phi\pi_t = \phi\kappa\widehat{y}_t$ and $\mathbb{E}_t\pi_{t+1} = \kappa\mathbb{E}_t\widehat{y}_{t+1}$ in the IS equation (14), we get

$$\widehat{y}_t = \psi_f \mathbb{E}_t \widehat{y}_{t+1} - \psi_c \frac{1}{\gamma}\left(\phi\kappa\widehat{y}_t - \kappa \mathbb{E}_t \widehat{y}_{t+1}\right),$$

which can be re-written as

$$\widehat{y}_t\left(1 + \psi_c \frac{1}{\gamma}\phi\kappa\right) = \mathbb{E}_t \widehat{y}_{t+1}\left(\psi_f + \psi_c \frac{1}{\gamma}\kappa\right).$$

Dividing by $\left(1 + \psi_c \frac{1}{\gamma}\phi\kappa\right)$ and plugging in for ψ_f and ψ_c yields

$$\widehat{y}_t = \frac{\bar{m}\delta + \frac{(1-\lambda)\kappa}{\gamma(1-\lambda\chi)}}{1 + \kappa\phi\frac{1}{\gamma}\frac{(1-\lambda)}{1-\lambda\chi}}\mathbb{E}_t \widehat{y}_{t+1}.$$

To obtain determinacy, the term in front of $\mathbb{E}_t\widehat{y}_{t+1}$ has to be smaller than 1. Solving this for ϕ yields

$$\phi > \phi^* = 1 + \frac{\bar{m}\delta - 1}{\frac{\kappa}{\gamma}\frac{1-\lambda}{1-\lambda\chi}}, \tag{31}$$

which is the condition in Proposition 3. This illustrates how bounded rationality raises the likelihood that the Taylor principle ($\phi^* = 1$) is sufficient for determinacy, as the Taylor principle can only hold if

$$\bar{m}\delta \leq 1.$$

In the rational model, this boils down to $\delta \leq 1$. However, the Taylor principle can be sufficient under bounded rationality, i.e., $\bar{m} < 1$, even when $\delta > 1$, thus, even when allowing for amplification. Note that we could also express condition (31) as

$$\phi > \phi^* = 1 + \frac{\psi_f - 1}{\frac{\kappa}{\gamma}\psi_c}.$$

Proposition 3 can be extended to allow for Taylor rules of the form

$$\widehat{i}_t = \phi_\pi \pi_t + \phi_y \widehat{y}_t$$

and in which the behavioral agents do not have rational expectations about the real interest rate but rather perceive the real interest rate to be equal to

$$\widehat{r}_t^{BR} \equiv \widehat{i}_t - \bar{m}^r \mathbb{E}_t \pi_{t+1},$$

where $\bar{m}^r$ can be equal to $\bar{m}$ or can potentially differ from it (if it equals 1, we are back to the case in which the behavioral agent is rational with respect to real interest rates).

Combining the static Phillips Curve with the generalized Taylor rule and the behavioral HANK IS equation, it follows that

$$\widehat{y}_t = \frac{\omega_f + \frac{\kappa}{\gamma}\omega_c \bar{m}^r}{1 + \frac{\omega_c}{\gamma}\left(\kappa\phi_\pi + \phi_y\right)} \mathbb{E}_t \widehat{y}_{t+1}. \tag{32}$$

From equation (32), it follows that we need

$$\phi_\pi > \bar{m}^r - \phi_y + \frac{\omega_f - 1}{\omega_c \frac{\kappa}{\gamma}} = \bar{m}^r - \phi_y + \frac{\bar{m}\delta - 1}{\frac{1-\lambda}{1-\chi\lambda}\frac{\kappa}{\gamma}} \tag{33}$$

for the model to feature a determinate, locally unique equilibrium. Condition (33) shows that both, $\bar{m}^r < 1$ and $\phi_y > 0$, weaken the condition in Proposition 3. Put differently, bounded rationality with respect to the real rate or a Taylor rule that responds to changes in output, both relax the condition on ϕ_π to yield determinacy.

A.5 IS Curve with Government Spending

Government spending is financed by uniform taxes, $\tau_t^H = \tau_t^U = G_t$, household H's net income is:

$$\widehat{c}_t^H = \widehat{w}_t + \widehat{n}_t^H + \frac{\mu^D}{\lambda}\widehat{d}_t - g_t, \tag{34}$$

where $g_t = log(G_t/Y)$.

We first derive households H consumption as a function of total income $\widehat{y}_t$. The good markets clearing condition is now

$$\widehat{y}_t = \lambda \widehat{c}_t^H + (1-\lambda)\widehat{c}_t^U + g_t. \tag{35}$$

Plugging this and the labor market clearing condition into (26), yields:

$$\varphi\widehat{n}_t^H + \gamma\widehat{c}_t^H = (\varphi+\gamma)\widehat{y}_t - \gamma g_t. \tag{36}$$

Replacing wages and the dividends in the households' budget constraint yields:

$$\widehat{c}_t^H = \left(\varphi\widehat{n}_t^H + \gamma\widehat{c}_t^H\right)\left(1 - \frac{\mu^D}{\lambda}\right) + \widehat{n}_t^H - g_t. \tag{37}$$

and using (36) yields:

$$\widehat{c}_t^H = \left(\varphi\widehat{n}_t^H + \gamma\widehat{c}_t^H\right)\left(1 - \frac{\mu^D}{\lambda}\right) + \widehat{n}_t^H - g_t. \tag{38}$$

Finally, consumption of H is given by:

$$\widehat{c}_t^H = \chi\widehat{y}_t - \left[\frac{\chi-1}{1+\frac{\varphi}{\gamma}} + 1\right] g_t \tag{39}$$

which is

$$\widehat{c}_t^H = \chi\left(\widehat{y}_t - g_t\right) + \left[\frac{\chi-1}{1+\frac{\gamma}{\varphi}}\right] g_t. \tag{40}$$

The consumption of unconstrained households is then given by (using the market clearing condition):

$$\widehat{c}_t^U = \frac{1-\lambda\chi}{1-\lambda}\left(\widehat{y}_t - g_t\right) - \frac{\lambda}{1-\lambda}\frac{\chi-1}{1+\frac{\gamma}{\varphi}} g_t. \tag{41}$$

The IS curve in terms of aggregate consumption is then obtained by plugging the consumption of the hand-to-mouth and of unconstrained households into the Euler equation of unconstrained households and using $\widehat{c}_t = (\widehat{y}_t - g_t)$.

$$\begin{aligned}\frac{1-\lambda\chi}{1-\lambda}\widehat{c}_t - \frac{\lambda}{1-\lambda}\frac{\chi-1}{1+\frac{\gamma}{\varphi}} g_t = s\mathbb{E}_t^{BR}&\left[\frac{1-\lambda\chi}{1-\lambda}\widehat{c}_{t+1} - \frac{\lambda}{1-\lambda}\frac{\chi-1}{1+\frac{\gamma}{\varphi}} g_{t+1}\right] \\ &+(1-s)\mathbb{E}_t^{BR}\left[\chi\widehat{c}_{t+1} + \left[\frac{\chi-1}{1+\frac{\gamma}{\varphi}}\right] g_{t+1}\right] - \frac{1}{\gamma}\mathbb{E}_t(\widehat{i}_t - \pi_{t+1}),\end{aligned}$$

which can be re-written as (using similar derivations as in Appendix A.2)

$$\begin{aligned}\widehat{c}_t &= \psi_f\mathbb{E}_t\widehat{c}_{t+1} - \frac{1}{\gamma}\psi_c\mathbb{E}_t(\widehat{i}_t - \pi_{t+1}) + \frac{\lambda}{1+\frac{\gamma}{\varphi}}\frac{\chi-1}{1-\lambda\chi} g_t \\ &- \left[s\frac{\lambda}{1-\lambda\chi}\frac{\chi-1}{1+\frac{\gamma}{\varphi}} + (1-s)\frac{\chi-1}{1+\frac{\gamma}{\varphi}}\frac{1-\lambda}{1-\chi\lambda}\right]\mathbb{E}_t^{BR} g_{t+1} \\ &= \psi_f\mathbb{E}_t\widehat{c}_{t+1} - \frac{1}{\gamma}\psi_c\mathbb{E}_t(\widehat{i}_t - \pi_{t+1}) + \zeta\left[\frac{\lambda(\chi-1)}{1-\lambda\chi}(g_t - \bar{m}\mathbb{E}_t g_{t+1}) + (\delta-1)\bar{m}\mathbb{E}_t g_{t+1}\right]\end{aligned}$$

with $\zeta = \frac{1}{1+\frac{\gamma}{\varphi}}$. Replacing the expectations and taking the derivative with respect to g_t yields the consumption multiplier.

A.6 Derivation of Lemma 1

Let us first state a few auxiliary results that will prove helpful later. First, in log-linearized terms, the stochastic discount factor is given by

$$\frac{1}{\gamma}\mathbb{E}_t^{BR}\widehat{q}_{t,t+1}^U = \widehat{c}_t^U - s\bar{m}\mathbb{E}_t\widehat{c}_{t+1}^U - (1-s)\bar{m}\mathbb{E}_t\widehat{c}_{t+1}^H$$

and for i periods ahead:

$$\frac{1}{\gamma}\mathbb{E}_t^{BR}\widehat{q}_{t,t+i}^U = \widehat{c}_t^U - s\bar{m}^i\mathbb{E}_t\widehat{c}_{t+i}^U - (1-s)\bar{m}^i\mathbb{E}_t\widehat{c}_{t+i}^H.$$

Furthermore, we have:

$$\begin{aligned}\frac{1}{\gamma}\mathbb{E}_t^{BR}\widehat{q}_{t+1,t+2}^U &= \mathbb{E}_t^{BR}\left[\widehat{c}_{t+1}^U - s\widehat{c}_{t+2}^U - (1-s)\widehat{c}_{t+2}^H\right]\\ &= \bar{m}\mathbb{E}_t\widehat{c}_{t+1}^U - s\bar{m}^2\mathbb{E}_t\widehat{c}_{t+2}^U - (1-s)\bar{m}^2\mathbb{E}_t\widehat{c}_{t+2}^H\end{aligned}$$

and the stochastic discount factor has the property

$$\mathbb{E}_t^{BR}\left[\widehat{q}_{t,t+i}^U\right] = \mathbb{E}_t^{BR}\left[\widehat{q}_{t,t+1}^U + \widehat{q}_{t+1,t+2}^U + ... + \widehat{q}_{t+i-1,t+i}^U\right].$$

Using these results, $\mathbb{E}_t^{BR}\left[\widehat{q}_{t,t+i}^U\right]$ can be written as

$$\begin{aligned}\frac{1}{\gamma}\mathbb{E}_t^{BR}\widehat{q}_{t,t+i}^U = \widehat{c}_t^U &+ (1-s)\bar{m}\mathbb{E}_t\left[\widehat{c}_{t+1}^U - \widehat{c}_{t+1}^H\right]\\ &+ (1-s)\bar{m}^2\mathbb{E}_t\left[\widehat{c}_{t+2}^U - \widehat{c}_{t+2}^H\right] + ...+\\ &+ (1-s)\bar{m}^i\mathbb{E}_t\left[\widehat{c}_{t+i}^U - \widehat{c}_{t+1}^H\right] - \bar{m}^i\mathbb{E}_t\widehat{c}_{t+i}^U,\end{aligned}$$

or put differently

$$\frac{1}{\gamma}\mathbb{E}_t^{BR}\widehat{q}_{t,t+i}^U + \bar{m}^i\mathbb{E}_t\widehat{c}_{t+i}^U = \widehat{c}_t^U + (1-s)\mathbb{E}_t\sum_{k=1}^{i}\bar{m}^k\left(\widehat{c}_{t+k}^U - \widehat{c}_{t+k}^H\right). \tag{42}$$

The (linearized) budget constraint can be written as

$$\mathbb{E}_t^{BR}\sum_{i=0}^{\infty}\beta^i\left(\frac{1}{\gamma}\widehat{q}_{t,t+i}^U + \widehat{c}_{t+i}^U\right) = \mathbb{E}_t^{BR}\sum_{i=0}^{\infty}\beta^i\left(\frac{1}{\gamma}\widehat{q}_{t,t+i}^U + \widehat{y}_{t+i}^U\right)$$

$$\Leftrightarrow \mathbb{E}_t^{BR}\sum_{i=0}^{\infty}\beta^i\left(\frac{1}{\gamma}\widehat{q}_{t,t+i}^U\right) + \mathbb{E}_t\sum_{i=0}^{\infty}(\beta\bar{m})^i\widehat{c}_{t+i}^U = \mathbb{E}_t^{BR}\sum_{i=0}^{\infty}\beta^i\left(\frac{1}{\gamma}\widehat{q}_{t,t+i}^U\right) + \mathbb{E}_t\sum_{i=0}^{\infty}(\beta\bar{m})^i\widehat{y}_{t+i}^U.$$

Now, focus on the left-hand side and notice that the sum $\mathbb{E}_t\sum_{i=0}^{\infty}(\beta\bar{m})^i\widehat{c}_{t+i}^U$ cancels with the $\bar{m}^i\mathbb{E}_t\widehat{c}_{t+i}^U$ terms in equation (42) when summing them up. The left-hand side of the budget constraint can thus be written as

$$\begin{aligned}&\mathbb{E}_t\sum_{i=0}^{\infty}\beta^i\left(\widehat{c}_t^U + (1-s)\sum_{k=1}^{i}\bar{m}^k\left(\widehat{c}_{t+k}^U - \widehat{c}_{t+k}^H\right)\right)\\ &= \frac{1}{1-\beta}\widehat{c}_t^U + (1-s)\mathbb{E}_t\sum_{i=0}^{\infty}\beta^i\sum_{k=1}^{i}\bar{m}^k\left(\widehat{c}_{t+k}^U - \widehat{c}_{t+k}^H\right)\end{aligned}$$

$$= \frac{1}{1-\beta}\widehat{c}_t^U + \frac{1-s}{1-\beta}\mathbb{E}_t \sum_{i=1}^{\infty} (\beta\bar{m})^i \left(\widehat{c}_{t+i}^U - \widehat{c}_{t+i}^H\right).$$

Note, from the Euler equation of the unconstrained households, we obtain the real interest rate

$$-\frac{1}{\gamma}\widehat{r}_t = \widehat{c}_t^U - s\mathbb{E}_t^{BR}\widehat{c}_{t+1}^U - (1-s)\mathbb{E}_t^{BR}\widehat{c}_{t+1}^H$$
$$= \frac{1}{\gamma}\mathbb{E}_t^{BR}\widehat{q}_{t,t+1}^U,$$

and similarly,

$$-\frac{1}{\gamma}\bar{m}^i\mathbb{E}_t\widehat{r}_{t+i} = \frac{1}{\gamma}\mathbb{E}_t^{BR}\widehat{q}_{t+i,t+i+1}^U,$$

where $\widehat{r}_t$ is the (linearized) real interest rate.

Combining these results, we see that

$$\mathbb{E}_t^{BR}\sum_{i=0}^{\infty}\beta^i\frac{1}{\gamma}\widehat{q}_{t,t+i}^U = -\frac{1}{1-\beta}\frac{1}{\gamma}\beta\mathbb{E}_t\sum_{i=0}^{\infty}(\beta\bar{m})^i\,\widehat{r}_{t+i}.$$

Plugging this into the right-hand side of the budget constraint and multiplying both sides by $1-\beta$ yields

$$\widehat{c}_t^U = -\frac{1}{\gamma}\beta\widehat{r}_t + (1-\beta)\widehat{y}_t^U - (1-s)\mathbb{E}_t\sum_{i=1}^{\infty}(\beta\bar{m})^i\left(\widehat{c}_{t+i}^U - \widehat{c}_{t+i}^H\right)$$
$$-\frac{1}{\gamma}\beta\mathbb{E}_t\sum_{i=1}^{\infty}(\beta\bar{m})^i\,\widehat{r}_{t+i} + (1-\beta)\mathbb{E}_t\sum_{i=1}^{\infty}(\beta\bar{m})^i\,\widehat{y}_{t+i}^U,$$

or written recursively

$$\widehat{c}_t^U = -\frac{1}{\gamma}\beta\widehat{r}_t + (1-\beta)\widehat{y}_t^U + \beta\bar{m}s\mathbb{E}_t\widehat{c}_{t+1}^U + \beta\bar{m}(1-s)\mathbb{E}_t\widehat{c}_{t+1}^H.$$

Now, aggregating, i.e., multiplying the expression for $\widehat{c}_t^U$ by $(1-\lambda)$, adding $\lambda\widehat{c}_t^H$ and using $\widehat{c}_t^H = \chi\widehat{y}_t$ as well as $\widehat{y}_t^U = \frac{1-\lambda\chi}{1-\lambda}\widehat{y}_t$, yields the consumption function

$$\widehat{c}_t = \left[1-\beta(1-\lambda\chi)\right]\widehat{y}_t - \frac{(1-\lambda)\beta}{\gamma}\widehat{r}_t + \beta\bar{m}\delta(1-\lambda\chi)\mathbb{E}_t\widehat{c}_{t+1}, \tag{43}$$

as stated in the main text.

To obtain the share of indirect effects, note that the model does not feature any endogenous state variables and hence, endogenous variables inherit the persistence of the exogenous variables, ρ. Thus, $\mathbb{E}_t\widehat{c}_{t+1} = \rho\widehat{c}_t$. Plugging this into the consumption function (43), we get

$$\widehat{c}_t = \frac{1-\beta(1-\lambda\chi)}{1-\beta\bar{m}\delta\rho(1-\lambda\chi)}\widehat{y}_t - \frac{(1-\lambda)\beta}{\gamma(1-\beta\bar{m}\delta\rho(1-\lambda\chi))}\widehat{r}_t.$$

The term in front of $\widehat{y}_t$ is the share of indirect effects.

A.7 Derivation of Proposition 5

To prove Proposition 5, we start from the Euler equation (12). Plugging in for $\widehat{c}_t^U$, $\widehat{c}_{t+1}^U$ and $\widehat{c}_{t+1}^H$ from equations (8) and (10), we get

$$\widehat{y}_t = s\mathbb{E}_t^{BR}\left[\widehat{y}_{t+1}\right] + (1-s)\frac{1-\lambda}{1-\lambda\chi}\mathbb{E}_t^{BR}\left[\widehat{y}_{t+1}\right] - \psi_c\left(\widehat{i}_t - \mathbb{E}_t\pi_{t+1}\right),$$

which can be re-written as

$$\widehat{y}_t = \delta\mathbb{E}_t^{BR}\left[\widehat{y}_{t+1}\right] - \psi_c\left(\widehat{i}_t - \mathbb{E}_t\pi_{t+1}\right).$$

Now, using the expectations setup from Proposition 5, we get $\delta\mathbb{E}_t^{BR}\left[\widehat{y}_{t+1}\right] = (1-\bar{m})\delta\widehat{y}_{t-1} + \bar{m}\delta\mathbb{E}_t\left[\widehat{y}_{t+1}\right]$ which proves Proposition 5.

A.8 Cognitive Discounting of the State Vector

In Section 2, we assume that cognitive discounting applies to all variables, which differs slightly from the assumption in Gabaix (2020) who assumes that cognitive discounting applies to the *state* of the economy (exogenous shocks as well as announced monetary and fiscal policies). He then proves (Lemma 1 in Gabaix (2020)) how cognitive discounting applies as a result (instead of as an assumption) to all future variables, including future consumption choices. For completeness, we show in this section how our results are unaffected when following the approach in Gabaix (2020).

Let X_t denote the (de-meaned) state vector which evolves as

$$X_{t+1} = G^X\left(X_t, \varepsilon_{t+1}\right), \tag{44}$$

where G^X denotes the transition function of X in equilibrium and ε are zero-mean innovations. Linearizing equation (44) yields

$$X_{t+1} = \Gamma X_t + \varepsilon_{t+1}, \tag{45}$$

where ε_{t+1} might have been renormalized. The assumption in Gabaix (2020) is that the behavioral agent perceives the state vector to follow

$$X_{t+1} = \bar{m}G^X(X_t, \varepsilon_{t+1}), \tag{46}$$

or in linearized terms

$$X_{t+1} = \bar{m}\left(\Gamma X_t + \varepsilon_{t+1}\right). \tag{47}$$

The expectation of the boundedly-rational agent of X_{t+1} is thus $\mathbb{E}_t^{BR}\left[X_{t+1}\right] = \bar{m}\mathbb{E}_t\left[X_{t+1}\right] = \bar{m}\Gamma X_t$. Iterating forward, it follows that $\mathbb{E}_t^{BR}\left[X_{t+k}\right] = \bar{m}^k\mathbb{E}_t\left[X_{t+k}\right] = \bar{m}^k\Gamma^k X_t$.

Now, consider any variable $z(X_t)$ with $z(0) = 0$ (e.g., demeaned consumption of unconstrained households $C^U(X_t)$). Linearizing $z(X)$, we obtain $z(X) = b_X^z X$ for some b_X^z and thus

$$\begin{aligned}
\mathbb{E}_t^{BR}\left[z(X_{t+k})\right] &= \mathbb{E}_t^{BR}\left[b_X^z X_{t+k}\right] \\
&= b_X^z\mathbb{E}_t^{BR}\left[X_{t+k}\right] \\
&= b_X^z\bar{m}^k\mathbb{E}_t\left[X_{t+k}\right]
\end{aligned}$$

$$= \bar{m}^k \mathbb{E}_t \left[b_X^z X_{t+k} \right]$$
$$= \bar{m}^k \mathbb{E}_t \left[z(X_{t+k}) \right].$$

For example, expected consumption of unconstrained households tomorrow (in linearized terms) is given by

$$\mathbb{E}_t^{BR} \left[\widehat{c}^U(X_{t+1}) \right] = \bar{m} \mathbb{E}_t \left[\widehat{c}^U(X_{t+1}) \right], \tag{48}$$

which we denote in the main text as

$$\mathbb{E}_t^{BR} \left[\widehat{c}_{t+1}^U \right] = \bar{m} \mathbb{E}_t \left[\widehat{c}_{t+1}^U \right]. \tag{49}$$

Now, take the linearized Euler equation (12) of unconstrained households:

$$\widehat{c}_t^U = s \mathbb{E}_t^{BR} \left[\widehat{c}_{t+1}^U \right] + (1-s) \mathbb{E}_t^{BR} \left[\widehat{c}_{t+1}^H \right] - \frac{1}{\gamma} \widehat{r}_t, \tag{50}$$

where $\widehat{r}_t \equiv \widehat{i}_t - \mathbb{E}_t \pi_{t+1}$.

Using the notation in Gabaix (2020), we can write the Euler equation as

$$\widehat{c}^U(X_t) = s \mathbb{E}_t^{BR} \left[\widehat{c}^U(X_{t+1}) \right] + (1-s) \mathbb{E}_t^{BR} \left[\widehat{c}^H(X_{t+1}) \right] - \frac{1}{\gamma} \widehat{r}(X_t). \tag{51}$$

Now, applying the results above, we obtain

$$\widehat{c}^U(X_t) = s \bar{m} \mathbb{E}_t \left[\widehat{c}^U(X_{t+1}) \right] + (1-s) \bar{m} \mathbb{E}_t \left[\widehat{c}^H(X_{t+1}) \right] - \frac{1}{\gamma} \widehat{r}(X_t), \tag{52}$$

which after writing $\widehat{c}^U(X_t)$, $\widehat{c}^U(X_{t+1})$ and $\widehat{c}^H(X_{t+1})$ in terms of total output yields exactly the IS equation in Proposition 1.

A.9 Microfounding $\bar{m}$

Gabaix (2020) shows how to microfound $\bar{m}$ from a noisy signal extraction problem in the case of a representative agent. Following these lines, we show how this signal-extraction problem generates a setup in which the family head behaves as if she was boundedly rational.

The (linearized) law of motion of the state variable, X_t, is given by $X_{t+1} = \Gamma X_t + \varepsilon_{t+1}$ (a similar reasoning extends to the non-linearized case), where X has been demeaned. Now assume that every agent j within the family of unconstrained households (the expectations of the hand-to-mouth agents are irrelevant) receives a noisy signal of X_{t+1}, S_{t+1}^j, given by

$$S_{t+1}^j = \begin{cases} X_{t+1} & \text{with probability } p \\ X'_{t+1} & \text{with probability } 1-p \end{cases}$$

where X'_{t+1} is an i.i.d. draw from the unconditional distribution of X_{t+1}, which has an unconditional mean of zero. In words, with probability p the agent j receives perfectly precise information and with probability $1-p$ agent j receives a signal realization that is completely uninformative. A fully-informed rational agent would have $p = 1$.

The conditional mean of X_{t+1}, given the signal S^j_{t+1}, is given by

$$X^e_{t+1} \equiv \mathbb{E}\left[X_{t+1}|S_{t+1} = s^j_{t+1}\right] = p \cdot s^j_{t+1}.$$

To see this, note that the joint distribution of (X_{t+1}, S^j_{t+1}) is

$$f(x_{t+1}, s^j_{t+1}) = pg(s^j_{t+1})\delta_{s^j_{t+1}}(x_{t+1}) + (1-p)g(s^j_{t+1})g(x_{t+1}),$$

where $g(X_{t+1})$ denotes the distribution of X_{t+1} and δ is the Dirac function. Given that the unconditional mean of X_{t+1} is 0, i.e., $\int x_{t+1}g(x_{t+1})dx_{t+1} = 0$, it follows that

$$\begin{aligned} \mathbb{E}_t\left[X_{t+1}|S^j_{t+1} = s^j_{t+1}\right] &= \frac{\int x_{t+1}f(x_{t+1}, s^j_{t+1})dx_{t+1}}{\int f(x_{t+1}, s^j_{t+1})dx_{t+1}} \\ &= \frac{pg(s^j_{t+1})s^j_{t+1} + (1-p)g(s^j_{t+1})\int x_{t+1}g(x_{t+1})dx_{t+1}}{g(s^j_{t+1})} \\ &= ps^j_{t+1}. \end{aligned}$$

The intuition is that the signal distribution is such that the agent either receives a perfectly precise signal or a completely uninformative signal. As the perfectly-precise signal arrives with probability p and the unconditional mean is zero, it follows that the agent puts a weight p on the signal.

Furthermore, we have

$$\mathbb{E}\left[S_{t+1}|X_{t+1}\right] = pX_{t+1} + (1-p)\mathbb{E}\left[X'_{t+1}\right] = pX_{t+1}.$$

So, it follows that the *average* expectation of X_{t+1} within the family is given by

$$\begin{aligned} \mathbb{E}\left[X^e_{t+1}(S_{t+1})|X_{t+1}\right] &= \mathbb{E}\left[p \cdot S_{t+1}|X_{t+1}\right] \\ &= p \cdot \mathbb{E}\left[S_{t+1}|X_{t+1}\right] \\ &= p^2 X_{t+1}. \end{aligned}$$

Defining $\bar{m} \equiv p^2$ and since $X_{t+1} = \Gamma X_t + \varepsilon_{t+1}$, we have that the family head perceives the law of motion of X to equal

$$X_{t+1} = \bar{m}\left(\Gamma X_t + \varepsilon_{t+1}\right), \tag{53}$$

as imposed in equation (47). The boundedly-rational expectation of X_{t+1} is then given by

$$\mathbb{E}^{BR}_t\left[X_{t+1}\right] = \bar{m}\mathbb{E}_t\left[X_{t+1}\right].$$

B Calibration

Our baseline calibration is summarized in Table 1. The values for γ and κ are directly taken from Bilbiie (2021, 2020) and are quite standard in the literature. Gabaix (2020), however, sets $\kappa = 0.11$ and $\gamma = 5$. Even though these coefficients differ quite substantially from our baseline calibration, note that our results would barely be affected by this. To see this, note that *amplification* is only determined by λ and χ, both independent of κ and γ. The determinacy condition on the other hand depends on both, κ and γ, but what ultimately matters is the fraction $\frac{\kappa}{\gamma}$ (see Proposition 3). As κ and γ are both approximately five times larger in Gabaix (2020) compared to Bilbiie (2021) and our baseline calibration, the fraction is approximately the same and thus, the determinacy region under an interest-rate peg remains unchanged.

Table 1: Tractable Model: Baseline Calibration

Parameter	Description	Value
γ	Risk Aversion	1
κ	Slope of NKPC	0.02
χ	Business-Cycle Exposure of H	1.35
λ	Share of H	0.33
s	Type-Switching Probability	$0.8^{1/4}$
β	Time Discount Factor	0.99
$\bar{m}$	Cognitive Discounting Parameter	0.85

The household heterogeneity parameters, χ, λ and s are also standard in the analytical HANK literature (see Bilbiie (2020)). The most important assumption for our qualitative results in Section 3 is $\chi > 1$. With $\lambda = 0.33$, a χ of 1.35 implies an amplification of conventional monetary policy shocks of 20% compared to RANK, as estimated by Patterson (2019).

For figure 6, i.e., to compute the iMPCs we choose a yearly calibration with $s = 0.8$ and $\beta = 0.95$.

The cognitive discounting parameter $\bar{m}$. The cognitive discounting parameter $\bar{m}$ is set to 0.85, as in Gabaix (2020) and Benchimol and Bounader (2019). Fuhrer and Rudebusch (2004), for example, estimate an IS equation and find that $\psi_f \approx 0.65$, which together with $\delta > 1$, would imply a $\bar{m}$ much lower than 0.85 and especially our determinacy results would be even stronger under such a calibration.

Another way to calibrate $\bar{m}$ (as pointed out in Gabaix (2020)) is to interpret the estimates in Coibion and Gorodnichenko (2015) through the "cognitive-discounting lens". They regress forecast errors on forecast revisions

$$x_{t+h} - F_t x_{t+h} = c + b^{CG} \left(F_t x_{t+h} - F_{t-1} x_{t+h} \right) + u_t,$$

where $F_t x_{t+h}$ denotes the forecast at time t of variable x, h periods ahead. Focusing on inflation, they find that $b^{CG} > 0$ in consensus forecasts, pointing to *underreaction* (similar results are, for example, found in Angeletos et al. (2021) and Adam et al. (2022) for other variables).

In the model, the law of motion of x is $x_{t+1} = \Gamma(x_t + \varepsilon_{t+1})$ whereas the behavioral agents perceive it to be $x_{t+1} = \bar{m}\Gamma(x_t + \varepsilon_{t+1})$. It follows that $F_t x_{t+h} = (\bar{m}\Gamma)^h x_t$ and thus, forecast revisions are equal to

$$\begin{aligned} F_t x_{t+h} - F_{t-1} x_{t+h} &= (\bar{m}\Gamma)^h x_t - (\bar{m}\Gamma)^{h+1} x_{t-1} \\ &= (\bar{m}\Gamma)^h \Gamma(1-\bar{m}) x_{t-1} + (\bar{m}\Gamma)^h \varepsilon_t. \end{aligned}$$

The forecast error is given by

$$x_{t+h} - F_t x_{t+h} = \Gamma^h(1-\bar{m}^h)\Gamma x_{t-1} + \Gamma^h(1-\bar{m}^h)\varepsilon_t + \sum_{j=0}^{h-1} \Gamma^j \varepsilon_{t+h-j},$$

where $\sum_{j=0}^{h-1} \Gamma^j \varepsilon_{t+h-j}$ is the rational expectations forecast error. Gabaix (2020) shows that b^{CG} is bounded below $b^{CG} \geq \frac{1-\bar{m}^h}{\bar{m}^h}$, showing that $\bar{m} < 1$ yields $b^{CG} > 0$, as found empirically. When replacing the weak inequality with an equality, we get

$$\bar{m}^h = \frac{1}{1+b^{CG}}.$$

Most recently, Angeletos et al. (2021) estimate b^{CG} (focusing on a horizon $h = 3$) to lie between $b^{CG} \in [0.74, 0.81]$ for unemployment forecasts and $b^{CG} \in [0.3, 1.53]$ for inflation, depending on the considered period (see their Table 1). These estimates imply $\bar{m} \in [0.82, 0.83]$ for unemployment and $\bar{m} \in [0.73, 0.92]$ for inflation, and are thus close to our preferred value of 0.85. Note, however, that these estimates pertain to professional forecasters and should therefore be seen as upper bounds on $\bar{m}$. We now turn to direct evidence on $\bar{m}$ for households (of different income groups). We find that households are less rational than professional forecasters.

B.1 Estimating $\bar{m}$ for different Household Groups

To test for heterogeneity in the degree of cognitive discounting, we follow Coibion and Gorodnichenko (2015) and regress forecast errors on forecast revisions as follows

$$x_{t+4} - \mathbb{E}_t^{e,BR} x_{t+4} = c^e + b^{e,CG}\left(\mathbb{E}_t^{e,BR} x_{t+4} - \mathbb{E}_{t-1}^{e,BR} x_{t+4}\right) + \epsilon_t^e, \tag{54}$$

to estimate $b^{e,CG}$ for different groups of households, indexed by e. As shown above, $b^{e,CG} > 0$ is consistent with underreaction and the corresponding cognitive discounting parameter is approximately given by

$$\bar{m}^e = \left(\frac{1}{1+b^{e,CG}}\right)^{1/4}. \tag{55}$$

Ideally, we would use actual data and expectations data about future marginal utilities of consumption which, however, are not available. Instead, we focus on expectations about future unemployment. The Survey of Consumers from the University of Michigan provides 1-year ahead unemployment expectations and we use the unemployment rate from the FRED database as our measure of actual unemployment. Consistent with the model, we split the households into three groups based on their income. The bottom and top income groups each contain the 25% households with the lowest and highest income, respectively, and the remaining 50% are assigned to the middle income group.

The Michigan Survey asks households whether they expect unemployment to increase, decrease or to remain about the same over the next twelve months. We follow Carlson and Parkin (1975), Mankiw (2000) and Bhandari et al. (2019) to translate these categorical unemployment expectation into numerical expectations.

Focus on group $e \in \{L, M, H\}$ and let $q_t^{e,D}$, $q_t^{e,S}$ and $q_t^{e,U}$ denote the shares within income group e reported at time t that think unemployment will go down, stay roughly the same, or go up over the next year, respectively. We assume that these shares are drawn from a cross-sectional distribution of responses that are normally distributed according to $\mathcal{N}\left(\mu_t^e, (\sigma_t^e)^2\right)$ and a threshold a such that when a household expects unemployment to remain within the range $[-a, a]$ over the next year, she responds that unemployment will remaine "about the same". We thus have

$$q_t^{e,D} = \Phi\left(\frac{-a - \mu_t^e}{\sigma_t^e}\right) \qquad q_t^{e,U} = 1 - \Phi\left(\frac{a - \mu_t^e}{\sigma_t^e}\right),$$

which after some rearranging yields

$$\sigma_t^e = \frac{2a}{\Phi^{-1}\left(1 - q_t^{e,U}\right) - \Phi^{-1}\left(q_t^{e,D}\right)}$$
$$\mu_t^e = a - \sigma_t^e \Phi^{-1}\left(1 - q_t^{e,U}\right).$$

This leaves us with one degree of freedom, namely a. We make two assumptions. First, a is independent of the income group. The second assumption is that we set $a = 0.5$ which means that if a household expects the change in unemployment to be less than half a percentage point (in absolute terms), she reports that she expects unemployment to be about the same as it is at the time of the survey. We discuss different a later on.

As the question in the survey is about the expected change in unemployment, we add the actual unemployment rate at the time of the survey to μ_t^e to construct a time-series of unemployment expectations, as in Bhandari et al. (2019). That said, we will also report the case of expected unemployment *changes*.

Given the so-constructed expectations, we can compute forecast revisions as

$$\mu_t^e - \mu_{t-1}^e$$

and four-quarter-ahead forecast errors using the actual unemployment rate u_t obtained from FRED as

$$u_{t+4} - \mu_t^e. \tag{56}$$

For the case of expected unemployment changes, we replace u_{t+4} with $(u_{t+4} - u_t)$ in equation (56).

Following Coibion and Gorodnichenko (2015), we then regress forecast errors on forecast revisions

$$u_{t+4} - \mu_t^e = c^e + b^{e,CG} \left(\mu_t^e - \mu_{t-1}^e \right) + \epsilon_t^e, \tag{57}$$

to estimate $b^{e,CG}$ for each income group e. Note, however, that the expectations in the forecast revisions are about unemployment at different points in time. To account for this, we instrument forecast revisions by the *main business cycle shock* obtained from Angeletos et al. (2020).

Table 2: Regression Results of Equation (54)

	IV Regression			OLS		
	Bottom 25%	Middle 50%	Top 25%	Bottom 25%	Middle 50%	Top 25%
$\widehat{b}^{e,CG}$	0.85	0.75	0.63	1.22	1.10	0.90
s.e.	(0.471)	(0.453)	(0.401)	(0.264)	(0.282)	(0.247)
F-stat.	24.76	18.74	17.86	-	-	-
N	152	152	152	157	157	157

Note: This table provides the estimated $\widehat{b}^{e,CG}$ from regression (54) for different income groups. The first three columns show the results when the right-hand side in equation (54) is instrumented using the *main business cycle shock* from Angeletos et al. (2020) and the last three columns using OLS. Standard errors are robust with respect to heteroskedasticity and are reported in parentheses. The row "F-stat." reports the first-stage F-statistic for the IV regressions.

Table 2 shows the results. The first three columns report the estimated $b^{e,CG}$ from the IV regressions and the last three columns the same coefficients estimated via OLS. Standard errors are robust with respect to heteroskedasticity and are reported in parentheses. The row "F-stat." reports the first-stage F-statistic for the IV regressions. We see that in all cases $\widehat{b}^{e,CG}$ is positive, suggesting that households of all income groups tend to underreact, consistent with our assumption of $\bar{m} < 1$.

Using equation (55) we obtain $\bar{m}^e$ equal to 0.86, 0.87 and 0.88 for the bottom 25%, the middle 50% and the top 25%, respectively for the estimates from the IV regressions and 0.82, 0.83 and 0.85 for the OLS estimates. When estimating $\bar{m}^e$ using expected unemployment *changes* instead of the level, the estimated $\bar{m}^e$ equal 0.57, 0.59 and 0.64 for the IV regressions and 0.77, 0.80 and 0.86 for the OLS regressions.

There are two main take-aways from this empirical exercise: first, it further confirms that $\bar{m} = 0.85$ is a reasonable (but rather conservative) deviation from rational expectations. Second, the data suggests that there is heterogeneity in the degree of rationality conditional on households income. In particular, households with higher income tend to exhibit higher degrees of rationality.[21]

If we consider inflation expectations instead of unemployment expectations, we obtain estimated cognitive discounting parameters of 0.70, 0.75 and 0.78 for the bottom 25%, the middle 50% and the top 25%, respectively. Thus, somewhat lower than for unemployment and the differ-

[21] This is consistent with other empirical findings on heterogeneous deviations from FIRE. Broer et al. (2021a), for example, document that wealthier households tend to have more accurate beliefs, as measured by forecast errors.

ences across income groups are larger. In particular, higher-income households tend to be more rational (they discount less) than lower-income households. The differences, however, are overall rather small.

C Figures to Section 3

C.1 Resolving the Catch-22

We graphically illustrate the Catch-22 (Bilbiie (2021)) of the rational model and the resolution of it in the behavioral HANK model in Figure 9. The figure shows on the vertical axis the response of contemporaneous output relative to the initial response in the RANK model with rational expectations for anticipated i.i.d. monetary policy shocks occurring at different times k on the horizontal axis.[22]

The orange-dotted line represents the baseline calibration of the rational HANK model. We see that this model is able to generate contemporaneous amplification of monetary policy shocks, that is, an output response that is relatively stronger than in RANK. Put differently the GE effects amplify the effects of monetary policy shocks. Yet, at the same time, it exacerbates the forward guidance puzzle as shocks occurring in the future have even stronger effects on today's output than contemporaneous shocks.

The black-dashed-dotted line shows how the forward guidance puzzle can be resolved by allowing for $\chi < 1$. Yet, this comes at the cost that the model is unable to generate amplification of contemporaneous monetary policy shocks. Recent empirical findings, however, document that GE effects indeed amplify monetary policy changes (Patterson (2019), Auclert (2019)).

Figure 9: Resolving the Catch-22

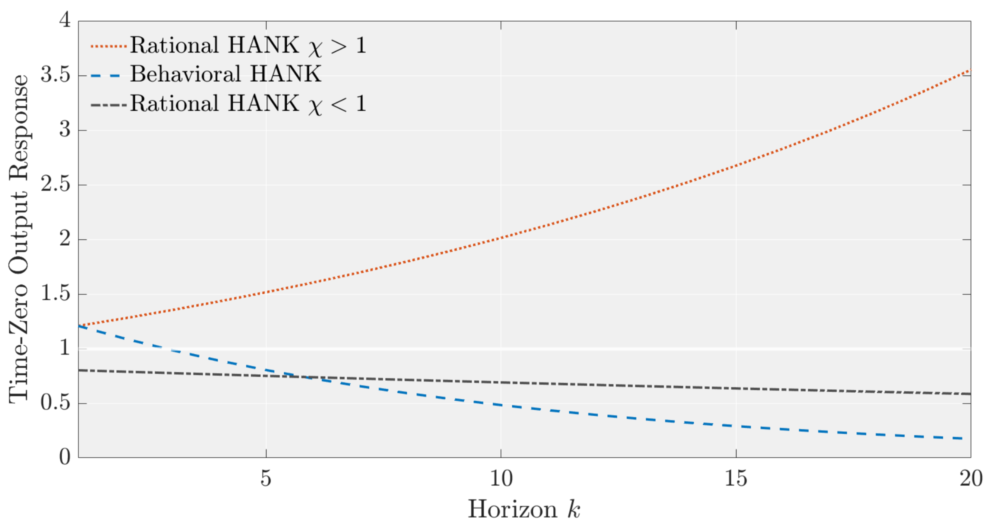

Note: This figure shows the response of total output in period 0 to anticipated i.i.d. monetary policy shocks occurring at different horizons k (horizontal axis), relative to the initial response in the RANK model under rational expectations (equal to 1).

The blue-dashed line shows that the behavioral HANK model, on the other hand, generates

[22]Under fully-rigid prices (i.e., $\kappa = 0$) the RANK model would deliver a constant response for all k. The same is true for two-agent NK models (TANK), i.e., tractable HANK models without type switching. Whether the constant response would lie above or below its RANK counterpart depends on $\chi \lesseqgtr 1$ in the same way the initial response depends on $\chi \lesseqgtr 1$.

Figure 10: The Effective Lower Bound Problem

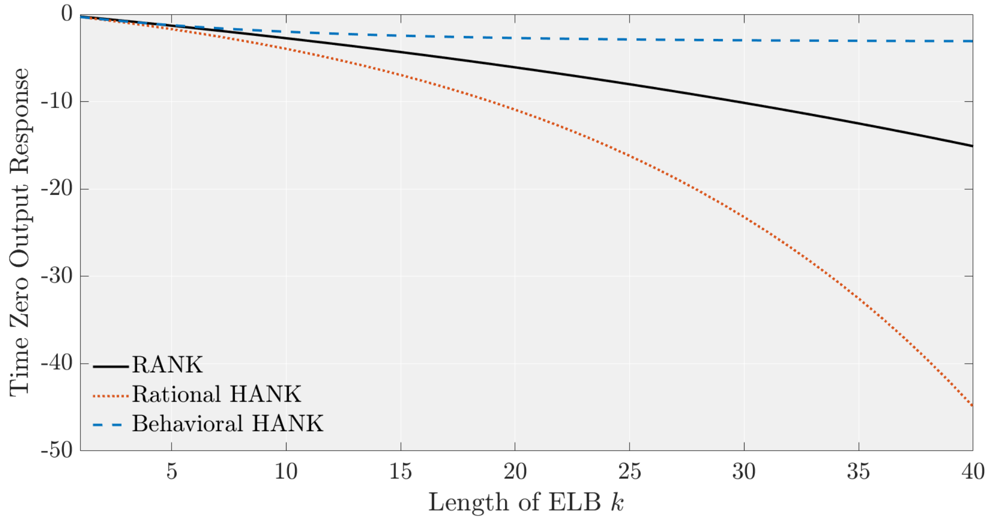

Note: This figure shows the contemporaneous output response for different lengths of a binding ELB k (horizontal axis) and compares the responses across different models.

both: amplification of contemporaneous monetary policy and a resolution of the forward guidance puzzle, both consistent with the empirical facts.

C.2 Stability at the Effective Lower Bound

We illustrate the stability of the behavioral HANK model at the lower bound graphically in Figure 10. The figure shows the output response in RANK, the rational HANK and the behavioral HANK to different lengths of a binding ELB (depicted on the horizontal axis). The shortcoming of monetary policy due to the ELB, i.e., the gap $\left(\widehat{i}_{ELB} - \widetilde{r}^n\right) > 0$, is set to a relatively small value of 0.25% (1% annually), and we set $\bar{m} = 0.85$. Figure 10 shows the implosion of output in the rational RANK (back-solid line) and even more so in the rational HANK model (orange-dotted line): an ELB that is expected to bind for 40 quarters would decrease today's output in the rational RANK by 15% and in the rational HANK model by 45%. On the other hand—and consistent with recent experiences in advanced economies—output in the behavioral HANK model remains quite stable and drops by a mere 3%, as illustrated by the blue-dashed line.

D Further Extensions

D.1 Fiscal Policy

We now show that the sufficient statistic for amplification of contemporaneous monetary policy is also a sufficient statistic to generate positive consumption multipliers of fiscal policy under constant real rates, as estimated empirically. Dupor et al. (2021) and Galí et al. (2007), for example, provide empirical evidence for positive effects of government spending on private consumption. Furthermore, Nakamura and Steinsson (2014) and Chodorow-Reich (2019) document fiscal multipliers above 1, which through the lens of our model is equivalent to saying that consumption responds positively to government spending.

To characterize fiscal multipliers, we assume government spending g_t to follow an AR(1) with persistence $\rho_g \geq 0$, and to be 0 in steady state. The government taxes all agents uniformly to finance g_t.

The behavioral HANK IS equation with government spending is given by:

$$\widehat{c}_t = \psi_f \mathbb{E}_t \widehat{c}_{t+1} - \psi_c \frac{1}{\gamma}\left(\widehat{i}_t - \mathbb{E}_t \pi_{t+1}\right) + \zeta \left[\frac{\lambda(\chi-1)}{1-\lambda\chi}\left(g_t - \bar{m}\mathbb{E}_t g_{t+1}\right) + (\psi_f - \bar{m})\mathbb{E}_t g_{t+1}\right],$$

where $\zeta \equiv \frac{\varphi}{\gamma\left(1+\frac{\varphi}{\gamma}\right)}$ (see appendix A.5). The static Phillips Curve in this setting is given by $\pi_t = \kappa c_t + \kappa \zeta g_t$. The following Proposition characterizes the fiscal multiplier in the behavioral HANK model.

Proposition 6. *The fiscal multiplier in the behavioral HANK model is given by*

$$\frac{\partial \widehat{c}_t}{\partial g_t} = \frac{1}{1-\nu\rho_g}\frac{\zeta}{1+\frac{1}{\gamma}\psi_c\phi\kappa}\left[\frac{\chi-1}{1-\lambda\chi}\left[\lambda(1-\bar{m}\rho_g) + \bar{m}\rho_g(1-s)\right] - \kappa\frac{1}{\gamma}\psi_c\left(\phi-\rho_g\right)\right],$$

where $\nu \equiv \frac{\psi_f + \kappa\frac{1}{\gamma}\psi_c}{1+\frac{1}{\gamma}\psi_c\phi\kappa}$.

To make the argument as clear as possible, we assume prices to be fully rigid, $\kappa = 0$ which, given our Taylor rule, implies that the real interest rate is held constant after the government spending shock. This is a useful benchmark as in this case the consumption response in RANK is 0 (see Bilbiie (2011) and Woodford (2011)).[23]

From Proposition 6, we derive the constant-real-rate multiplier in the behavioral HANK model:

$$\frac{\partial \widehat{c}_t}{\partial g_t} = \frac{1}{1-\nu\rho_g}\zeta\left[\frac{\chi-1}{1-\lambda\chi}\left[\lambda(1-\bar{m}\rho_g) + \bar{m}\rho_g(1-s)\right]\right].$$

As $\chi > 1$ the fiscal multiplier is bounded from below by 0 irrespective of the persistence ρ_g. In other words, the constant-real-rate multiplier in the behavioral HANK model is strictly positive. With $\chi > 1$ the high MPC households benefit disproportionately more from the extra income out of the increase in government spending which increases the fiscal multiplier through a Keynesian type argument.

[23] Auclert et al. (2018) also use a constant real interest rate case to show that their HANK model can generate (output) fiscal multipliers larger than one.

Figure 11 illustrates the effect of bounded rationality on the fiscal multiplier by plotting the fiscal multiplier in the behavioral HANK model for varying degrees of $\bar{m}$ (blue-solid line) and compares it to the multiplier in the rational HANK model and RANK. For this exercise, we set the persistence parameter to an intermediate value $\rho_g = 0.6$. It shows that the fiscal multiplier decreases with decreasing $\bar{m}$. Yet, even for the extreme case of $\bar{m} = 0$, in which households fully discount all future increases in government spending the fiscal multiplier is still substantially above zero even though it is somewhat weaker than under rational expectations. The results are consistent—although somewhat lower—with recent empirical estimates in Dupor et al. (2021) who estimate the non-durable consumption response to lie between 0.2 and 0.29.

Figure 11: Consumption Response to Government Spending

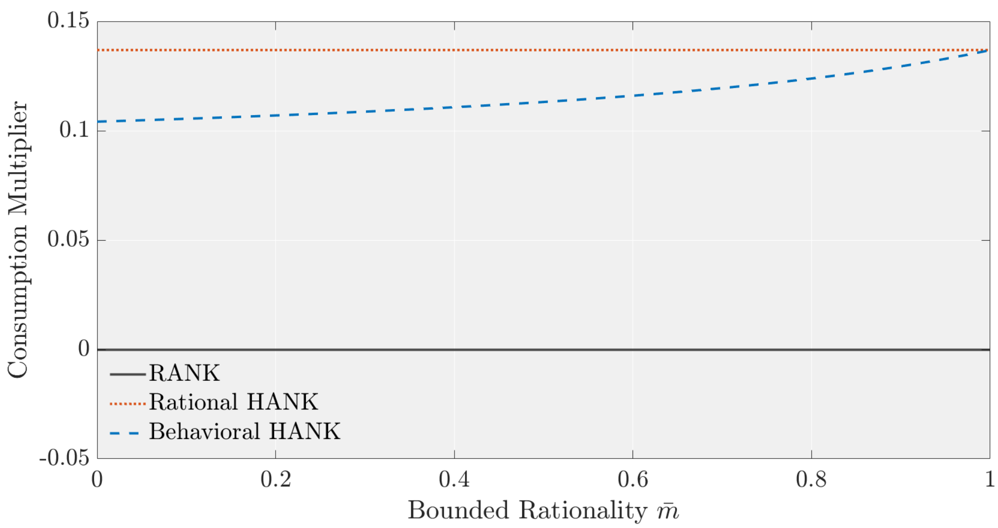

Note: This figure shows the consumption multipliers (the consumption response to government spending) for different degrees of bounded rationality (blue-dashed line). The orange-dotted line plots the multiplier in the rational version of the model and the black-solid line shows the zero-multiplier in the RANK model.

The behavioral HANK model does not rely on a specific financing type to achieve positive consumption responses to fiscal spending. This is in contrast to the behavioral RANK model in Gabaix (2020). In the behavioral RANK model, bounded rationality can also increase the multiplier but only if the government delays taxing the agents to finance the contemporaneous spending as boundedly-rational agents will then discount the future increases in taxes. In HANK models, on the other hand, the fiscal multiplier can in principle be larger than one with $\chi \leq 1$ if the hand-to-mouth households pay relatively less of the fiscal spending's cost than unconstrained households (see Bilbiie (2020) or Ferriere and Navarro (2018)). Both of these channels would also push up the multiplier in the behavioral HANK model, yet it does not depend on any of these two to achieve fiscal multipliers larger than 0.

A corollary of Propositon 6 is that with persistent government spending, $\rho_g > 0$, and with $\chi > 1$, more bounded rationality, i.e., a lower $\bar{m}$, leads to a lower fiscal multiplier.[24] Bounded

[24] We focus on the case in which $\nu\rho_g < 1$, which holds in the behavioral HANK model even for $\rho_g = 1$, and we

rationality decreases the fiscal multiplier as boundedly-rational agents discount the fact that an increase in government spending today has a positive effect on future spending as well. In the case of an i.i.d. spending shock the fiscal multiplier is independent of $\bar{m}$. Furthermore, the fiscal multiplier is bounded from above in the behavioral HANK model as $\nu\rho_g < 1$ even for highly persistent shocks. In the rational model, on the other hand, this is not the case. The fiscal multiplier approaches infinity as $\nu\rho_g \to 1$, which can occur because in the rational HANK model $\nu > 1$. As $\nu\rho_g > 1$ the multiplier even becomes negative. The behavioral HANK model, on the other hand, rules out these undesirable model implications.

D.2 Allowing for Steady State Inequality

So far, we have assumed that there is no steady state inequality, i.e., $C^H = C^U$. In the following, we relax this assumption and denote steady state inequality by $\Omega \equiv \frac{C^U}{C^H}$. Recall the Euler equation of unconstrained households

$$\left(C_t^U\right)^{-\gamma} = \beta R_t \mathbb{E}_t^{BR}\left[s\left(C_t^U\right)^{-\gamma} + (1-s)\left(C_t^H\right)^{-\gamma}\right],$$

from which we can derive the steady state real rate

$$R = \frac{1}{\beta(s + (1-s)\Omega^\gamma)}.$$

Log-linearizing the Euler equation yields

$$\widehat{c}_t^U = \beta R \bar{m}\left[s\mathbb{E}_t \widehat{c}_{t+1}^U + (1-s)\Omega^\gamma \mathbb{E}_t \widehat{c}_{t+1}^H\right] - \frac{1}{\gamma}\left(\widehat{i}_t - \mathbb{E}_t \pi_{t+1}\right).$$

Combining this with the consumption functions and the steady state real rate yields the IS equation

$$\widehat{y}_t = \bar{m}\tilde{\delta}\mathbb{E}_t \widehat{y}_{t+1} - \frac{1}{\gamma}\frac{1-\lambda}{1-\lambda\chi}\left(\widehat{i}_t - \mathbb{E}_t \pi_{t+1}\right), \tag{58}$$

with

$$\tilde{\delta} \equiv 1 + (\chi - 1)\frac{(1-s)\Omega^\gamma}{s + (1-s)\Omega^\gamma}\frac{1}{1-\lambda\chi}.$$

From a qualitative perspective, the whole analysis in Section 3 could be carried out with $\tilde{\delta}$ instead of δ. Quantitatively the differences are small as well. For example, if we set $\Omega = 1.5$, we get $\tilde{\delta} = 1.05$ instead of $\delta = 1.034$. Thus, we need $\bar{m} < 0.93$ instead of $\bar{m} < 0.94$ for determinacy under a peg.

D.3 Persistent Monetary Policy Shocks

In the main text in Section 3, we illustrated the resolution of the Catch-22 by considering i.i.d. monetary policy shocks (following Bilbiie (2021)). The behavioral HANK model delivers initial amplification of these monetary shocks but the effects decrease with the horizon of the shock, i.e., the behavioral HANK model resolves the forward guidance puzzle. Another way to see this is by considering persistent shocks.

assume $1 - s - \lambda < 0$, which holds under all reasonable parameterizations.

Figure 12 illustrates this. The figure shows the response of output in period t to a shock in period t for different degrees of persistence (x-axis). The black-solid line shows the output response in RANK and the blue-dashed line in the behavioral HANK. The forward guidance puzzle in RANK manifests itself in the sense that highly persistent shocks have stronger effects in RANK than in the behavioral HANK. Persistent shocks are basically a form of forward guidance and thus, with high enough persistence in the shocks, the RANK model predicts stronger effects than the behavioral HANK model.

Figure 12: Initial Output Response for Varying Degrees of the Persistence

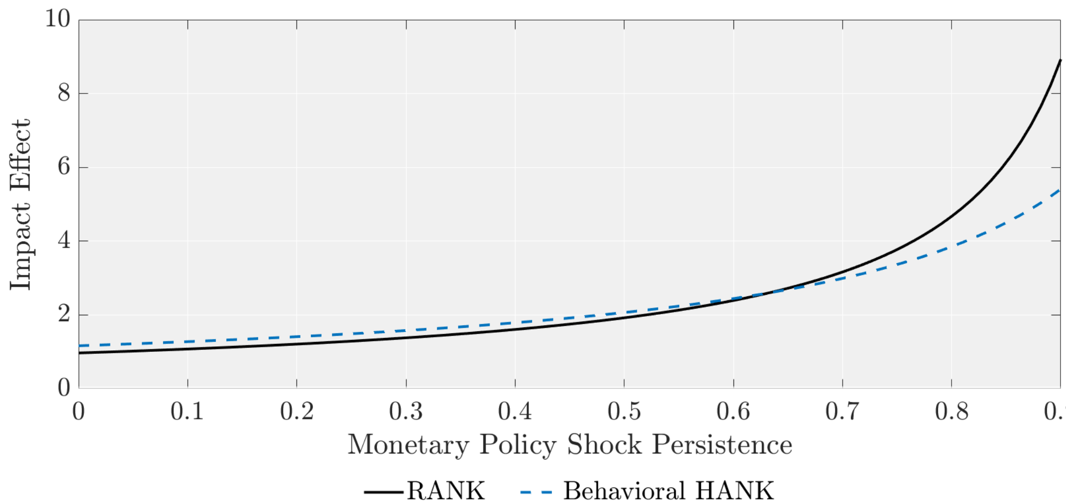

Note: This figure shows the initial output response to monetary policy shocks with different degrees of persistence.

As the persistence of the monetary policy shock approaches unity, the rational model leads to the paradoxical finding that an exogenous increase in the nominal interest rate leads to an expansion in output. To see this, note that we can write output as

$$\widehat{y}_t = -\frac{\frac{\psi_c}{\gamma}}{1 + \frac{\psi_c}{\gamma}\phi\kappa - \left(\psi_f + \psi_c \frac{\kappa}{\gamma}\right)\rho}\varepsilon_t^{MP}. \tag{59}$$

Given our baseline calibration and a Taylor coefficient of $\phi = 1$, the rational model would produce these paradoxical findings for $\rho > 0.967$. The behavioral HANK model, on the other hand, does not suffer from this as the denominator is always positive, even when $\phi = 0$ and $\rho = 1$.

D.4 Forward-Looking NKPC and Real Interest Rates

In the tractable model, we made the assumption that agents are rational with respect to real interest rates (as in Gabaix (2020)) and assumed a static Phillips Curve for simplicity. We now show that the results are barely affected when considering a forward-looking New Keynesian Phillips Curve (NKPC) and that agents are also boundedly rational with respect to real rates.

Gabaix (2020) derives the NKPC under bounded rationality and shows that it takes the form:

$$\pi_t = \beta M^f \mathbb{E}_t \pi_{t+1} + \kappa \widehat{y}_t,$$

with

$$M^f \equiv \bar{m} \left(\theta + \frac{1-\beta\theta}{1-\beta\theta\bar{m}} (1-\theta) \right),$$

where $1-\theta$ captures the Calvo probability of price adjustment.

Taking everything together (including the bounded rationality with respect to real interest rates), the model can be summarized by the following three equations:

$$\begin{aligned} \widehat{y}_t &= \psi_f \mathbb{E}_t \widehat{y}_{t+1} - \psi_c \frac{1}{\gamma} \left(\widehat{i}_t - \bar{m} \mathbb{E}_t \pi_{t+1} \right) \\ \pi_t &= \beta M^f \mathbb{E}_t \pi_{t+1} + \kappa \widehat{y}_t \\ \widehat{i}_t &= \phi \pi_t. \end{aligned}$$

Plugging the Taylor rule into the IS equation, we can write everything in matrix form:

$$\begin{pmatrix} \mathbb{E}_t \pi_{t+1} \\ \mathbb{E}_t \widehat{y}_{t+1} \end{pmatrix} = \underbrace{\begin{pmatrix} \frac{1}{\beta M^f} & -\frac{\kappa}{\beta M^f} \\ \frac{\psi_c}{\gamma \psi_f} \left(\phi - \frac{\bar{m}}{\beta M^f} \right) & \frac{1}{\psi_f} \left(1 + \frac{\psi_c \bar{m} \kappa}{\gamma \beta M^f} \right) \end{pmatrix}}_{\equiv A} \begin{pmatrix} \pi_t \\ \widehat{y}_t \end{pmatrix}. \tag{60}$$

For determinacy, we need

$$det(A) > 1; \quad det(A) - tr(A) > -1; \quad det(A) + tr(A) > -1.$$

The last condition is always satisfied. The first two conditions are satisfied if and only if

$$\phi > \max \left\{ \frac{\beta\delta M^f \bar{m} - 1}{\frac{\kappa}{\gamma} \frac{1-\lambda}{1-\lambda\chi}}, \bar{m} + \frac{(\delta\bar{m} - 1)(1-\beta M^f)}{\frac{\kappa}{\gamma} \frac{1-\lambda}{1-\lambda\chi}} \right\}.$$

In the case of a static Phillips curve but bounded rationality with respect to the real rate, the second condition is the crucial one. To capture the static Phillips curve, we can simply set $M^f = 0$. We can see that bounded rationality with respect to the real rate relaxes the determinacy condition whereas a forward-looking NKPC tightens it. But even in the case of a forward-looking NKPC (rational or behavioral), cognitive discounting relaxes the determinacy condition and thus, all our results from the static Phillips curve are qualitatively unchanged. Under our baseline calibration and $\theta = 0.375$ as in Gabaix (2020), the model still features determinacy under a peg, even when real interest rate expectations are rational (and therefore, also when they are behavioral).

E Quantitative Behavioral HANK Model

Table 3 shows how we calibrate the quantitative model introduced in Section 4. The calibration is quite standard in the literature. As in McKay et al. (2016), we assume that high productivity households pay all the taxes. In Section 5, we discuss less-progressive tax systems as well.

Patterson (2019) estimates that the unequal incidence—high-MPC households' incomes are more sensitive to aggregate income changes—leads to an amplification of 20% compared to the equal-incidence case. We target this moment by calibrating the dividend distribution accordingly. This implies that the highest productivity households receive 60% of the dividends, the middle-productivity households 40% and the low-productivity households 0. The fact that wealthier households (which tend to be the more productive households) receive a larger share of dividend income is consistent with the empirical findings in Kuhn et al. (2020).

We set the government debt level to match the quarterly MPCs of 0.16, as in Kaplan et al. (2018). This results in a debt-to-annual-GDP ratio of 50%. The rest of the calibration is as in McKay et al. (2016).

Table 3: Baseline Calibration Of Quantitative HANK Model

Parameter	Description	Value
R	Steady State Real Rate (annualized)	2%
γ	Risk aversion	2
φ	Inverse of Frisch elasticity	2
μ	Markup	1.2
θ	Calvo Price Stickiness	0.15
ρ_e	Autocorrelation of idiosyncratic risk	0.966
σ_e^2	Variance of idiosyncratic risk	0.033
$\tau(e)$	Tax shares	[0, 0, 1]
$d(e)$	Dividend shares	$[0, \frac{0.4}{0.5}, \frac{0.6}{0.25}]$
$\frac{B^G}{4Y}$	Government debt	0.5

E.1 Stability at the ELB and Fiscal Multipliers

Figure 13 shows the output and nominal interest rate response after a shock to the discount factor in the quantitative behavioral HANK model and in its rational counterpart. In particular, the discount factor jumps on impact by 0.8% for 12 quarters before it returns to steady state.

Figure 13: ELB recession in the quantitative behavioral HANK model

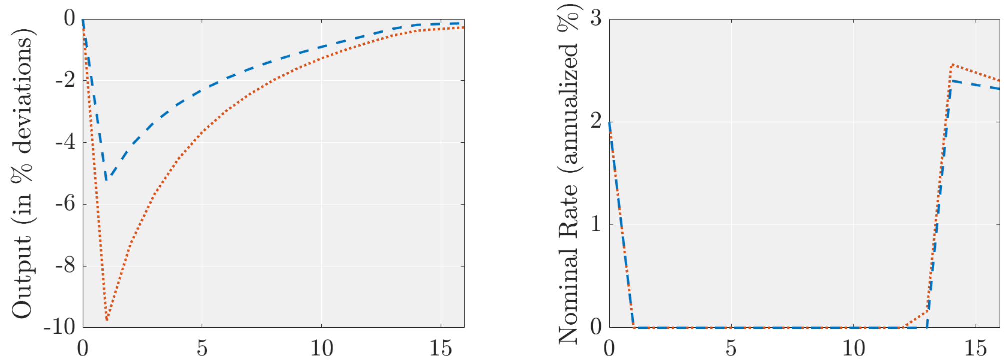

Note: This figure shows the impulse responses of total output and of the nominal interest rate after a discount factor shock that brings the economy to the ELB for 12 quarters.

We see that while the interest-rate path is quite similar across the two models, the output drop in the rational model is about twice as deep as in the behavioral HANK model. The intuition is as in the tractable model (Section 3). The binding ELB acts like a contractionary monetary policy shock because the nominal interest rate cannot keep up with the drop in the natural rate due to the ELB. Under rational expectations, households fully account for this and thus, cut back their consumption quite strongly on impact. Thus, the ELB leads to a large recession. Under cognitive discounting, on the other hand, households discount these future shocks and hence, decrease their consumption by less, leading to a milder recession.

Fiscal multiplier. To verify that the quantitative behavioral HANK model generates positive consumption multiplier under a constant real rate, we redo the experiments from Section D.1: the government exogeneously increases government consumption (which is assumed to be zero in steady state), which follows an AR(1)-process. The increase in government consumption is immediately financed by lump-sum taxes and households are taxed uniformly. Figure 14 shows the impact multiplier on consumption for various degrees of persistence, ρ_G. It shows that while the multiplier increases in persistence, it is bounded from below by zero. In other words, also the quantitative behavioral HANK model generates positive consumption multipliers, as documented empirically (Dupor et al. (2021)).

Figure 14: Consumption multiplier in the quantitative behavioral HANK

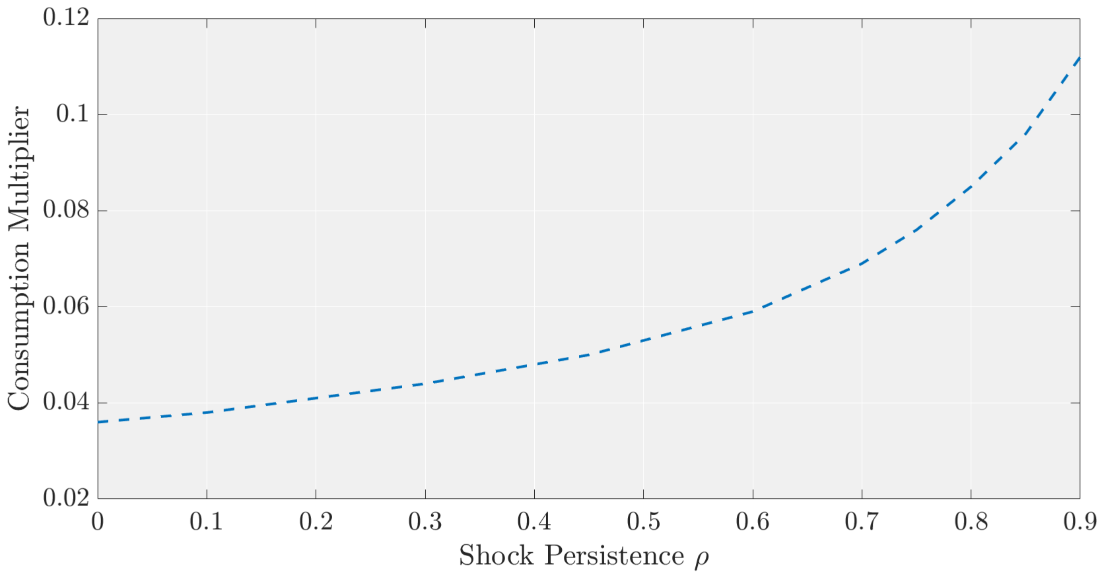

Note: This figure shows the impact consumption multiplier after an exogenous increase in government consumption which is financed by lump-sum taxes for various degrees of persistence.

E.1.1 Heterogeneous $\bar{m}$: Alternative Scenarios

The estimated differences in households' underreaction across different income groups are rather small. Nevertheless, one might argue that some agents (financial markets, for example) closely track what the Fed is doing and that they are usually well informed about its actions. To mirror this, we assume that the highest-productivity households are fully rational, i.e., their $\bar{m}$ is equal to 1. To keep the average $\bar{m}$ at 0.85, we then assume that the lowest-productivity households have a $\bar{m}$ of 0.7 and the middle-productivity households of 0.85.

The black-dashed-dotted line in Figure 15 shows the time zero output response (vertical axis) to an announced monetary policy shock taking place at different horizons (horizontal axis).

We see that forward guidance is more powerful than in the baseline calibration as the agents that tend to be more forward looking because they are not at their borrowing constraint are also more rational. Overall, however, our results remain robust. Even when the high-productivity households are fully rational and the other households have a cognitive discounting parameter of $\bar{m} = 0.85$ (such that the average $\bar{m}$ is above 0.85), the model resolves the forward guidance puzzle (not shown). Thus, even when a subpopulation of all households is fully rational, the behavioral HANK model can simultaneously generate amplification of conventional monetary policy through indirect effects and rule out the forward guidance puzzle.

Figure 15: Heterogeneous $\bar{m}$ and Monetary Policy

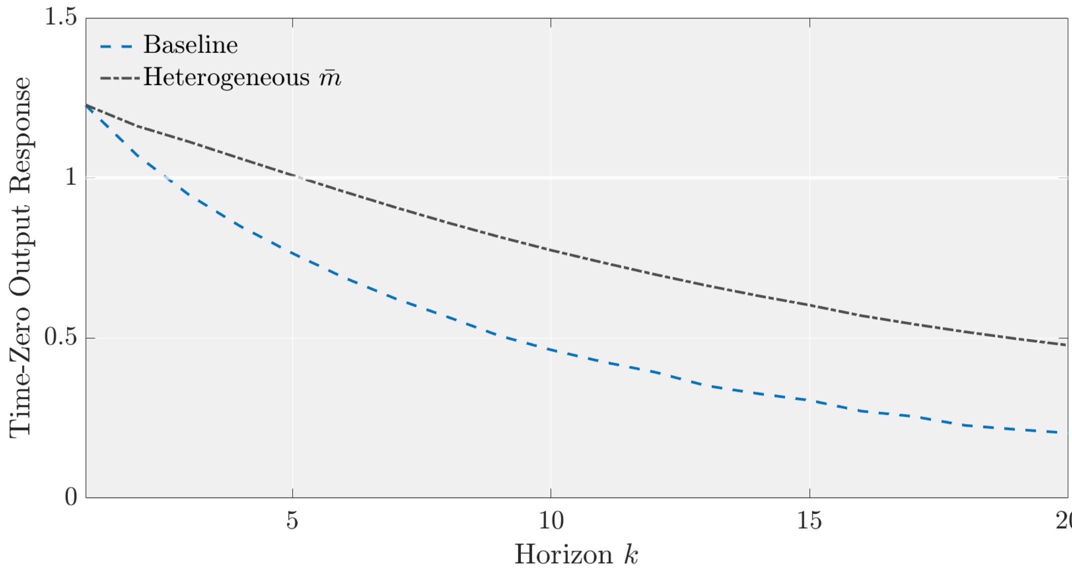

Note: This figure shows the response of total output in period 0 to anticipated i.i.d. monetary policy shocks occurring at different horizons k for the baseline calibration with $\bar{m} = 0.85$ for all households (blue-dashed line), and for the model in which high productivity households have $\bar{m} = 1$, medium-level productivity households have $\bar{m} = 0.85$ and low-productivity households have $\bar{m} = 0.7$ (black-dashed-dotted line).

F Additional Results and Figures to Section 5

F.1 Decomposition of Amplification Channel: Alternative Calibration

As shown in Section 5, household heterogeneity and cognitive discounting interact in such a way that productivity shocks get amplified through both features as well as their interaction. Given our baseline calibration, the interaction accounts for about 10% of the additional increase compared to RANK. We now consider an alternative calibration. We set the cognitive discounting parameter $\bar{m}$ to 0.6 instead of 0.85. Thus, somewhat closer to the lower bound of empirical estimates (see Section 2). Furthermore, we set the dividend shares in such a way that conventional, completely transitory, monetary policy shocks are amplified by 30% compared to the equal-incidence case, whereas our baseline calibration implies an amplification of 20%. This implies a dividend distribution where the lowest-productivity households do not receive any dividends, the medium-productivity households receive 30% and the high-productivity households 70% of the dividends. Figure 16 shows the decomposition of the additional amplification of negative productivity shocks under this alternative calibration.

Two things stand out. First, the overall additional increase is more twice as large. Given our discussion in Section 5, this is no surprise. Both—the underlying heterogeneity and cognitive discounting—induce a larger increase in inflation after a negative productivity shock. Given that they now both differ more from their counterparts in RANK, amplification becomes stronger. Second, the interaction becomes more important. In fact, the interaction alone accounts for more than the underlying heterogeneity itself. It amounts to more than 70% of the impact inflation response in RANK (1 percentage point) or about 25% of the *additional increase*.

Figure 16: Decomposition of the Additional Inflation Increase: Alternative Calibration

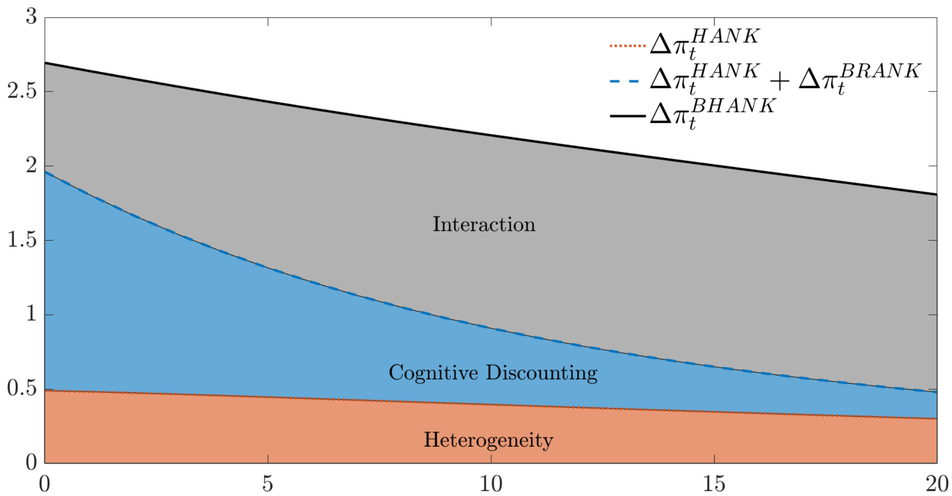

Note: This figure shows the additional inflation increase in the rational HANK model compared to the rational RANK model (orange-dotted line), the sum of the additional increase in the rational HANK and the behavioral RANK compared to the rational RANK model (blue-dashed line) and the additional increase in behavioral HANK model compared to the rational RANK model (black-solid line).

F.2 Procyclical Inequality

In response to a negative productivity shock, monetary policy has to increase nominal interest rates much more strongly according to the behavioral HANK model compared to the rational one. The reason is that future expected interest-rate hikes are less effective. The resolution of the forward guidance puzzle in the behavioral model illustrated this already in Section 4. However, also the rational HANK model can resolve the forward guidance puzzle when (counterfactually) higher-MPC households are on average *less* exposed to changes in aggregate income induced by monetary policy. The reason is that in expectation of a future nominal rate change, households increase their precautionary savings because if they would experience a decrease in their idiosyncratic productivity, they would benefit less from the overall increase in output. So, one might expect that the policy implications should be quite similar than in the behavioral model. This, however, is not the case.

Figures 17 and 18 illustrate this. The blue-dashed lines show the results from the behavioral HANK model (with countercyclical inequality, conditional on monetary policy) and the orange-dotted lines the results for the rational HANK model with procyclical inequality. From figure 17, we see that the cyclicality of inequality barely matters in the rational model when monetary policy fully stabilizes inflation (compare the orange-dotted lines with the ones in Figure 3). The reason is that by stabilizing inflation, the monetary authority closes the output gap and hence, shuts down the distributional effects explained above.

Furthermore, when monetary policy follows a Taylor rule and the economy overheats after a negative productivity shock, inequality increases strongly in the model with procyclical inequality (see Figure 18). The reason is that the positive output gap redistributes mainly to higher-productivity households, thus, increasing inequality. Put differently, the monetary authority does not really face a trade-off in that scenario, as by fully stabilizing inflation (and thus, closing the output gap) it also keeps inequality relatively low. This illustrates that accounting for the underlying heterogeneity matters for the policy prescriptions and it matters *how* we solve the forward guidance puzzle.

Figure 17: Procyclical inequality, strict inflation-targeting

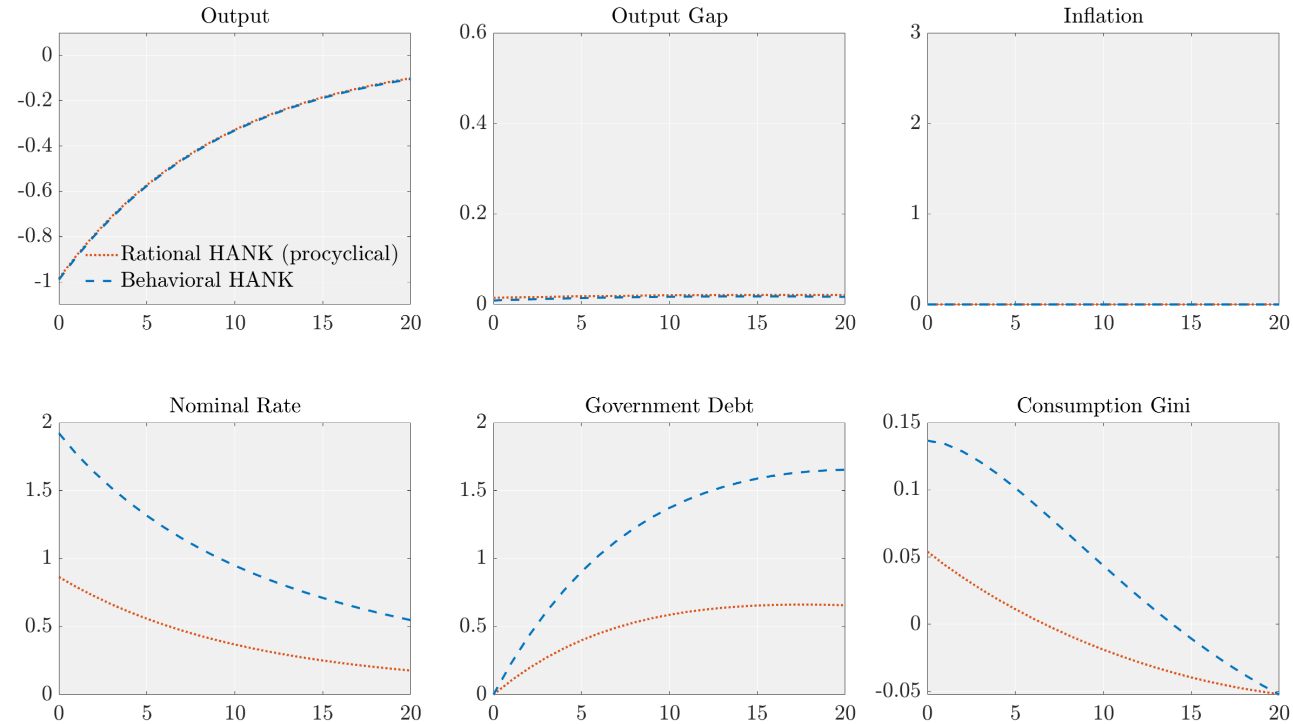

Note: This figure shows the impulse responses after a TFP shock that decreases potential output by 1% in the inflation-stabilizing monetary policy regime for the behavioral HANK model (blue-dashed lines) and for the rational HANK model with procyclical inequality (orange-dotted lines). Output and the output gap are shown as percentage deviations from steady state output, the nominal interest rate and inflation as annualized percentage points and the government debt level as percentage point deviations of the debt-per-annual GDP level. The lower-right figure shows the change in the consumption Gini index as a percentage deviation from the stationary equilibrium.

Figure 18: Procyclical inequality, Taylor rule

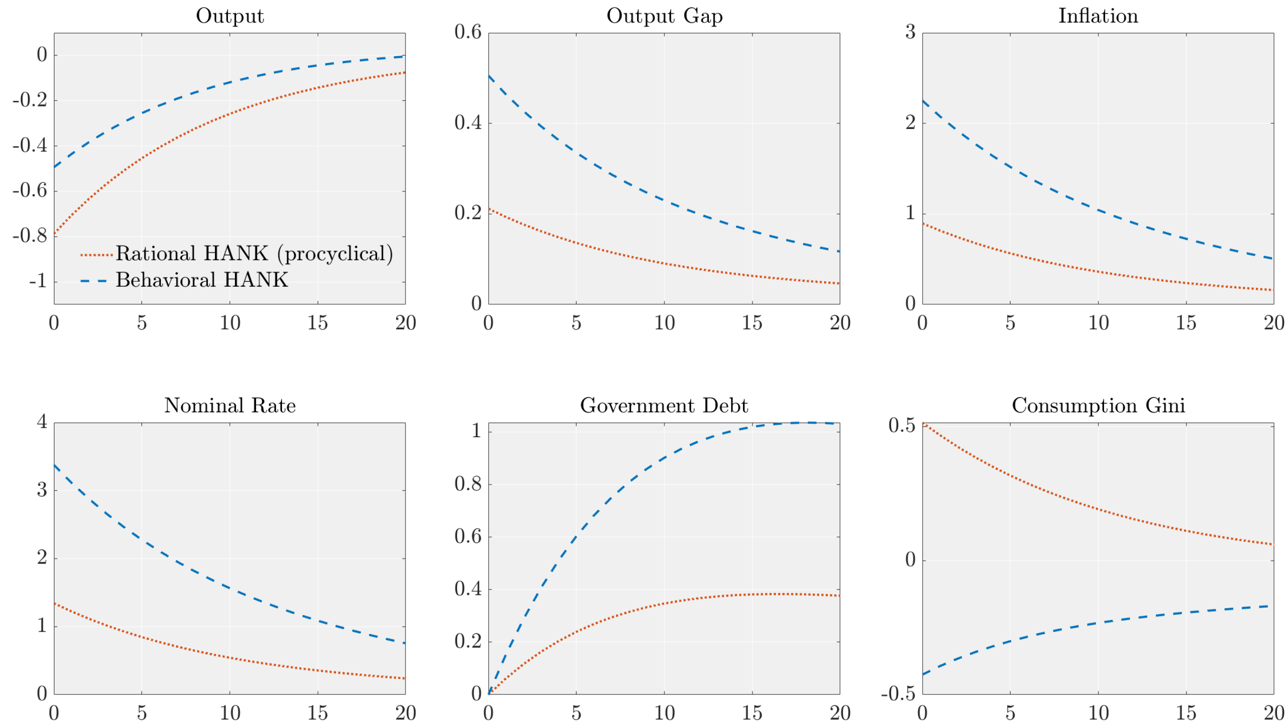

Note: This figure shows the impulse responses after a TFP shock that decreases potential output by 1% when monetary policy follows a Taylor rule for the behavioral HANK model (blue-dashed lines) and for the rational HANK model with procyclical inequality (orange-dotted lines). Output and the output gap are shown as percentage deviations from steady state output, the nominal interest rate and inflation as annualized percentage points and the government debt level as percentage point deviations of the debt-per-annual GDP level. The lower-right figure shows the change in the consumption Gini index as a percentage deviation from the stationary equilibrium.

F.3 Representative Household Models

In this section, we show the results for the same policy experiments when abstracting from household heterogeneity. The behavioral RANK model is exactly the same as in Gabaix (2020), but calibrated to be consistent with the behavioral HANK model. In particular, risk aversion and the inverse labor elasticity are set to 2, the Calvo probability of a price adjustment is set to 0.15, the steady state real interest rate is 2% (annualized), the shock persistence is 0.9 and the shock size is set in order to decrease potential output by 1% on impact.

Figures 19 and 20 show the results for the case in which monetary policy fully stabilizes inflation and in which case it follows a simple Taylor rule, respectively. We see that the policy implications differ from the behavioral HANK model. First of all, there is no trade-off between price stability and inequality, as there is no inequality in models with a representative household. Second, when monetary policy follows the same Taylor rule in these models, the spike in inflation is less pronounced than in the behavioral HANK model. The reason is that both features in our model—the underlying heterogeneity and cognitive discounting—amplify the overheating of the economy following a negative supply shock with a standard Taylor rule. Due to the heterogeneous exposure of households, higher-MPC households benefit relatively more from the overheating of the economy, which further reinforces the increase in inflation. Cognitive discounting dampens the response of the economy to the monetary-policy reaction, further fueling the inflation increase.

Figure 19: Representative Household, strict inflation targeting

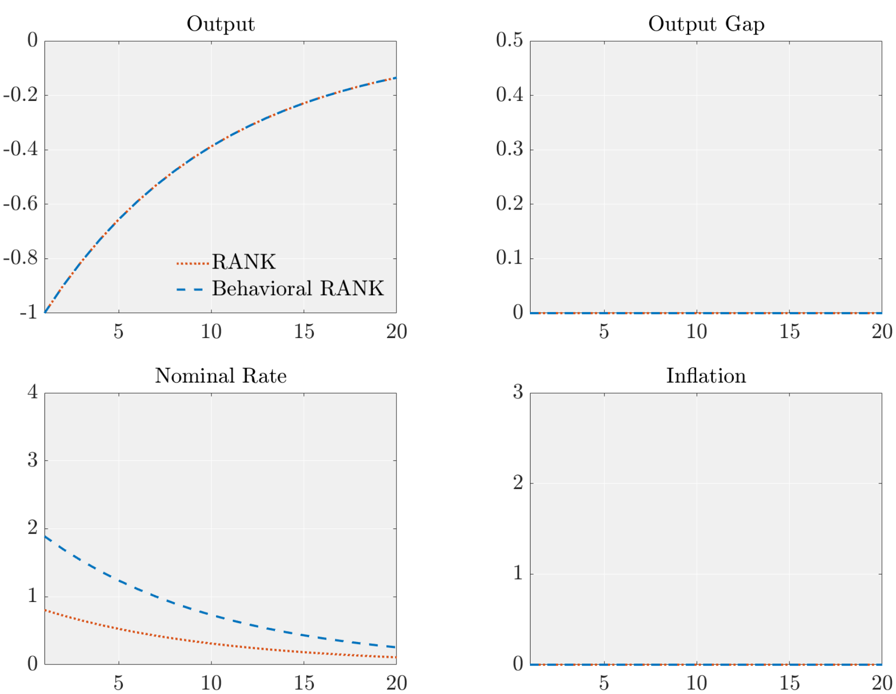

Note: This figure shows the impulse responses after a TFP shock that decreases potential output by 1% in the inflation-stabilizing monetary policy regime for the behavioral RANK model (blue-dashed lines) and for the rational RANK model (orange-dotted lines). Output and the output gap are shown as percentage deviations from steady state output, the nominal interest rate and inflation as annualized percentage points.

Figure 20: Representative Household, Taylor rule

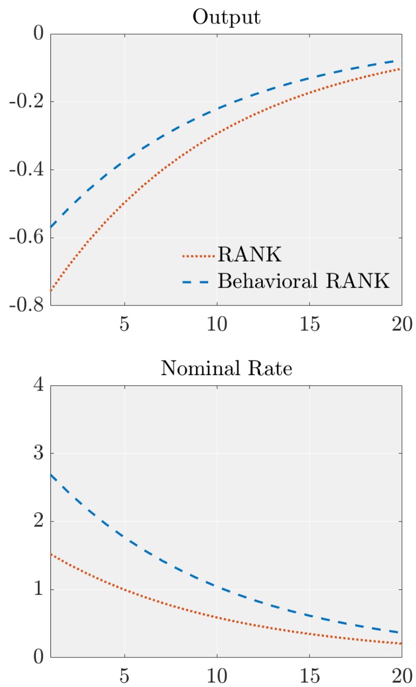

Note: This figure shows the impulse responses after a TFP shock that decreases potential output by 1% when monetary policy follows a Taylor rule for the behavioral RANK model (blue-dashed lines) and for the rational RANK model (orange-dotted lines). Output and the output gap are shown as percentage deviations from steady state output, the nominal interest rate and inflation as annualized percentage points.

F.4 Heterogeneous $\bar{m}$

In the main text, we consider the case of homogeneous cognitive discounting. As discussed in Section 4.2, the data however suggests that richer households tend to exhibit a somewhat smaller behavioral bias. If we incorporate heterogeneous $\bar{m}$ along the lines described in Section 4.2, we find that the trade-off between price stability and distributional considerations is slightly weaker.

Figures 21 and 22 show this graphically. The black-dashed-dotted lines show the impulse-response functions for the case with heterogenous $\bar{m}$. We see that inequality, at least on impact, rises less with heterogeneous $\bar{m}$ in the regime in which monetary policy fully stabilizes inflation compared to the homogeneous $\bar{m}$ case.

The reason is that high-productivity households tend to be more likely to be away from the borrowing constraint and as these households are now closer to rational expectations, they respond more strongly to expected future higher interest rates. Therefore, monetary policy has to react slightly less than in the case with homogeneous degrees of cognitive discounting. Overall, however, the differences are small and the trade-off that arises due to households' cognitive discounting remains substantial.

Figure 21: Heterogeneous $\bar{m}$, strict inflation-targeting

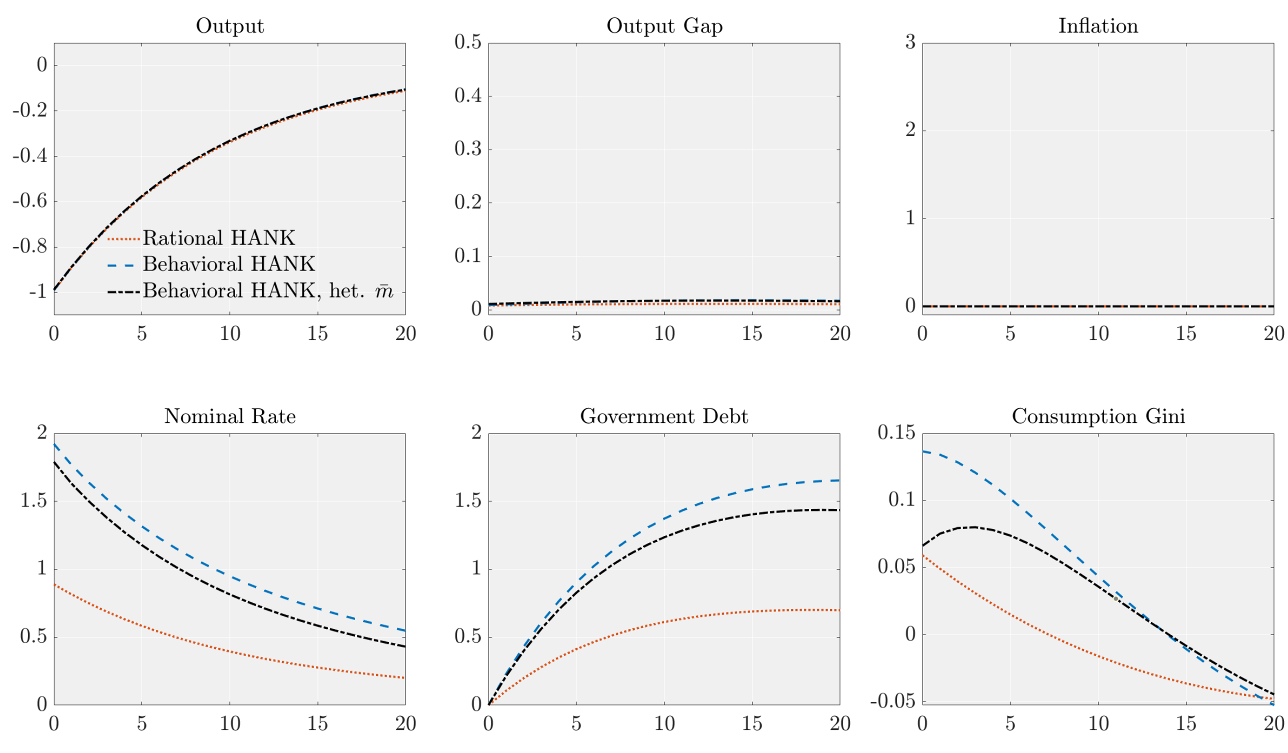

Note: This figure shows the impulse responses after a TFP shock that decreases potential output by 1% in the inflation-stabilizing monetary policy regime and heterogeneous $\bar{m}$ (black-dashed-dotted lines). Output and the output gap are shown as percentage deviations from steady state output, the nominal interest rate and inflation as annualized percentage points and the government debt level as percentage point deviations of the debt-per-annual GDP level. The lower-right figure shows the change in the consumption Gini index as a percentage deviation from the stationary equilibrium.

Figure 22: Heterogeneous $\bar{m}$, Taylor rule

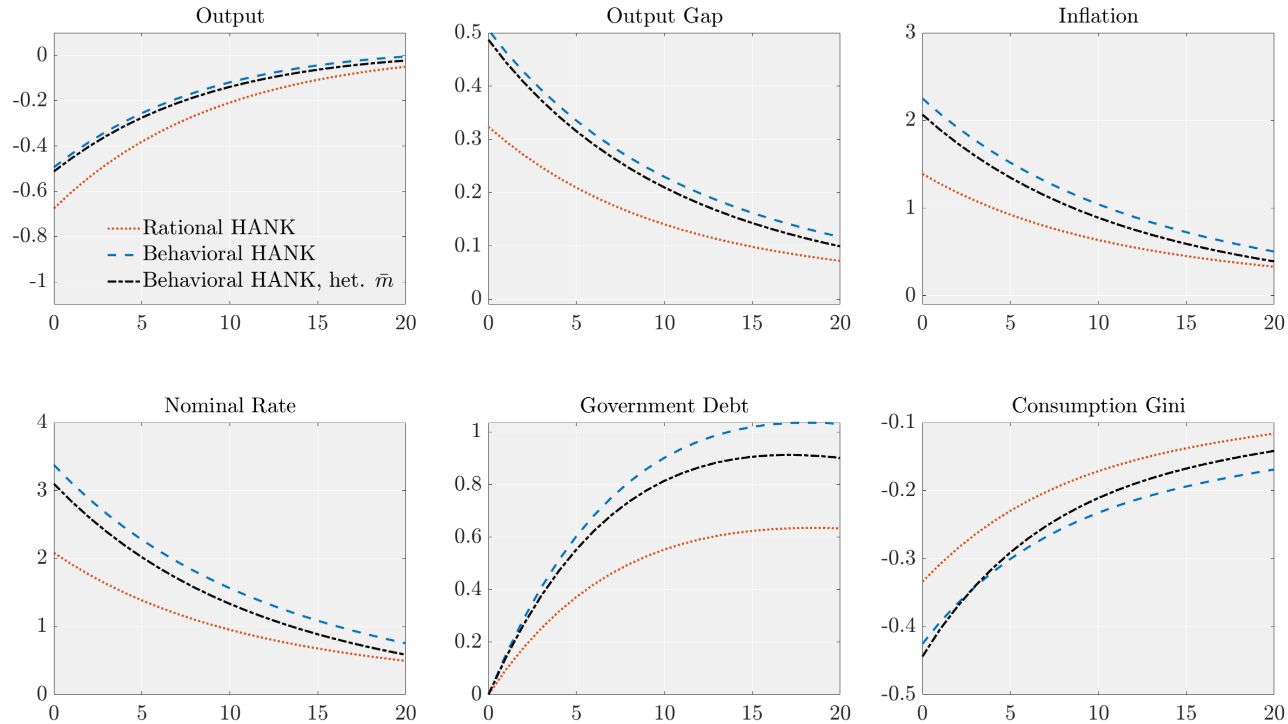

Note: This figure shows the impulse responses after a TFP shock that decreases potential output by 1% when monetary policy follows a Taylor rule and heterogeneous $\bar{m}$ (black-dashed-dotted lines).. Output and the output gap are shown as percentage deviations from steady state output, the nominal interest rate and inflation as annualized percentage points and government debt level as percentage point deviations of the debt-per-annual GDP level. The lower-right figure shows the change in the consumption Gini index as a percentage deviation from the stationary equilibrium.

F.5 Different Tax Systems

In the main text, we have assumed that only the most productive households pay taxes. We now examine how the predictions from Section 5 change if, instead, all households pay taxes, according to their productivity. Recall the budget constraint of a household with idiosyncratic productivity e:

$$C_{i,t} + \frac{B_{i,t+1}}{R_t} = B_{i,t} + W_t e_{i,t} N_{i,t} + D_t d(e) - \tau_t(e).$$

In the baseline calibration, we assume that $\tau_t(e) = \frac{T_t}{0.25}$ if $e = e_{high}$ and 0 otherwise (0.25 denotes the share of high-productivity households that actually pay taxes). To model *less progressive taxes*, we assume that $\tau_t(e) = e_t \frac{T_t}{\bar{e}}$.

Figures 23 and 24 show the impulse-response functions in that scenario for the case in which monetary policy fully stabilizes inflation and in which it follows a simple Taylor rule, respectively. The blue-dashed lines show the results for the behavioral HANK model, the orange-dotted lines for the rational HANK model and the black-dashed-dotted lines for the behavioral model with heterogeneous $\bar{m}$.

We see that qualitatively the results remain unchanged: there is a strong trade-off in the behavioral model between price stability and aggregate efficiency on the one hand and fiscal sustainability and inequality on the other hand. The implications for inequality, however, are even stronger than in the baseline fiscal regime. As the monetary authority raises interest rates more strongly in the behavioral model, the spillovers to fiscal policy are larger. Therefore, the tax system or the design of fiscal policy becomes more important.

Therefore, the increase in inequality in the full-inflation-stabilization case is much more pronounced in the behavioral model when taxes are less progressive. Furthermore, inequality keeps increasing for quite some time. The reason is that more productive households tend be less borrowing constrained and thus, adjust their consumption on impact in expectation of higher future taxes. Households at the borrowing constraint, however, reduce their consumption once taxes actually increase. As these households tend to consume relatively little, consumption inequality increases over time as taxes increase. With heterogeneous degrees of cognitive discounting, these effects become even more pronounced.

Figure 23: Heterogeneous $\bar{m}$, strict inflation-targeting and less progressive taxes

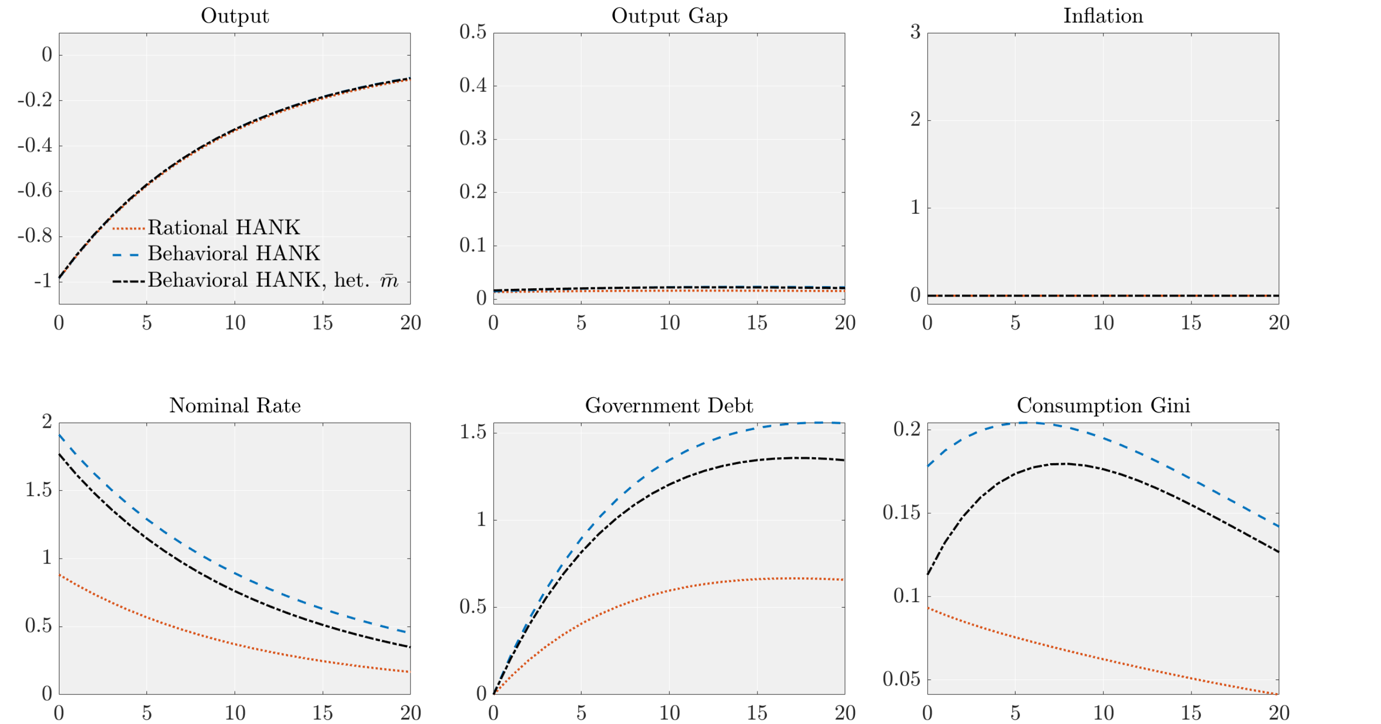

Note: This figure shows the impulse responses after a TFP shock that decreases potential output by 1% in the inflation-stabilizing monetary policy regime with a less-progressive tax system than in the baseline calibration. Output and the output gap are shown as percentage deviations from steady state output, the nominal interest rate and inflation as annualized percentage points and the government debt level as percentage point deviations of the debt-per-annual GDP level. The lower-right figure shows the change in the consumption Gini index as a percentage deviation from the stationary equilibrium.

Figure 24: Heterogeneous $\bar{m}$, Taylor rule and less progressive taxes

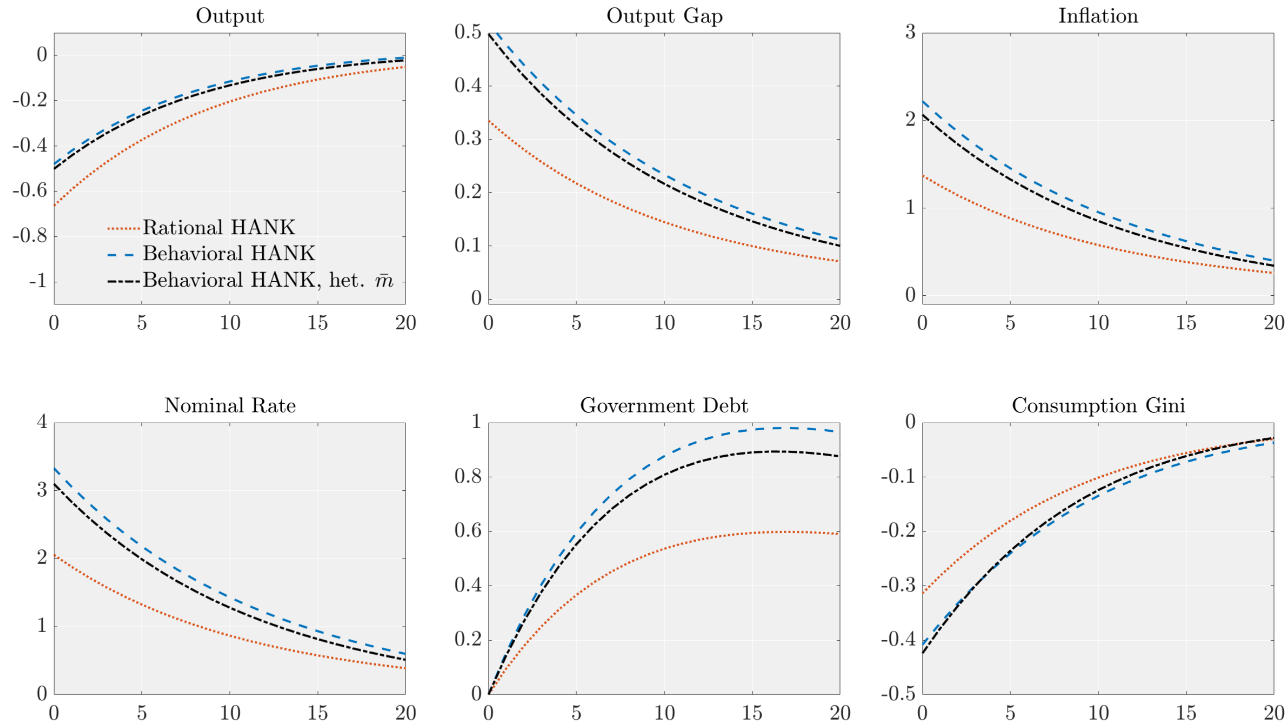

Note: This figure shows the impulse responses after a TFP shock that decreases potential output by 1% when monetary policy follows a Taylor rule with a less-progressive tax system than in the baseline calibration. Output and the output gap are shown as percentage deviations from steady state output, the nominal interest rate and inflation as annualized percentage points and the government debt level as percentage point deviations of the debt-per-annual GDP level. The lower-right figure shows the change in the consumption Gini index as a percentage deviation from the stationary equilibrium.

F.6 Productivity Shock with High Initial Debt Levels

Figure 25 shows the impulse responses after a negative productivity shock when monetary policy fully stabilizes inflation when the initial debt level is 90% of annual GDP instead of 50%.

Compared to Figure 3, we see that government debt increases more strongly when initial debt levels are higher. The increase in the real interest rate is more costly for the government which is financed by an additional increase in debt. Consumption inequality also increases more than at lower levels of government debt, even tough the differences are rather small.

Overall, higher government debt levels exacerbate the trade-off the monetary authority faces as the fiscal implications of fighting inflation become more severe at higher debt levels.

Figure 25: Inflationary productivity shock with high debt: strict inflation targeting

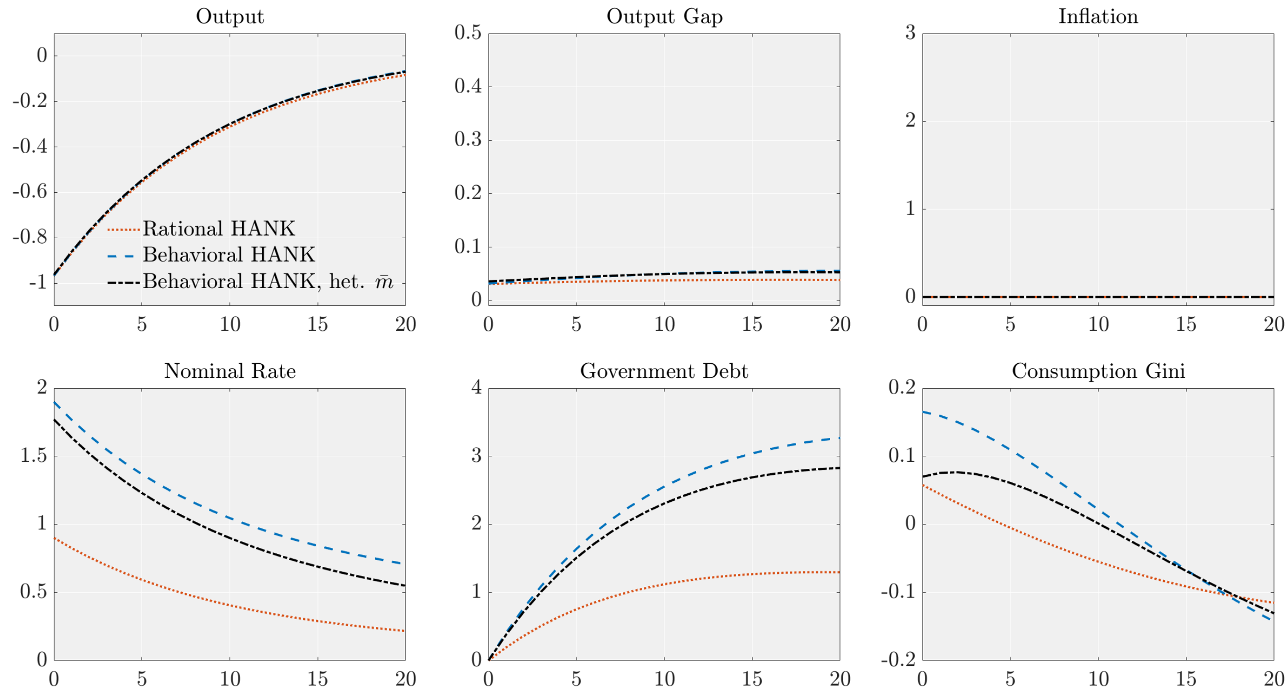

Note: This figure shows the impulse responses after a TFP shock that decreases potential output by 1% in the inflation-stabilizing monetary policy regime when the initial debt level is 90% instead of 50% of annual GDP. Output and the output gap are shown as percentage deviations from steady state output, the nominal interest rate and inflation as annualized percentage points and the government debt level as percentage point deviations of the debt-per-annual GDP level. The lower-right figure shows the change in the consumption Gini index as a percentage deviation from the stationary equilibrium.

Figure 26: Inflationary productivity shock with high debt: Taylor rule

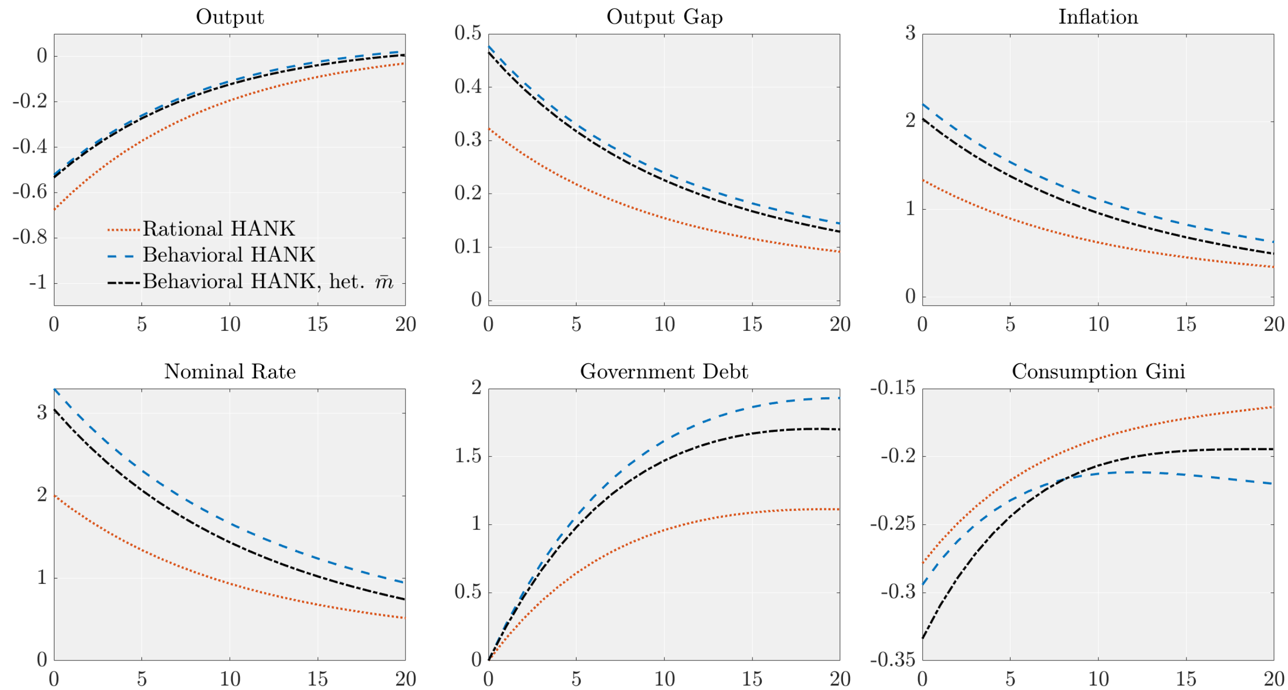

Note: This figure shows the impulse responses after a TFP shock that decreases potential output by 1% in the inflation-stabilizing monetary policy regime when the initial debt level is 90% instead of 50% of annual GDP. Output and the output gap are shown as percentage deviations from steady state output, the nominal interest rate and inflation as annualized percentage points and the government debt level as percentage point deviations of the debt-per-annual GDP level. The lower-right figure shows the change in the consumption Gini index as a percentage deviation from the stationary equilibrium.

F.7 Cost-Push Shocks

We now show that the fiscal and monetary implications are very similar for an inflationary cost-push shock. To introduce cost-push shocks, we assume that the desired mark-up of firms, μ_t follows an AR(1)-process, $\mu_t = (1-\rho_\mu)\bar{\mu} + \rho_\mu \mu_{t-1} + \varepsilon_t^\mu$, where ε_t^μ is an i.i.d. shock, $\bar{\mu}$ the steady-state level of the desired markup and ρ_μ the persistence of the shock process which we set to $\rho_\mu = 0.9$. The rest of the model is as in Section 5. Note, that the shock also applies to the model under flexible prices, thus moves potential output as well.

Figure 27 shows the impulse-response functions of output, the output gap, inflation, nominal interest rates, government debt and the consumption Gini index as a measure of consumption inequality following an inflationary cost-push shock. The blue-dashed lines and black-dashed-dotted lines show the responses in the behavioral HANK model with homogeneous and heterogeneous $\bar{m}$, respectively, and the orange-dotted lines in the rational HANK model. In all cases, monetary policy fully stabilizes inflation by assumption. Output drops, with the responses being practically identical across the two models. Again, the output gap is practically closed in all models. The required response of the nominal interest rate, however, differs substantially across the behavioral and the rational models, as was the case after a negative productivity shock, discussed in Section 5. In the behavioral HANK the monetary authority increases the nominal rate much more strongly and more persistently. The reason for this strong response is that households cognitively discount future (expected) interest rate hikes making them less effective for stabilizing inflation today. Thus, in order to achieve the same stabilization outcome in every period, the interest rate needs to increase by more.

Figure 27: Inflationary cost-push shock: strict inflation targeting

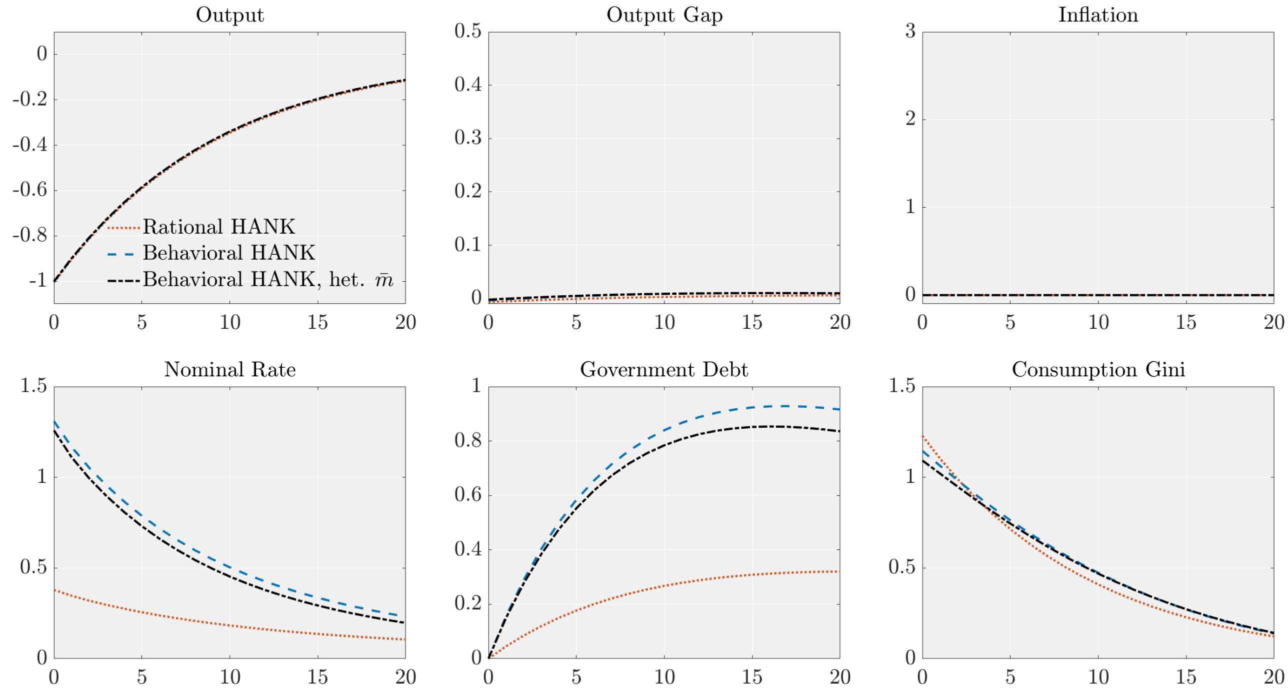

Note: This figure shows the impulse responses after a cost-push shock that decreases potential output by 1% in the inflation-stabilizing monetary policy regime. Output and the output gap are shown as percentage deviations from steady state output, the nominal interest rate and inflation as annualized percentage points and the government debt level as percentage point deviations of the debt-per-annual GDP level. The lower-right figure shows the change in the consumption Gini index as a percentage deviation from the stationary equilibrium.

Increasing the interest rates increases the cost of debt for the government which it finances in the short run by issuing more debt. The middle panel on the bottom line in Figure 27 shows that government debt in the behavioral model increases more than three times as much as in the rational model. Furthermore, consumption inequality increases in both models, somewhat stronger in the rational model. There are two channels: first and most important, the cost-push shock increases dividends and decreases wages which redistributes from low to high productivity households thereby pushing up consumption inequality. Second, the increase in the real interest rate redistributes towards high wealth households but it is the high productivity households who eventually pay the tax burden. This slightly decreases the consumption of high productivity households and increases the consumption of middle productivity households who hold some assets but do not face tax increases. Thus, the second channel slightly dampens the increase in inequality and, as real interest rates increase by more, this channel is stronger in the behavioral HANK model.

Figure 28 shows the impulse-response functions of output, the output gap, inflation, nominal interest rates, government debt (as a share of annual GDP) and consumption inequality for the same cost-push shock but for the case in which monetary policy follows a simple Taylor rule with a response coefficient of 1.5.

As in the case where monetary policy fully stabilizes inflation, the nominal interest rate increases more strongly in the behavioral HANK model than in its rational version. The difference across the two models, however, is somewhat smaller compared to the case in which inflation is completely stable. Inflation, however, increases more strongly in the behavioral model and also government debt increases more substantially.

Consumption inequality increases less strongly than with fully stabilizing inflation. The overheating economy—reflected in the positive output gap and increase in inflation—increases wages and decreases profits (relative to the inflation stabilizing regime) in the same way as expansionary policy shocks in Sections 3 and 4 do, thereby redistributing towards lower income households which dampens the increase in consumption inequality.

Figure 28: Inflationary cost-push shock: Taylor rule

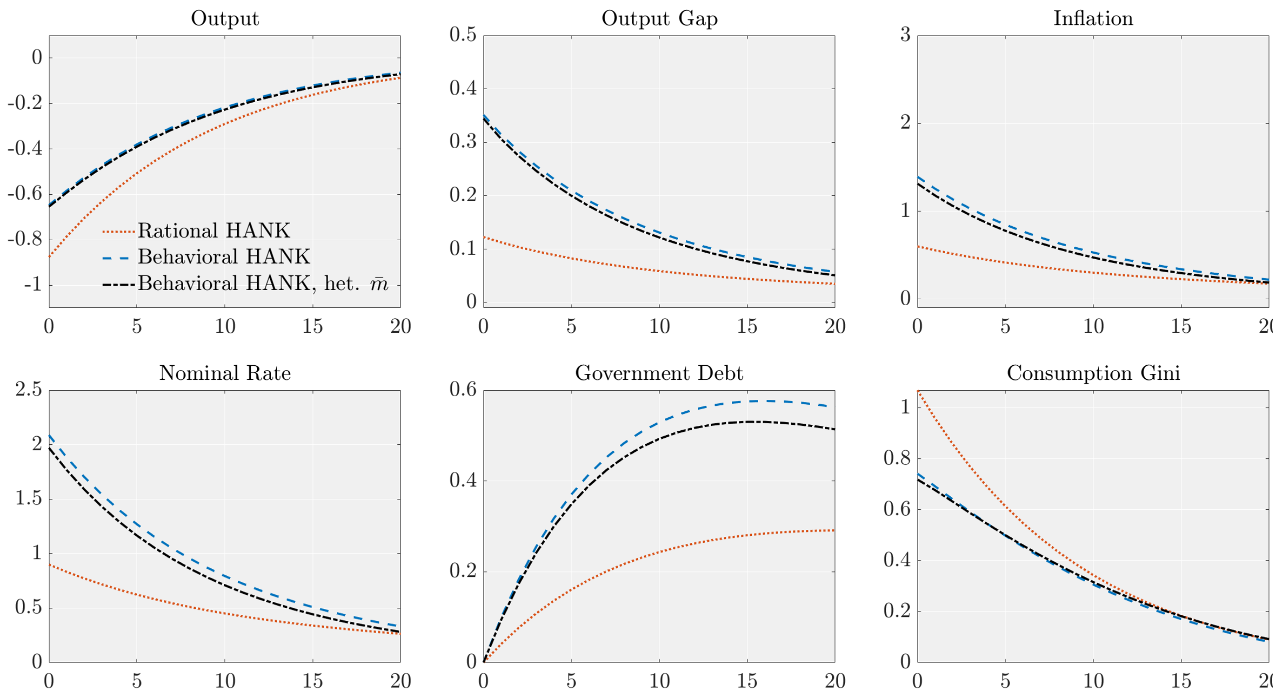

Note: This figure shows the impulse responses after a cost-push shock that decreases potential output by 1% in the Taylor rule monetary policy regime. Output and the output gap are shown as percentage deviations from steady state output, the nominal interest rate and inflation as annualized percentage points and the government debt level as percentage point deviations of the debt-per-annual GDP level. The lower-right figure shows the change in the consumption Gini index as a percentage deviation from the stationary equilibrium.

G Details on Intertemporal MPCs

In this section, we derive the iMPCs discussed in Section 6.1. Defining Y_t^j as type j's disposable income, we can write the households' budget constraints as

$$C_t^H = Y_t^H + \frac{1-s}{\lambda} R_t B_t$$
$$C_t^U + \frac{1}{1-\lambda} B_{t+1} = Y_t^U + \frac{s}{1-\lambda} R_t B_t,$$

where R_t denotes the real interest rate and B_t real bonds. Log-linearizing the two budget constraints around the zero-liquidity steady state and $R = \beta^{-1}$ yields

$$\widehat{c}_t^H = \widehat{y}_t^H + \frac{1-s}{\lambda} \beta^{-1} b_t \tag{61}$$
$$\widehat{c}_t^U + \frac{1}{1-\lambda} b_{t+1} = \widehat{y}_t^U + \frac{s}{1-\lambda} \beta^{-1} b_t, \tag{62}$$

where b_t denotes real bonds in shares of steady state output. Aggregating (61) and (62) delivers

$$\widehat{c}_t = \tilde{y}_t + \beta^{-1} b_t - b_{t+1}, \tag{63}$$

where $\tilde{y}_t$ denotes aggregate disposable income.

By plugging equations (61) and (62) into the Euler equation of unconstrained households (12), we can derive the dynamics of liquid assets b_t (ignoring changes in the real rate as this is a partial equilibrium exercise):

$$\mathbb{E}_t b_{t+2} - b_{t+1} \left[\frac{1}{s\bar{m}} + \beta^{-1} s + \frac{(1-s)^2 \beta^{-1} (1-\lambda)}{s\lambda} \right] + \frac{\beta^{-1}}{\bar{m}} b_t =$$
$$(1-\lambda) \mathbb{E}_t \widehat{y}_{t+1}^U + \frac{1-s}{s} (1-\lambda) \mathbb{E}_t \widehat{y}_{t+1}^H - \frac{1-\lambda}{s\bar{m}} \widehat{y}_t^U. \tag{64}$$

Note that a change in total disposable income by one changes the hand-to-mouth households' disposable income by χ and the disposable income of unconstrained households by $\frac{1-\lambda\chi}{1-\lambda}$.

Let us denote the right-hand side of equation (64) by $-\mathbb{E}_t \widehat{z}_t$. Factorizing the left-hand side and letting F denote the forward-operator, it follows that

$$(F - \mu_1)(F - \mu_2) \mathbb{E}_t b_t = -\mathbb{E}_t \widehat{z}_t, \tag{65}$$

where μ_1 and μ_2 denote the roots of the characteristic equation

$$\mathbb{E}_t b_{t+2} - \phi_1 b_{t+1} - \phi_2 b_t = 0, \tag{66}$$

where

$$\phi_1 \equiv \left[\frac{1}{s\bar{m}} + \beta^{-1} s + \frac{(1-s)^2 \beta^{-1} (1-\lambda)}{s\lambda} \right] \tag{67}$$

and

$$\phi_2 \equiv -\frac{\beta^{-1}}{\bar{m}}. \tag{68}$$

Thus, the roots are given by

$$\mu_{1,2} = \frac{\phi_1 \pm \sqrt{\phi_1^2 + 4\phi_2}}{2}. \tag{69}$$

It follows that

$$b_{t+1} = \mu_1 b_t - (F - \mu_2)^{-1}\mathbb{E}_t \widehat{z}_t$$
$$= \mu_1 b_t + \frac{\mu_2^{-1}}{1 - F\mu_2^{-1}}\mathbb{E}_t \widehat{z}_t.$$

Note that $\mathbb{E}_t \widehat{z}_t$ can be written as $\frac{1-\lambda\chi}{s}\left(\delta \mathbb{E}_t \widehat{y}_{t+1} - \frac{1}{\bar{m}}\widehat{y}_t\right)$. Without loss of generality, we let $\mu_2 > \mu_1$ and we have $\mu_2 > 1$. We have $(1 - F\mu_2^{-1})^{-1} = \sum_{l=0}^{\infty} \mu_2^{-l} F^l$. Thus, we end up with

$$b_{t+1} = \mu_1 b_t + \frac{1-\lambda\chi}{s}\sum_{l=0}^{\infty}\mu_2^{-(l+1)}\mathbb{E}_t\left(\frac{1}{\bar{m}}\widehat{y}_{t+l} - \delta\widehat{y}_{t+1+l}\right). \tag{70}$$

Plugging this in equation (63) and taking derivatives with respect to $\widehat{y}_{t+k}$ yields Proposition 4.

iMPCs for more than two periods. Figure 29 plots the MPCs for the year of the income windfall as well as the five consecutive years for different degrees of rationality. As discussed in section 6.1, under our benchmark calibration, the rational model predicts somewhat larger initial MPCs as behavioral, unconstrained households save relatively more. Over time, however, the MPCs in the behavioral model lie above their rational counterparts due to the fact that more and more of the initial unconstrained households become hand-to-mouth and start consuming their (higher) savings. As Figure 30 shows, the probability of type switching, $1 - s$, matters for when exactly the behavioral model starts to generate larger MPCs compared to the rational model.

Figure 29: Intertemporal MPCs

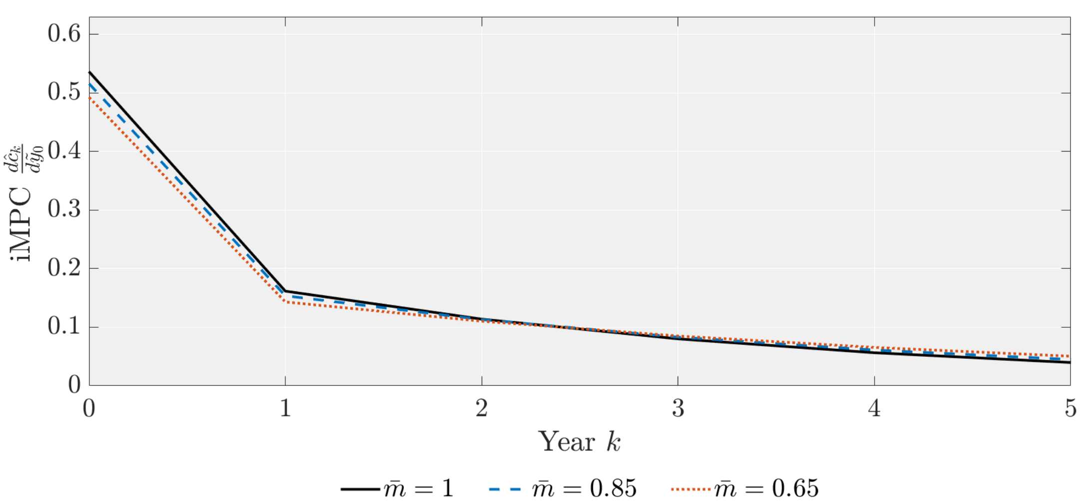

Note: This figure shows the aggregate intertemporal MPCs, i.e., the aggregate consumption response in year k to a change in aggregate disposable income in year 0 for different $\bar{m}$.

iMPCs and the role of idiosyncratic risk. In Figure 30, we plot he MPCs in the year of the income windfall (left panel) and the first year after the windfall (right panel) for a relatively high idiosyncratic risk of $1 - s = 0.5$. The high probability of becoming hand-to-mouth flips the role of $\bar{m}$ for the MPC_1 compared to our baseline calibration as discussed in Section 6.1.

Figure 30: Intertemporal MPCs, Bounded Rationality and Household Heterogeneity

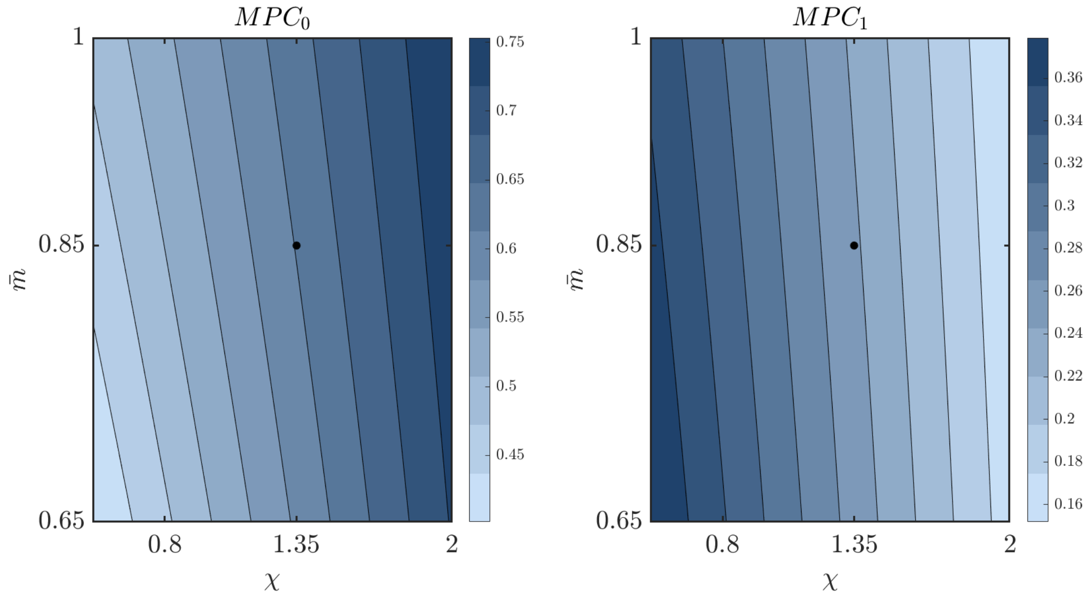

Note: This figure shows the aggregate intertemporal MPCs, i.e., the aggregate consumption response in year 0 (left) and year 1 (right) to a change in aggregate disposable income in year 0 for a transition probability $1 - s = 0.5$.

The reason being that the behavioral, unconstrained households save a relatively large amount of the received income windfall in period 0 as they cognitively discount the decrease in their future marginal utility. Thus, they end up with relatively more disposable income in year 1. Now, given the relatively high probability of type switching, there are many unconstrained households who end up being hand-to-mouth in year 1 after the income windfall. As they are hand-to-mouth, they consume their previously-accumulated savings which increases the MPC_1. The more behavioral unconstrained households are, i.e., the lower $\bar{m}$ is, the more pronounced this effect and hence, a lower $\bar{m}$ increases the MPC_1 in the case of a relatively high $1 - s$.

H Sticky Wages

In this section, we provide details on the sticky-wage extension presented in Section 6.2 as well as the calibration used to produce Figures 7 and 8. The way we introduce sticky wages follows Colciago (2011) and recently adopted by Bilbiie et al. (2021).[25]

In the household block, the only difference to our benchmark model is that we assume that there is a labor union pooling labor and setting wages on behalf of households. This leads to a condition similar to the labor-leisure conditions in Section 2. But instead of individual conditions, the condition is the same for every household:

$$\varphi \widehat{n}_t = \widehat{w}_t - \gamma \widehat{c}_t,$$

and $\widehat{n}_t = \widehat{n}_t^U = \widehat{n}_t^H$.

The labor union, however, is subject to wage rigidities. The nominal wage can only be re-optimized with a constant probability, which leads to a time-varying wage markup

$$\widehat{\mu}_t^w = \varphi \widehat{n}_t - \widehat{w}_t + \gamma \widehat{c}_t,$$

and a wage Phillips Curve

$$\pi_t^w = \beta \mathbb{E}_t \pi_{t+1}^w + \kappa_w \widehat{\mu}_t^w.$$

Wage inflation is given by

$$\pi_t^w = \widehat{w}_t - \widehat{w}_{t-1} + \pi_t.$$

The firm side is exactly the same as in the main text but we focus on the case with rational firms, which gives rise to a standard Phillips Curve:

$$\pi_t = \beta \mathbb{E}_t \pi_{t+1} + \kappa_\pi \widehat{mc}_t,$$

where $\widehat{mc}_t$ is a time-varying price markup. Table 4 summarizes all equilibrium equations.

The calibration of this extended model is presented in Table 5. The parameters γ, φ, s, β and $\bar{m}$ are as in our baseline calibration. The parameters of the Taylor rule, ρ_i and ϕ, are set as estimated in Auclert et al. (2020).

The slope of the wage Phillips curve, κ_w, is set as in Bilbiie et al. (2021) and we focus on the *no-redistribution* case $\mu^D = 0$. Note, that this leads to impact responses of consumption of the two household types that are very close to the ones in our baseline model: $\widehat{c}_t^H$ increases by about 1.4, whereas output increases by 1. The baseline calibration of $\chi = 1.35$ would predict that in the model without sticky wages, $\widehat{c}_t^H$ increases by 1.35 when output increases by 1. We focus on a relatively stable inflation and set κ_π to 0.01.

The only parameter that we change with respect to our baseline calibration is λ which we set to 0.37 instead of 0.33. A value of 0.37 is still in the range of often used values (see, for example Bilbiie (2020)). We increase λ somewhat compared to our baseline calibration in order to increase

[25] See also Erceg et al. (2000). Broer et al. (2020) and Broer et al. (2021b) discuss the role of sticky wages in (rational) TANK models for the analysis of monetary and fiscal policy, respectively.

Table 4: Sticky Wages, Equilibrium Equations

Name	Equation
Wage Markup	$\widehat{\mu}_t^w = \gamma\widehat{c}_t + \varphi\widehat{n}_t - \widehat{w}_t$
Wage Phillips Curve	$\pi_t^w = \beta\mathbb{E}_t\pi_{t+1}^w + \kappa_w\widehat{\mu}_t^w$
Wage Inflation	$\pi_t^w = \widehat{w}_t - \widehat{w}_{t-1} + \pi_t$
Bond Euler	$\widehat{c}_t^U = s\bar{m}\mathbb{E}_t\widehat{c}_{t+1}^U + (1-s)\bar{m}\mathbb{E}_t\widehat{c}_{t+1}^H - \frac{1}{\gamma}(\widehat{i}_t - \mathbb{E}_t\pi_{t+1})$
H Budget Constraint	$\widehat{c}_t^H = \widehat{w}_t + \widehat{n}_t + \widehat{t}_t^H$
H Transfer	$\widehat{t}_t^H = \frac{\mu^D}{\lambda}D_t$
Profits	$\widehat{d}_t = \widehat{y}_t - (\widehat{w}_t + \widehat{n}_t)$
Labor Demand	$\widehat{w}_t = \widehat{mc}_t + \widehat{y}_t - \widehat{n}_t$
Phillips Curve	$\pi_t = \beta\mathbb{E}_t\pi_{t+1} + \kappa_\pi\widehat{mc}_t$
Production	$\widehat{y}_t = \widehat{n}_t$
Consumption	$\widehat{c}_t = \lambda\widehat{c}_t^H + (1-\lambda)\widehat{c}_t^U$
Resource Constraint	$\widehat{y}_t = \widehat{c}_t$
Taylor Rule	$\widehat{i}_t = \rho_i\widehat{i}_{t-1} + (1-\rho_i)\phi\pi_t + \varepsilon_t^{MP}$

Table 5: Sticky Wage Model Calibration

Parameter	γ	κ_π	λ	s	φ	μ^D	κ_w	β	ρ_i	ϕ
Value	1	0.01	0.37	$0.8^{1/4}$	1	0	0.075	0.99	0.89	1.5

the role of hand-to-mouth households in the response to monetary policy shocks and thus, allows the model to generate the pronounced hump-shaped responses. Setting $\lambda = 0.33$ still produces hump-shaped responses but those are somewhat less pronounced.