Three Essays in Macroeconomics

Inauguraldissertation zur Erlangung des akademischen Grades
eines Doktors der Wirtschaftswissenschaften
der Universität Mannheim

vorgelegt von

Oliver Pfäuti

im Frühjahrs-/Sommersemester 2023

Abteilungssprecher:	Prof. Klaus Adam, Ph.D.
Referent:	Prof. Klaus Adam, Ph.D
Koreferent:	Prof. Dr. Matthias Meier
Vorsitzender der Disputation:	Prof. Dr. Antonio Ciccone
Tag der Disputation:	24.05.2023

Acknowledgements

First and foremost, I want to thank my advisor, Klaus Adam. He taught me more about economics and what it means to do research than anybody else. Without him, I could not have written this dissertation the way I did and definitely would not have been able to enjoy the process as much as I did. Thank you, Klaus, for supporting me at all stages of my Ph.D. I also want to thank Zhen Huo who welcomed me at Yale with open arms, and made my time there so enjoyable and productive. It was in Zhen's class where I started working on my job market paper and it was his support that made me to continue pursuing the project. Further, I am very grateful to Eduardo Dávila who showed me the beauty of clean anaylitcal results and who taught me a lot about what it means to be an academic. Both, Eduardo and Zhen, kept advising and supporting me even after I returned back to Mannheim, for which I am immensely thankful.

I also want to thank Matthias Meier and Antonio Ciccone for serving on my dissertation committee. Matthias showed me that empirical macroeconomic research can be as much fun as macro theory (almost, at least) and Antonio was a great (and tough) sparring partner the very first time I presented my own research in a seminar. I still think back to that first presentation when preparing presentations now and make sure to avoid the mistakes I made back then.

I want to thank my coauthors, Timo Reinelt and Fabian Seyrich, for showing me how much more fun it is to work together. Timo always offered me his advice on all things somewhat related to economics, research and the life as a PhD student and I learned a lot from him. Fabian taught me a lot about HANK models and always supported me during times I doubted myself. When people say that you spend more time with your coauthors than your partner, this is definitely true with Fabian. Thank you for being such an amazing friend!

I benefitted a lot from the continuous feedback I received from the macro group in Mannheim, especially from Miren Azkarate-Askasua, Antoine Camous, Husnu Dalgic, Andreas Gulyas, Krzysztof Pytka, Michèle Tertilt, and Minchul Yum. I would also like to thank the faculty members I met at Yale who helped shaping me as an economist. I am particularly grateful to Marios Angeletos, Nick Barberis, Michael Peters, and Alp Simsek.

I am very grateful for the administrative support from Golareh Khalilpour, Ulrich Kehl, Marion Lehnert, Caroline Mohr, Susanne Putz, and Sylvia Rosenkranz. I gratefully acknowledge financial support from Stiftung Geld & Währung, the Swiss National Science Foundation, the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation), and CRC TR 224 (Project C02).

While working a lot during these five years, I also had a lot of fun and this is thanks to the people I met here, who lived through these years with me from the start, and who became my "Mannheim family": Giovanni Ballarin, Suzanne Bellue, Jacopo Gambato, Tommaso Gasparini, Lukas Mahler, Laura Montenbruck, Andrés Plúas Lopez, Jasmina Simon and Jonathan Stäbler. I also want to thank my friend and room-mate in the U.S., Carl-Christian Groh, who is probably the person who saw me present (or practicing to present) more than anybody else. Thanks also to Max Jager who was always there when I needed him. Thanks to my office buddies, Ulrich Roschitsch and Hannes Twieling, as well as fellow Ph.D. students who accompanied me on my journey: Alexander Matusche, Lukas Hack, Francesco Beraldi, Antzelos Kyriazis, and Chenping Yang.

None of this would have been possible without the unconditional support from my family. My parents taught me that you receive more when you give, that joy doubles when you share it, and that shared pain is half the pain. I cannot describe in words how grateful I am to both of you. I also want to thank my brother Yannik, to whom I always looked up to but never felt looked down upon. Lastly, I want to thank Julia. When I say that I would not be here without your love and support, I mean it. Without you, I probably would not have even gone to University ten years ago. But I definitely would not be the person I am today without you. You supported me when I decided to move to Germany to pursue my Ph.D., you supported me to go to the U.S. during a global pandemic and a travel ban in place, that made it impossible to know when we will see each other again, and throughout you motivated me to pursue my goals. Thank you for your unconditional love, your generosity and support. I cannot wait for the next adventure that lies ahead of us!

Contents

List of Figures

List of Tables

Introduction

This dissertation contains three chapters on macroeconomics and monetary economics with a particular focus on how subjective beliefs and behavioral biases matter for the conduct of monetary policy.

In the first chapter, which is joint work with Fabian Seyrich, we develop a new framework for business-cycle and policy analysis. The model features a New Keynesian core and we allow for household heterogeneity and incomplete markets as well as bounded rationality in the form of cognitive discounting. Under cognitive discounting, households' expectations underreact to macroeconomic news, consistent with what we find in the data. The model accounts for recent empirical findings on the transmission mechanisms of monetary policy. In particular, monetary policy affects household consumption to a large extent through indirect effects which tends to amplify the effects of conventional monetary policy on consumption as the incomes of households that exhibit higher marginal propensities to consume are more exposed to aggregate income fluctuations induced by monetary policy. Announcements of future monetary policy changes, in contrast, have relatively weak effects on current economic activity, and the economy remains stable at the effective lower bound. In contrast to existing models, we account for these facts simultaneously without having to rely on a specific monetary or fiscal policy. When abstracting from either household heterogeneity or bounded rationality the model fails to do so. Accounting for these facts simultaneously has important implications for monetary and fiscal policy. We uncover a novel amplification of negative productivity shocks through household heterogeneity, cognitive discounting and the interaction of the two, such that inflation increases substantially more than in existing models. Thus, to stabilize inflation after such a shock, the monetary authority needs to respond more aggressively which fuels inequality and leads to a large increase in government debt.

In the second chapter, which is joint work with Klaus Adam and Timo Reinelt, we document several deviations of U.S. households' housing price expectations from rational expectations. In particular, expectations are updated on average too sluggishly; following housing price changes, expectations initially underreact but subsequently overreact; and households are overly optimistic about capital gains when the price-to-rent ratio is high. We show that weak forms of capital gain extrapolation allow to simultaneously replicate the behavior of housing prices and the deviations from rational expectations as an equilibrium outcome. We embed capital gain extrapolation into a sticky price model featuring a lower-bound constraint on nominal interest rates and show that subjective beliefs about housing prices lead to a misallocation of resources. If households become optimistic about future housing prices they invest an inefficiently high amount into the housing sector rather than into non-durable consumption. This misallocation affects the natural rate of interest and we show that lower levels of natural rates increase the volatility of housing prices and thereby the volatility of the natural rate of interest. This channel exacerbates the relevance of the lower bound constraint and causes the optimal inflation target to increase strongly as the natural rate falls.

In the third chapter, I propose a new approach to quantify *attention to inflation* in the data. The approach is based on a model of optimal attention choice subject to information acquisition costs. The result is a law of motion for inflation expectations in which attention governs how strongly agents update their expectations following an inflation surprise. Using micro survey data of professional forecasters and consumers in the U.S., I then use this approach to estimate attention to inflation in the data and show that attention was very low after the 1980s. In the 1970s and 1980s, on the other hand, attention to inflation was substantially higher. Consistent with the underlying model, times of higher inflation volatility and persistence are characterized by higher attention to inflation. This decline in attention has important implications for optimal monetary policy. On the one hand, lower attention helps to stabilize inflation expectations and actual inflation. On the other hand, lower attention renders managing inflation expectations more difficult which can lead to prolonged periods of a binding lower bound and low inflation due to slowly-adjusting inflation expectations. To mitigate this loss of control, the optimal inflation target substantially increases as attention falls. Allowing for permanently higher inflation rates, however, is costly in terms of welfare. Thus, welfare decreases as attention declines whereas it would increase when abstracting from the lower bound constraint.

Chapter 1

A Behavioral Heterogeneous Agent New Keynesian Model

Joint with Fabian Seyrich

1.1 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.

[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 (2022) 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. (2012), D'Acunto et al. (2022), and Roth et al. (2023) 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.

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 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. (2022) or Roth et al. (2023).[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

[2]We show how to microfound cognitive discounting as a noisy-signal extraction problem of otherwise rational agents. Angeletos and Lian (2023) 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 (2015a), Coibion et al. (2022a) or Angeletos et al. (2021). Born et al. (2022) and Kučinskas and Peters (2022) show that even when agents overreact to micro news, they underreact to macro news, which is what we focus on in this paper.

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 as contemporaneous interest-rate changes (Del Negro et al. (2012), 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 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 (2022)). 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 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 households.

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.

[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.

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. (2021) 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. (2023) 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

[4]Bilbiie (2021) proposes an additional resolution: a pure risk channel which can, in theory, break the comovement 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 1.4.2 where we discuss heterogeneous degrees of cognitive discounting) and show that our results remain robust in that scenario.

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 1.2 and our analytical results in Section 1.3. In Section 1.4, 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 1.5. We discuss three extensions of the tractable model in Section 1.6 and Section 1.7 concludes.

1.2 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)).

1.2.1 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

$$\mathcal{U}(C_t^i, N_t^i) \equiv \begin{cases} \frac{(C_t^i)^{1-\gamma}}{1-\gamma} - \frac{(N_t^i)^{1+\varphi}}{1+\varphi}, & \text{if } \gamma \neq 1, \\ \log\left(C_t^i\right) - \frac{(N_t^i)^{1+\varphi}}{1+\varphi}, & \text{if } \gamma = 1, \end{cases} \tag{1.1}$$

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

[6] Wiederholt (2015), Angeletos and Lian (2018), Andrade et al. (2006.12019), 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.

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

$$C_t^U + B_{t+1}^U = W_t N_t^U + \frac{1-\mu^D}{1-\lambda} D_t + s\frac{1+i_{t-1}}{1+\pi_t} B_t^U, \tag{1.2}$$

and the hand-to-mouth households

$$C_t^H = W_t N_t^H + \frac{\mu^D}{\lambda} D_t + (1-s)\frac{1+i_{t-1}}{1+\pi_t}\frac{1-\lambda}{\lambda} B_t^U, \tag{1.3}$$

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

$$\left(C_t^U\right)^{-\gamma} \geq \beta\mathbb{E}_t^{BR}\left[R_t\left(s\left(C_{t+1}^U\right)^{-\gamma} + (1-s)\left(C_{t+1}^H\right)^{-\gamma}\right)\right], \tag{1.4}$$

[7] We show in Appendix A.1.9 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.

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

$$\mathbb{E}_t^{BR}\left[X_{t+1}\right] = \mathbb{E}_t^{BR}\left[X_t^d + \tilde{X}_{t+1}\right] \equiv X_t^d + \bar{m}\mathbb{E}_t\left[\tilde{X}_{t+1}\right], \tag{1.5}$$

where $\mathbb{E}_t\left[\cdot\right]$ 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 1.6.3. 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.1.9, note, that the exact microfoundation or underlying behavioral friction is not crucial for the rest of our analysis. Angeletos and Lian (2023) show how other forms of bounded rationality or lack of common knowledge can lead to observationally-equivalent expectations.

Log-linearizing equation (1.5) around the steady state yields

$$\mathbb{E}_t^{BR}\left[\widehat{x}_{t+1}\right] = (1-\bar{m})\widehat{x}_t^d + \bar{m}\mathbb{E}_t\left[\widehat{x}_{t+1}\right] \tag{1.6}$$

[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 (2022) 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.1.8 derives our results following the approach in Gabaix (2020). The results remain exactly the same.

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. (2022) or Roth et al. (2023). In Appendix A.2, 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 (2015a) 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 A.4 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 1.4.

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

$$\widehat{i}_t = \phi \pi_t + \epsilon_t^{MP}, \tag{1.7}$$

with ϵ_t^{MP} being a monetary policy shock (Appendix A.1 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.

1.2.2 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

$$\widehat{c}_t^H = \chi \widehat{y}_t, \tag{1.8}$$

with

$$\chi \equiv 1 + \varphi\left(1 - \frac{\mu^D}{\lambda}\right) \tag{1.9}$$

measuring the cyclicality of the H household's consumption (see appendix A.1.1). Patterson (2022) 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) and Hintermaier and Koeniger (2019) that poorer households' income is on average more exposed to monetary policy shocks, 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 (1.8) with the goods market clearing condition yields

$$\widehat{c}_t^U = \frac{1 - \lambda\chi}{1 - \lambda}\widehat{y}_t, \tag{1.10}$$

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

$$\widehat{c}_t^U - \widehat{c}_t^H = \frac{1-\chi}{1-\lambda}\widehat{y}_t. \tag{1.11}$$

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 (2023) 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:

$$\widehat{c}_t^U = s\mathbb{E}_t^{BR}\left[\widehat{c}_{t+1}^U\right] + (1-s)\mathbb{E}_t^{BR}\left[\widehat{c}_{t+1}^H\right] - \frac{1}{\gamma}\left(\widehat{i}_t - \mathbb{E}_t^{BR}\pi_{t+1}\right). \tag{1.12}$$

For the case without type-switching, i.e., for $s = 1$, equation (1.12) 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 (1.12). We show in Appendix A.4 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:

$$\pi_t = \kappa\widehat{y}_t, \tag{1.13}$$

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 A.4 we show that a forward-looking Phillips Curve (rational or behavioral) does not qualitatively affect our results.

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 particu-

[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 1.4 and 1.5, households' consumption levels will differ in the stationary equilibrium.

lar, 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 1.4 and show that our results presented in the following do not depend on these assumptions.

1.3 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.

1.3.1 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 (1.13), a rule for monetary policy (equation (1.7)), which together with the aggregate IS equation determines aggregate demand. To obtain the aggregate IS equation, we combine the hand-to-mouth households' consumption (1.8) with the consumption of unconstrained households (1.10) and their consumption Euler equation (1.12) (see appendix A.1 for all the derivations).

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

$$\widehat{y}_t = \psi_f \mathbb{E}_t \widehat{y}_{t+1} - \psi_c \frac{1}{\gamma} \left(\widehat{i}_t - \mathbb{E}_t \pi_{t+1} \right), \tag{1.14}$$

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 (1.14), 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 (1.14) 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 (2022). 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. (2022). The cognitive discounting parameter, $\bar{m}$ is set to 0.85, as explained in Section 1.2. Details on the calibration and a discussion of the robustness of our findings for different calibrations are presented in Appendix A.2. Note, that even when we vary certain parameters, we keep $\lambda < \chi^{-1}$.

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 (1.14) with the static Phillips Curve (1.13) and the Taylor rule (1.7) so that we can represent the model in a single first-order difference equation:

$$\widehat{y}_t = \frac{\psi_f + \psi_c \frac{\kappa}{\gamma}}{1 + \psi_c \phi \frac{\kappa}{\gamma}} \mathbb{E}_t \widehat{y}_{t+1} - \frac{\psi_c \frac{1}{\gamma}}{1 + \psi_c \phi \frac{\kappa}{\gamma}} \varepsilon_t^{MP}. \tag{1.15}$$

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. (2012), Roth et al. (2023)).

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*

$$\psi_c > 1 \Leftrightarrow \chi > 1, \tag{1.16}$$

and the forward guidance puzzle is ruled out if

$$\psi_f + \frac{\kappa}{\gamma}\psi_c < 1. \tag{1.17}$$

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 1.2.2, 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*

$$\widehat{c}_t = [1 - \beta(1 - \lambda\chi)]\,\widehat{y}_t - \frac{(1-\lambda)\beta}{\gamma}\widehat{r}_t + \beta\bar{m}\delta(1 - \lambda\chi)\mathbb{E}_t\widehat{c}_{t+1}. \tag{1.18}$$

[11]Detailed analyses of the forward guidance puzzle in RANK are provided by McKay et al. (2016) and Del Negro et al. (2012).

[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 (1.15) is smaller than 1. Given that we consider $\phi = 0$, this boils down to the condition stated in Proposition 2.

What determines whether condition (1.17) 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 1.2).[13] Figure A.1 in Appendix A.2 shows that the solution of the forward guidance in our model is very robust with respect to changes in the heterogeneity parameters.

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 A.3. 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.

[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 A.4.3.

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 principle is given by $\phi > 1$, where ϕ is the inflation-response coefficient in the Taylor rule (1.7). 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:*

$$\phi > \phi^* = 1 + \frac{\bar{m}\delta - 1}{\frac{\kappa}{\gamma}\frac{1-\lambda}{1-\lambda\chi}}. \tag{1.19}$$

We obtain Proposition 3 directly from the difference equation (1.15). 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.1.4 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

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

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 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

$$\widehat{y}_t = -\frac{1}{\gamma} \psi_c \underbrace{\left(\widehat{i}_{ELB} - \widetilde{r}^n \right)}_{>0} \sum_{j=0}^{k} \left(\psi_f + \frac{\kappa}{\gamma} \psi_c \right)^j , \tag{1.20}$$

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 A.3 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 A.4, 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. (2022)). 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.

[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 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.

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 policy mainly works through direct intertemporal substitution channels and these models 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 1.3.2. 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.

1.4 The Full Behavioral HANK Model

In this section, we move from the limited-heterogeneity setup described in Section 1.2 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 (1.1)).

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

$$\begin{aligned} C_{i,t}^{-\gamma} &\geq \beta R_t \mathbb{E}_t^{BR}\left[C_{i,t+1}^{-\gamma}\right] \\ &= \beta R_t \mathbb{E}_t^{BR}\left[\bar{C}_{i,t+1}^{-\gamma} + \left(C_{i,t+1}^{-\gamma} - \bar{C}_{i,t+1}^{-\gamma}\right)\right] \\ &= \beta R_t \mathbb{E}_t\left[\bar{C}_{i,t+1}^{-\gamma} + \bar{m}\left(C_{i,t+1}^{-\gamma} - \bar{C}_{i,t+1}^{-\gamma}\right)\right], \end{aligned} \tag{1.21}$$

where the rational expectations operator $\mathbb{E}_t[\cdot]$ denotes the expectations that a fully rational household would have in the behavioral economy. Equation (1.21) 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 (2022)). 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 A.5.

[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.

1.4.1 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 1.1: Monetary Policy and Forward Guidance

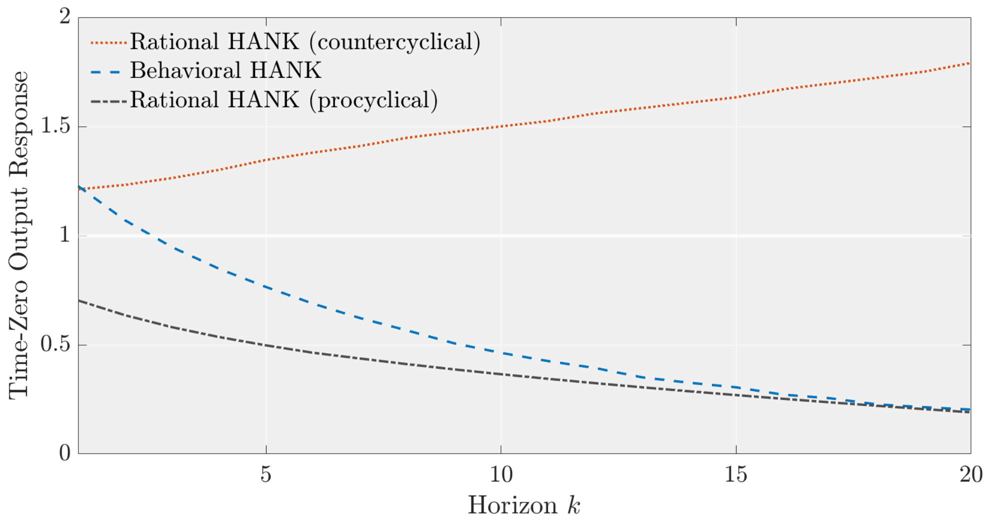

Note: This figure shows the response of total output in period 0 to anticipated one-time monetary policy shocks occurring at different horizons k, relative to the response in the representative agent model under rational expectations (normalized to 1). The blue-dashed line shows the results for the behavioral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

Figure 1.1 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 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 (2022), 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 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 1.5.

Sticky wages. So far, we have assumed that prices are sticky and wages are fully flexible. We show in Appendix A.5.1 that our results remain robust if instead we allow prices to be fully flexible but wages to be sticky.

Stronger amplification. In our baseline calibration, contemporaneous monetary policy shocks are 20% more effective compared to the representative agent model. In Appendix A.5.2, we show that cognitive discounting still resolves the forward guidance puzzle even when doubling the initial amplification to 40%. While the forward guidance puzzle gets substantially exacerbated in the rational model, the behavioral HANK model still resolves the puzzle and thus, our results remain robust.

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

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

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 A.5 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 A.5 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.

1.4.2 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 A.2.1, 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 A.5.

Figure 1.2 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

Figure 1.2: Heterogeneous $\bar{m}$ and Monetary Policy

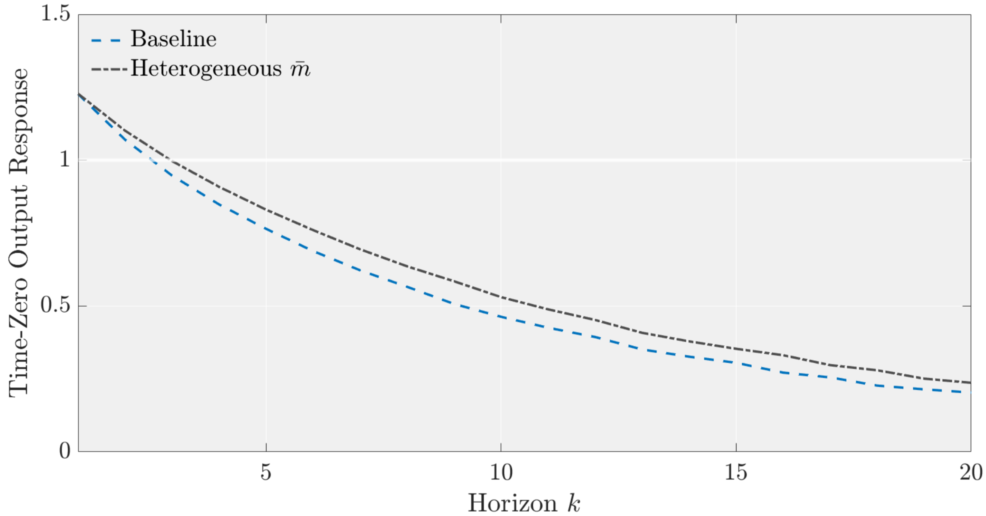

Note: This figure shows the response of total output in period 0 to anticipated i.i.d. monetary policy shocks occurring at different horizons k for the baseline calibration with $\bar{m} = 0.85$ for all households (blue-dashed line) and for the model in which households differ in their levels of cognitive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

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 A.2.1, 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.

1.5 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 1.3 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 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 1.3 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 1.3 and 1.4 after policy shocks are muted here.

Figure 1.3: Inflationary supply shock: strict inflation-targeting

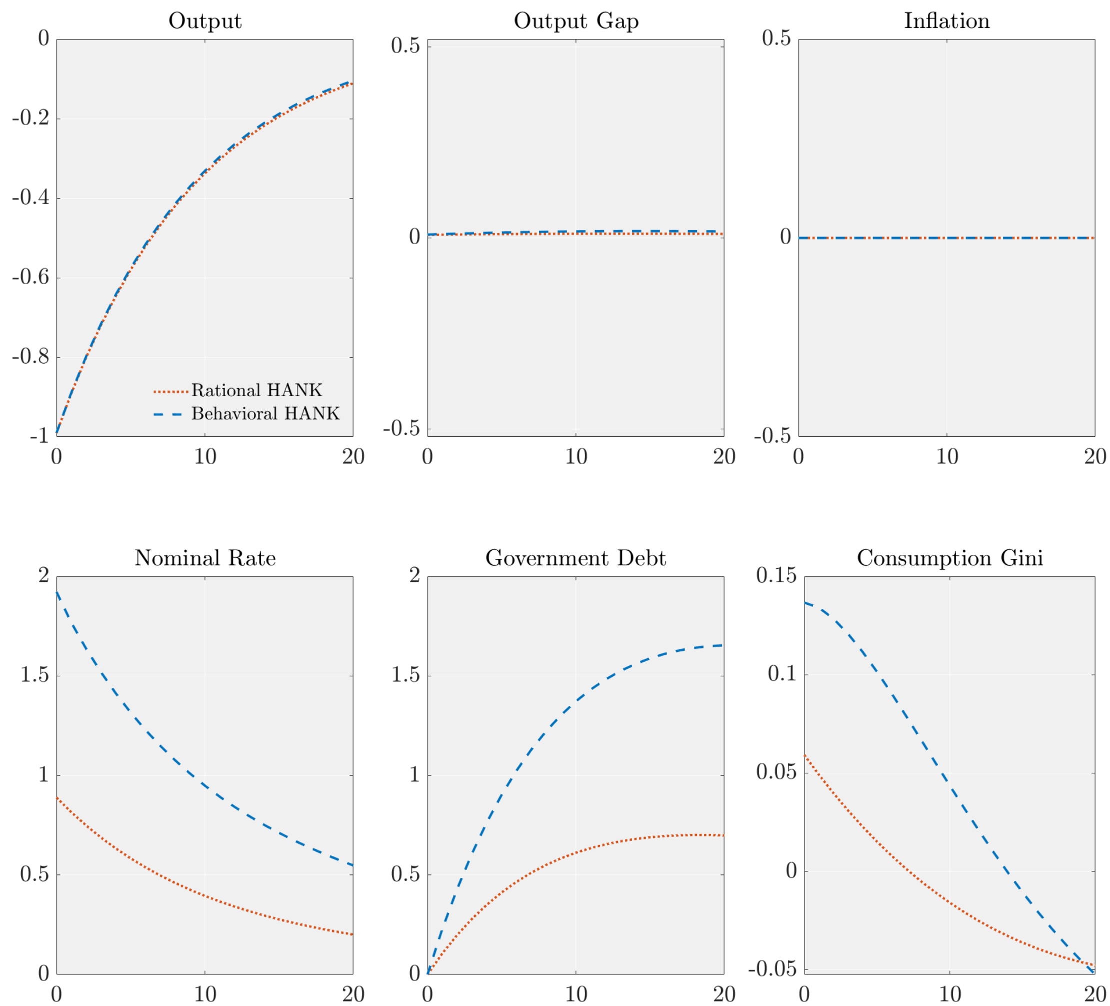

Note: This figure shows the impulse responses after a TFP shock that decreases potential output by 1% in the inflation-stabilizing monetary policy regime. Output and the output gap are shown as percentage deviations from steady state output, the nominal interest rate and inflation as annualized percentage 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 Appendix A.6.6, 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.

How do these results change when monetary policy follows a standard Taylor rule? Figure 1.4 shows the case when monetary policy follows a simple Taylor rule (1.7) 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

Figure 1.4: Inflationary supply shock: Taylor rule

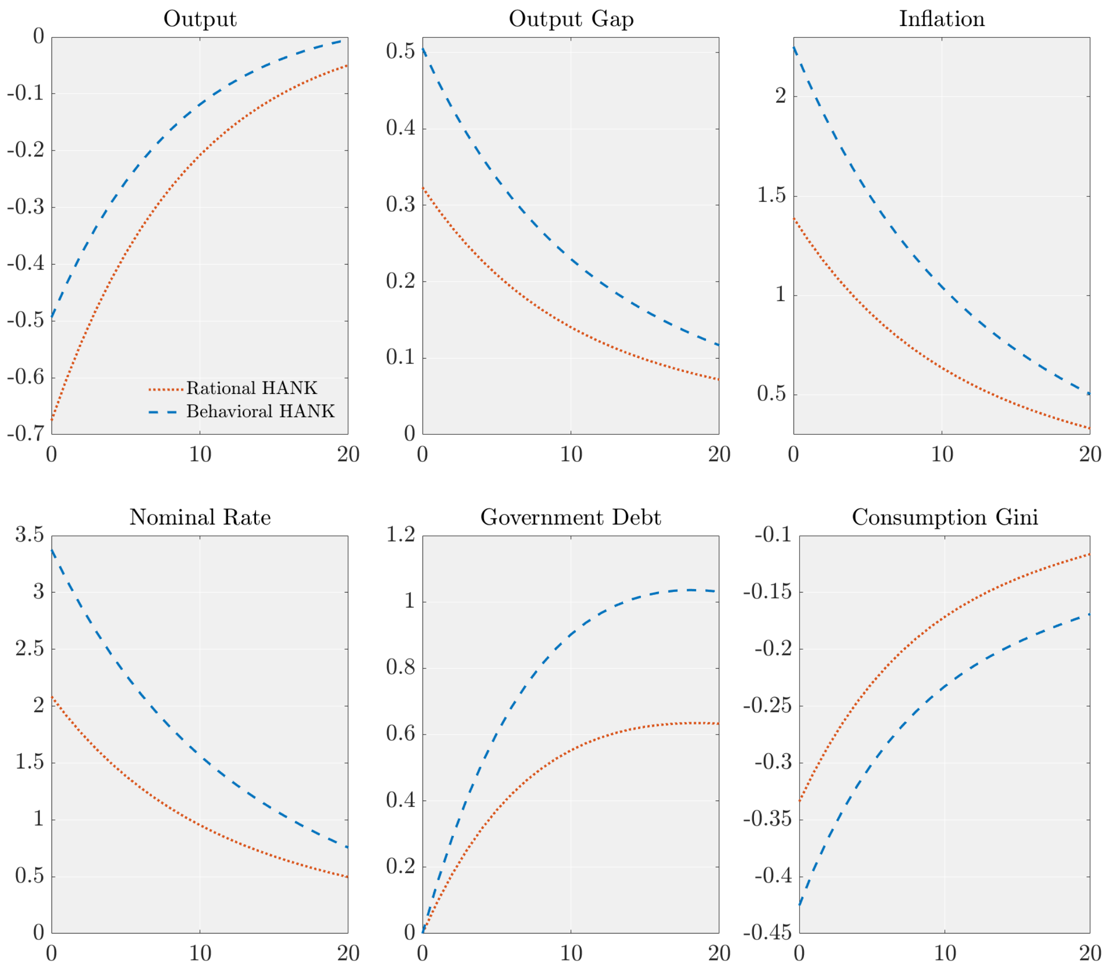

Note: This figure shows the impulse responses after a productivity shock for the case that monetary policy follows a Taylor rule. Output and the output gap are shown as percentage deviations from steady state output, the nominal interest rate and inflation as annualized percentage points and the government 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.

in Sections 1.3 and 1.4 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 would amplify and cognitive discounting dampen the effect. For example, a monetary shock of 1 percentage point

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 1.4). 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 A.6.3). 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 1.5 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 A.11 in Appendix A.6.1, the inflation increase that is due to the interaction becomes even larger than the one due

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.

to household heterogeneity under this alternative calibration.

Figure 1.5: Decomposition of the Additional Inflation Increase

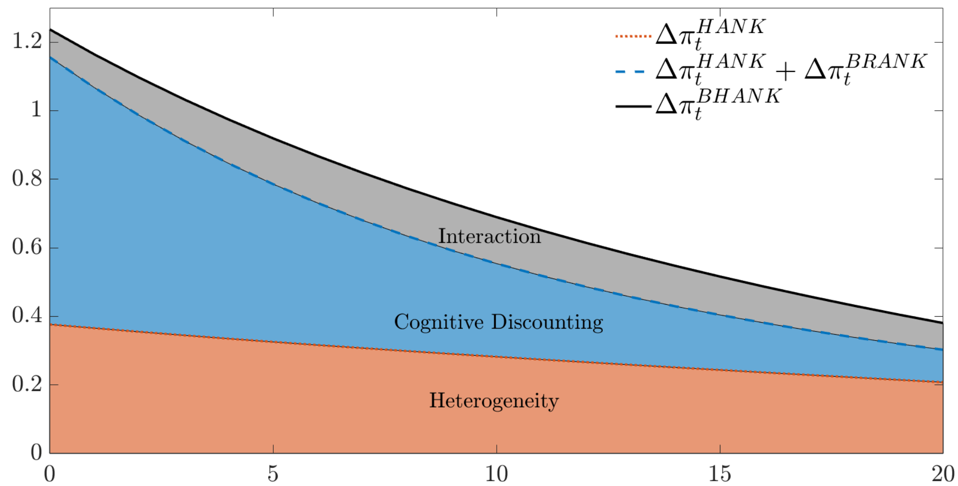

Note: This figure shows the additional inflation increase in the rational HANK model compared to the rational RANK model (orange-dotted line), the sum of the additional increase in the rational HANK and the behavioral RANK compared to the rational RANK model (blue-dashed line) and the additional increase 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 1.4, 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 A.13 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 (2015a), Coibion et al. (2022a), D'Acunto et al. (2022) or Roth et al. (2023)). 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 A.6.5 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 A.6.5 shows these results graphically.

Heterogeneous $\bar{m}$. When we incorporate heterogeneous $\bar{m}$ along the lines described in Section 1.4.2, 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 A.6.4). 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.

Behavioral firms. In Appendix A.6.7, we discuss the case in which firms cognitively discount the future in the same way as households. The increase in inflation when monetary policy follows a Taylor rule is somewhat muted whereas the increase in the output gap is amplified compared to the case in which firms are rational. The reason is that firms discount the increase in their future marginal costs and thus increase their prices not as strongly. According to the Taylor rule this then leads to a smaller increase in interest rates so that households consume more, leading to an increase in demand and thus, the output gap.

Cost-push shocks. So far, we have focused on the inflationary pressure coming from negative TFP shocks. We show in Appendix A.6.8 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.

1.6 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.

1.6.1 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 (2022)).[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 following Proposition characterizes the iMPCs in the behavioral HANK model (see Appendix A.7 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*

$$MPC_0 \equiv \frac{d\widehat{c}_0}{d\widetilde{y}_0} = 1 - \frac{1-\lambda\chi}{s\bar{m}}\mu_2^{-1}$$
$$MPC_k \equiv \frac{d\widehat{c}_k}{d\widetilde{y}_0} = \frac{1-\lambda\chi}{s\bar{m}}\mu_2^{-1}\left(\beta^{-1}-\mu_1\right)\mu_1^{k-1}, \quad \textit{for } k>0,$$

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 A.7.

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 1.3. Figure 1.6 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

[20]See, e.g., Lian (2023) or Boutros (2022) for MPC analyses in models deviating from FIRE.

Figure 1.6: Intertemporal MPCs, Bounded Rationality and Household Heterogeneity

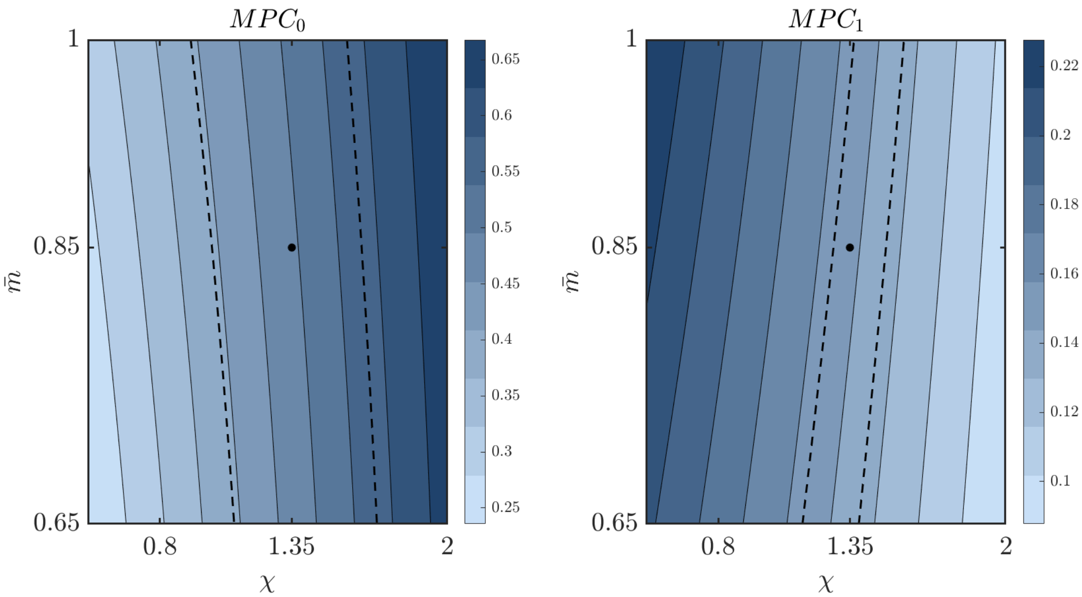

Note: This figure shows the aggregate intertemporal MPCs in our tractable model, i.e., the aggregate consumption response in year 0 (left) and year 1 (right) to a change in aggregate disposable income in year 0 for different χ (horizontal axis) and $\bar{m}$ (vertical axis). The dashed lines show the 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).

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 A.25 in Appendix A.7 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 A.26 in Appendix A.7).

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%.

1.6.2 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. (2022).

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 A.8.

Hump-shaped responses to monetary policy shocks. Figure 1.7 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.

Figure 1.7: Monetary Policy Shock with Sticky Wages

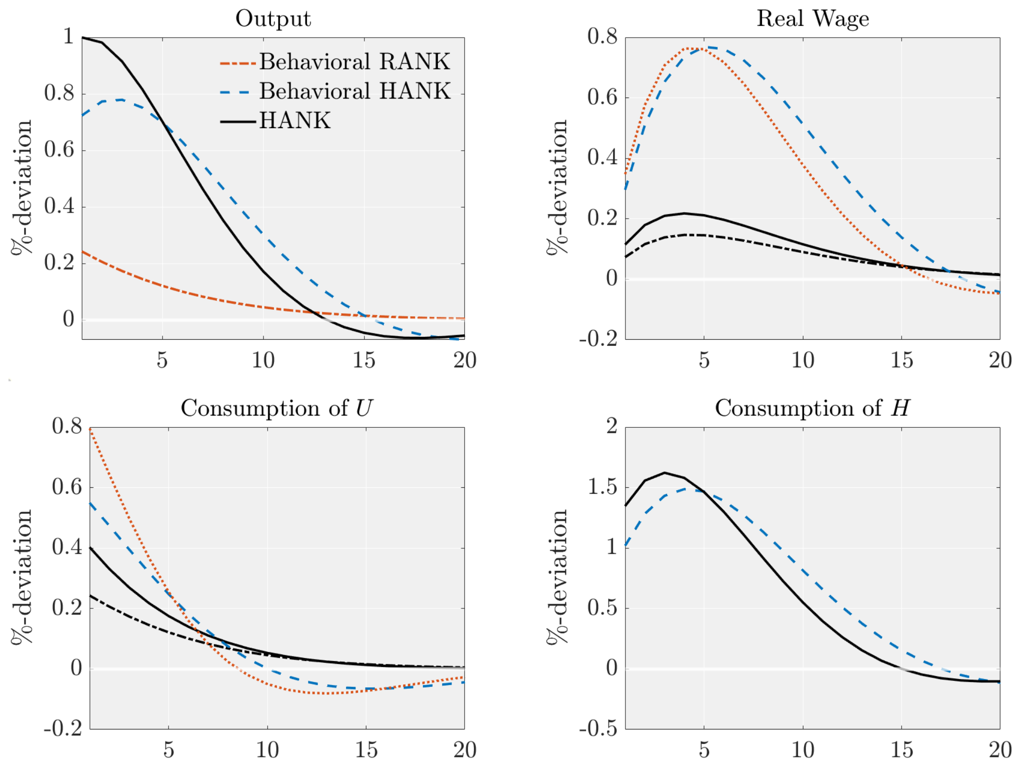

Note: This figure shows the impulse-response functions of output, real wages and consumption of the two household types to a monetary policy shock in the tractable behavioral HANK model, the rational and the behavioral RANK model and the rational HANK model with sticky wages. The shock size is normalized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

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 motive 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 1.6.1 as well as the results presented in Figure 1.7 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 1.8 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 1.7 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

Figure 1.8: Forecast Errors with Sticky Wages in Response to a Monetary Policy Shock

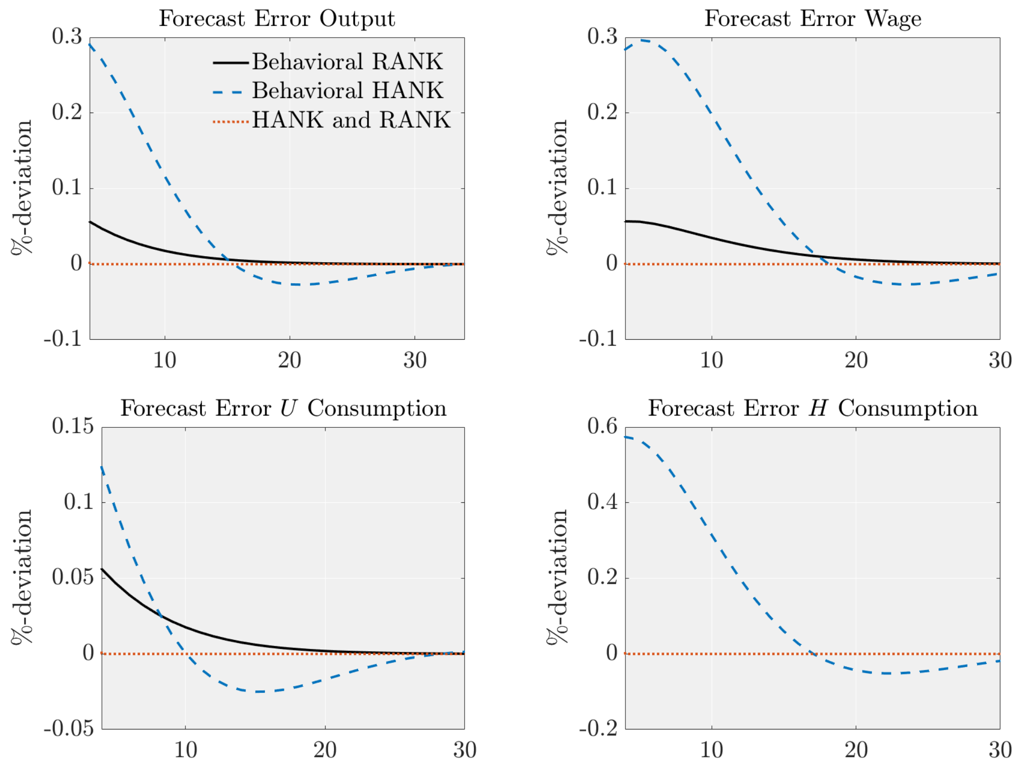

Note: This figure shows the forecast error dynamics of output, the real wage, consumption of unconstrained households and of hand-to-mouth households after an expansionary monetary policy shock in the tractable model.

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

1.6.3 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 (2023)). 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*

$$\mathbb{E}_t^{BR}\left[X_{t+1}\right] = (1-\bar{m})X_{t-1} + \bar{m}\mathbb{E}_t\left[X_{t+1}\right] \tag{1.22}$$

and the behavioral HANK IS equation is then given by

$$\widehat{y}_t = \psi_f \mathbb{E}_t \widehat{y}_{t+1} - \psi_c \frac{1}{\gamma}\left(\widehat{i}_t - \mathbb{E}_t \pi_{t+1}\right) + (1-\bar{m})\delta \widehat{y}_{t-1}. \tag{1.23}$$

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 (2023) 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.

1.7 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.

Appendix A

Appendix to Chapter 1

A.1 Model Details and Derivations

A.1.1 Derivation of χ

In Section 1.2, we stated that

$$\widehat{c}_t^H = \chi \widehat{y}_t, \tag{A.1}$$

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 (A.1) from the H-households' budget constraint, optimality conditions and market clearing.

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

$$(N_t^H)^{\varphi} = W_t (C_t^H)^{-\gamma}, \tag{A.2}$$

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

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

Since both households work for the same wage, we obtain

$$\varphi \widehat{n}_t^H + \gamma \widehat{c}_t^H = \varphi \widehat{n}_t^U + \gamma \widehat{c}_t^U \tag{A.3}$$

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 (A.3) then gives

$$\varphi\widehat{n}_t^H + \gamma\widehat{c}_t^H = (\varphi+\gamma)\widehat{y}_t. \tag{A.4}$$

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

$$C_t^H = W_t N_t^H + \frac{\mu^D}{\lambda} D_t. \tag{A.5}$$

In log-linearized terms, we get

$$\widehat{c}_t^H = \widehat{w}_t + \widehat{n}_t^H + \frac{\mu^D}{\lambda}\widehat{d}_t, \tag{A.6}$$

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

$$\widehat{c}_t^H = \left(\varphi\widehat{n}_t^H + \gamma\widehat{c}_t^H\right)\left(1 - \frac{\mu^D}{\lambda}\right) + \widehat{n}_t^H. \tag{A.7}$$

Using (A.4) to solve for $\widehat{n}_t^H$ and plugging it into (A.7) 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.1.2 Derivation of Proposition 1.

Combining equations (1.8) and (1.10) with the bounded-rationality setup in equation (1.6) for $\widehat{x}_t^d = 0$ as X_t^d is given by the steady state, we have

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

Plugging these two equations as well as equation (1.10) into the Euler equation of unconstrained households (1.12) 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.1.3 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 (1.13), 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.1.4 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 (1.14), 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

$$\phi > \phi^* = 1 + \frac{\bar{m}\delta - 1}{\frac{\kappa}{\gamma}\frac{1-\lambda}{1-\lambda\chi}}, \tag{A.8}$$

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 (A.8) 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

$$\widehat{y}_t = \frac{\psi_f + \frac{\kappa}{\gamma}\psi_c \bar{m}^r}{1 + \frac{\psi_c}{\gamma}\left(\kappa\phi_\pi + \phi_y\right)} \mathbb{E}_t \widehat{y}_{t+1}. \tag{A.9}$$

From equation (A.9), it follows that we need

$$\phi_\pi > \bar{m}^r - \phi_y + \frac{\psi_f - 1}{\psi_c \frac{\kappa}{\gamma}} = \bar{m}^r - \phi_y + \frac{\bar{m}\delta - 1}{\frac{1-\lambda}{1-\chi\lambda}\frac{\kappa}{\gamma}} \tag{A.10}$$

for the model to feature a determinate, locally unique equilibrium. Condition (A.10) 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.1.5 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:

$$\widehat{c}_t^H = \widehat{w}_t + \widehat{n}_t^H + \frac{\mu^D}{\lambda}\widehat{d}_t - g_t, \tag{A.11}$$

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

$$\widehat{y}_t = \lambda \widehat{c}_t^H + (1-\lambda)\widehat{c}_t^U + g_t. \tag{A.12}$$

Plugging this and the labor market clearing condition into (A.3), yields:

$$\varphi \widehat{n}_t^H + \gamma \widehat{c}_t^H = (\varphi + \gamma)\widehat{y}_t - \gamma g_t. \tag{A.13}$$

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

$$\widehat{c}_t^H = \left(\varphi \widehat{n}_t^H + \gamma \widehat{c}_t^H\right)\left(1 - \frac{\mu^D}{\lambda}\right) + \widehat{n}_t^H - g_t. \tag{A.14}$$

and using (A.13) yields:

$$\widehat{c}_t^H = \left(\varphi \widehat{n}_t^H + \gamma \widehat{c}_t^H\right)\left(1 - \frac{\mu^D}{\lambda}\right) + \widehat{n}_t^H - g_t. \tag{A.15}$$

Finally, consumption of H is given by:

$$\widehat{c}_t^H = \chi \widehat{y}_t - \left[\frac{\chi - 1}{1 + \frac{\varphi}{\gamma}} + 1\right] g_t \tag{A.16}$$

which is

$$\widehat{c}_t^H = \chi\left(\widehat{y}_t - g_t\right) + \left[\frac{\chi - 1}{1 + \frac{\gamma}{\varphi}}\right] g_t. \tag{A.17}$$

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

$$\widehat{c}_t^U = \frac{1 - \lambda\chi}{1 - \lambda}\left(\widehat{y}_t - g_t\right) - \frac{\lambda}{1-\lambda}\frac{\chi - 1}{1 + \frac{\gamma}{\varphi}} g_t. \tag{A.18}$$

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 $\widehat{c}_t = (\widehat{y}_t - g_t)$.

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

which can be re-written as (using similar derivations as in Appendix A.1.2)

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

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.1.6 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

$$\frac{1}{\gamma}\mathbb{E}_t^{BR}\widehat{q}_{t,t+i}^U + \bar{m}^i\mathbb{E}_t\widehat{c}_{t+i}^U = \widehat{c}_t^U + (1-s)\mathbb{E}_t\sum_{k=1}^{i}\bar{m}^k\left(\widehat{c}_{t+k}^U - \widehat{c}_{t+k}^H\right). \tag{A.19}$$

The (linearized) budget constraint can be written as

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

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 (A.19) when summing them up. The left-hand side of the budget constraint can thus be written as

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

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

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

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

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

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

$$\widehat{c}_t = \left[1-\beta(1-\lambda\chi)\right]\widehat{y}_t - \frac{(1-\lambda)\beta}{\gamma}\widehat{r}_t + \beta\bar{m}\delta(1-\lambda\chi)\mathbb{E}_t\widehat{c}_{t+1}, \tag{A.20}$$

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 (A.20), 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.1.7 Derivation of Proposition 5

To prove Proposition 5, we start from the Euler equation (1.12). Plugging in for $\widehat{c}_t^U$, $\widehat{c}_{t+1}^U$ and $\widehat{c}_{t+1}^H$ from equations (1.8) and (1.10), 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.1.8 Cognitive Discounting of the State Vector

In Section 1.2, 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

$$X_{t+1} = G^X\left(X_t, \varepsilon_{t+1}\right), \tag{A.21}$$

where G^X denotes the transition function of X in equilibrium and ε are zero-mean innovations. Linearizing equation (A.21) yields

$$X_{t+1} = \Gamma X_t + \varepsilon_{t+1}, \tag{A.22}$$

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

$$X_{t+1} = \bar{m} G^X(X_t, \varepsilon_{t+1}), \tag{A.23}$$

or in linearized terms

$$X_{t+1} = \bar{m}\left(\Gamma X_t + \varepsilon_{t+1}\right). \tag{A.24}$$

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

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

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

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

$$\mathbb{E}_t^{BR} \left[\widehat{c}^U (X_{t+1}) \right] = \bar{m} \mathbb{E}_t \left[\widehat{c}^U (X_{t+1}) \right], \tag{A.25}$$

which we denote in the main text as

$$\mathbb{E}_t^{BR} \left[\widehat{c}_{t+1}^U \right] = \bar{m} \mathbb{E}_t \left[\widehat{c}_{t+1}^U \right]. \tag{A.26}$$

Now, take the linearized Euler equation (1.12) of unconstrained households:

$$\widehat{c}_t^U = s \mathbb{E}_t^{BR} \left[\widehat{c}_{t+1}^U \right] + (1-s) \mathbb{E}_t^{BR} \left[\widehat{c}_{t+1}^H \right] - \frac{1}{\gamma} \widehat{r}_t, \tag{A.27}$$

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

$$\widehat{c}^U (X_t) = s \mathbb{E}_t^{BR} \left[\widehat{c}^U (X_{t+1}) \right] + (1-s) \mathbb{E}_t^{BR} \left[\widehat{c}^H (X_{t+1}) \right] - \frac{1}{\gamma} \widehat{r}(X_t). \tag{A.28}$$

Now, applying the results above, we obtain

$$\widehat{c}^U (X_t) = s \bar{m} \mathbb{E}_t \left[\widehat{c}^U (X_{t+1}) \right] + (1-s) \bar{m} \mathbb{E}_t \left[\widehat{c}^H (X_{t+1}) \right] - \frac{1}{\gamma} \widehat{r}(X_t), \tag{A.29}$$

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.1.9 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

$$X_{t+1} = \bar{m}\left(\Gamma X_t + \varepsilon_{t+1}\right), \tag{A.30}$$

as imposed in equation (A.24). The boundedly-rational expectation of X_{t+1} is then given by

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

A.2 Calibration

Our baseline calibration is summarized in Table A.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 A.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 1.3 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 (2022).

For figure 1.6, 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 (2015a) 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

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

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.

A.2.1 Estimating $\bar{m}$ for different Household Groups

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

$$x_{t+4} - \mathbb{E}_t^{e,BR} x_{t+4} = c^e + b^{e,CG}\left(\mathbb{E}_t^{e,BR} x_{t+4} - \mathbb{E}_{t-1}^{e,BR} x_{t+4}\right) + \epsilon_t^e, \tag{A.31}$$

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

$$\bar{m}^e = \left(\frac{1}{1+b^{e,CG}}\right)^{1/4}. \tag{A.32}$$

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. (2022) to translate these categorical unemployment expectations 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. Setting $a = 0.1$, for example, barely affects our results.

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. (2022). 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

$$u_{t+4} - \mu_t^e. \tag{A.33}$$

For the case of expected unemployment changes, we replace u_{t+4} with $(u_{t+4} - u_t)$ in equation (A.33).

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

$$u_{t+4} - \mu_t^e = c^e + b^{e,CG} \left(\mu_t^e - \mu_{t-1}^e \right) + \epsilon_t^e, \tag{A.34}$$

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 A.2: Regression Results of Equation (A.31)

	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 (A.31) for different income groups. The first three columns show the results when the right-hand side in equation (A.31) 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 A.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 (A.32) 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.[1]

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 differences 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.

A.2.2 Robustness of Forward Guidance Puzzle Solution

In our baseline calibration, we set χ such that we obtain an amplification of conventional monetary policy shocks of 20% compared to the case in which all households are equally exposed to monetary policy (i.e., $\psi_c = 1.2$) for a given share of hand-to-mouth, λ. In particular, we set $\lambda = 0.33$. This results in $\chi = 1.35$. But directly calibrating χ is difficult, as it is hard to measure in the data. Instead, we indirectly calibrate it in order to obtain some desired ψ_c. However, ψ_c depends on χ and λ, and both matter for whether the model resolves the forward guidance puzzle or not.

We therefore show in Figure A.1 for different ψ_c (on the horizontal axis) the highest $\bar{m}$ (on the vertical axis) that still resolves the forward-guidance puzzle. The blue-dashed line shows this for $\lambda = 0.33$ and the orange-dotted line for $\lambda = 0.147$ which is the share of borrowing-constrained households in Farhi and Werning (2019).

We see that a $\bar{m}$ of 0.85 (as indicated by the black-solid line) rules out the forward-guidance puzzle in almost all cases. Only at the relatively low λ of 0.147 and a high $\psi_c > 1.48$, we would require a $\bar{m}$ of about 0.84 instead of 0.85 to rule out the forward-guidance puzzle. Given that the empirical estimates point towards values of $\bar{m} \in [0.6, 0.85]$, we conclude the resolution of the forward guidance puzzle in the behavioral HANK model with countercyclical income risk is quite robust.

[1] 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.

Figure A.1: Robustness of FG Puzzle Solution

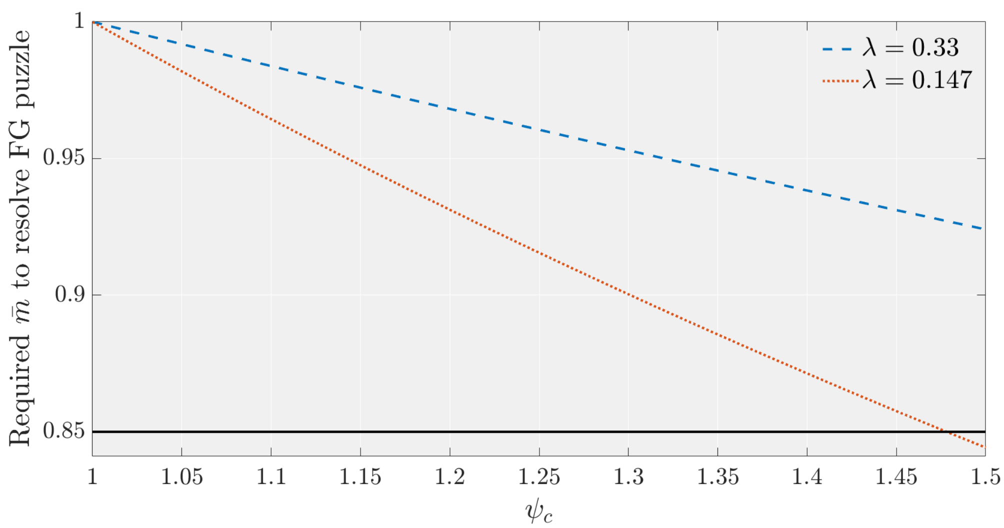

Note: This figures show for different ψ_c (horizontal axis) the required $\bar{m}$ to resolve the forward-guidance puzzle on the vertical axis. The blue-dashed line shows this for our benchmark calibration of $\lambda = 0.33$ and the orange-dotted line for $\lambda = 0.147$.

A.3 Figures to Section 1.3

A.3.1 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 A.2. 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.[2]

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 (2022), Auclert (2019)).

Figure A.2: Resolving the Catch-22

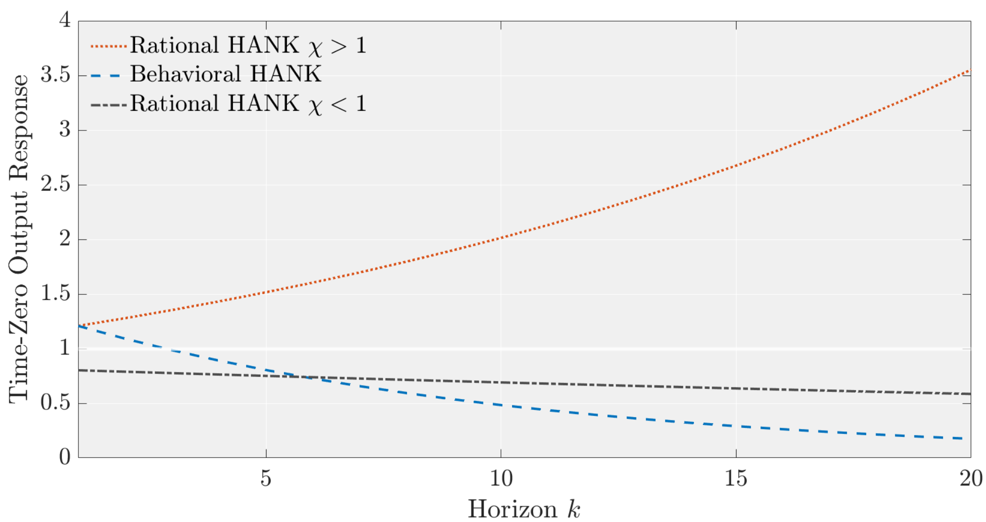

Note: This figure shows the response of total output in period 0 to anticipated i.i.d. monetary policy shocks occurring at different horizons k (horizontal axis), relative to the initial response in the RANK model under rational expectations (equal to 1).

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

[2] 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 A.3: The Effective Lower Bound Problem

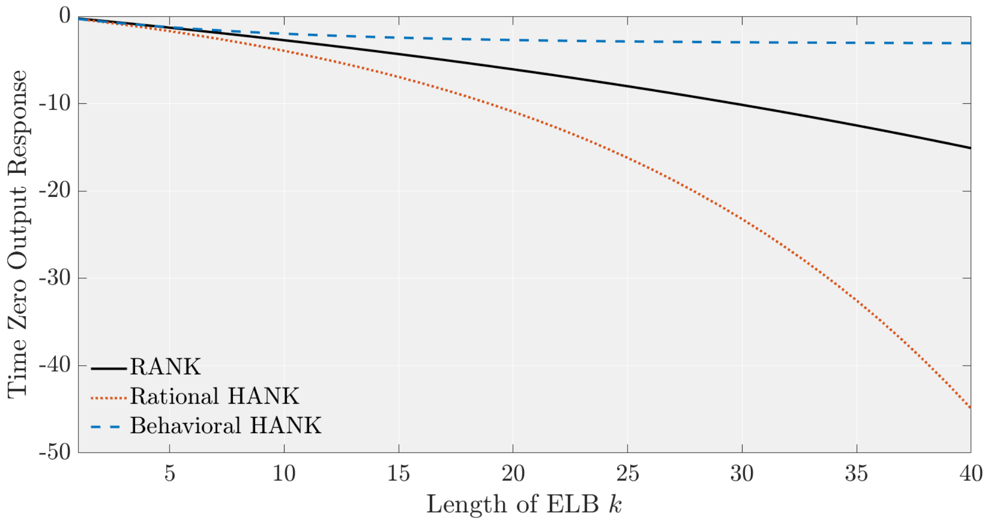

Note: This figure shows the contemporaneous output response for different lengths of a binding ELB k (horizontal axis) and compares the responses across different models.

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

A.3.2 Stability at the Effective Lower Bound

We illustrate the stability of the behavioral HANK model at the lower bound graphically in Figure A.3. 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 $\bar{m} = 0.85$.[3] Figure A.3 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.

[3]Note, that the size of the gap, $\left(\widehat{i}_{ELB} - \widetilde{r}^n\right) > 0$, does not matter qualitatively for the results.

A.4 Further Extensions

A.4.1 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. (2022) 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.1.5). 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)).[4]

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.

[4]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.

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.

Figure A.4 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. (2022) who estimate the non-durable consumption response to lie between 0.2 and 0.29.

Figure A.4: Consumption Response to Government Spending

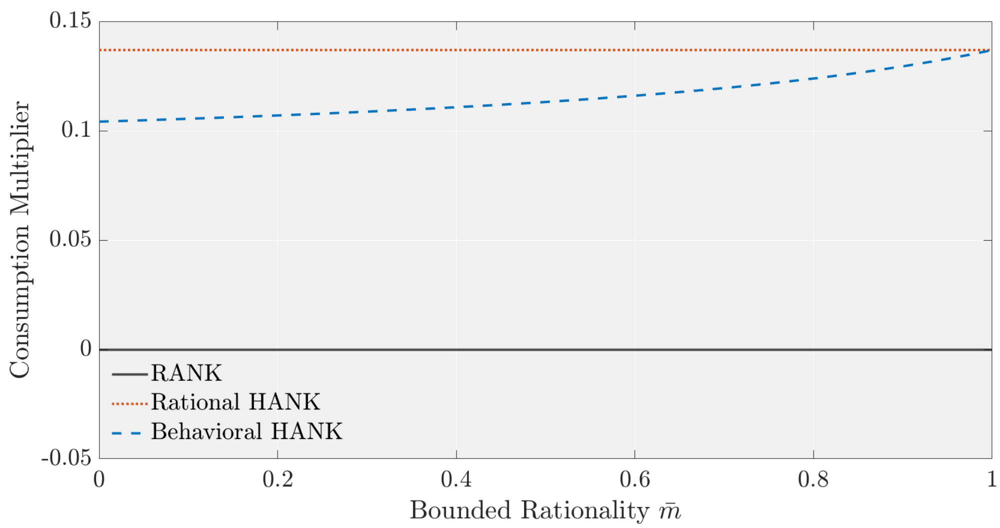

Note: This figure shows the consumption multipliers (the consumption response to government spending) for different degrees of bounded rationality (blue-dashed line). The orange-dotted line plots the multiplier in the rational version of the model and the black-solid line shows the zero-multiplier in 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 (2022)). 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.[5] Bounded 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.

A.4.2 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

$$\widehat{y}_t = \bar{m}\tilde{\delta}\mathbb{E}_t\widehat{y}_{t+1} - \frac{1}{\gamma}\frac{1-\lambda}{1-\lambda\chi}\left(\widehat{i}_t - \mathbb{E}_t\pi_{t+1}\right), \tag{A.35}$$

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 1.3 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.

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

A.4.3 Persistent Monetary Policy Shocks

In the main text in Section 1.3, 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.

Figure A.5 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 A.5: Initial Output Response for Varying Degrees of the Persistence

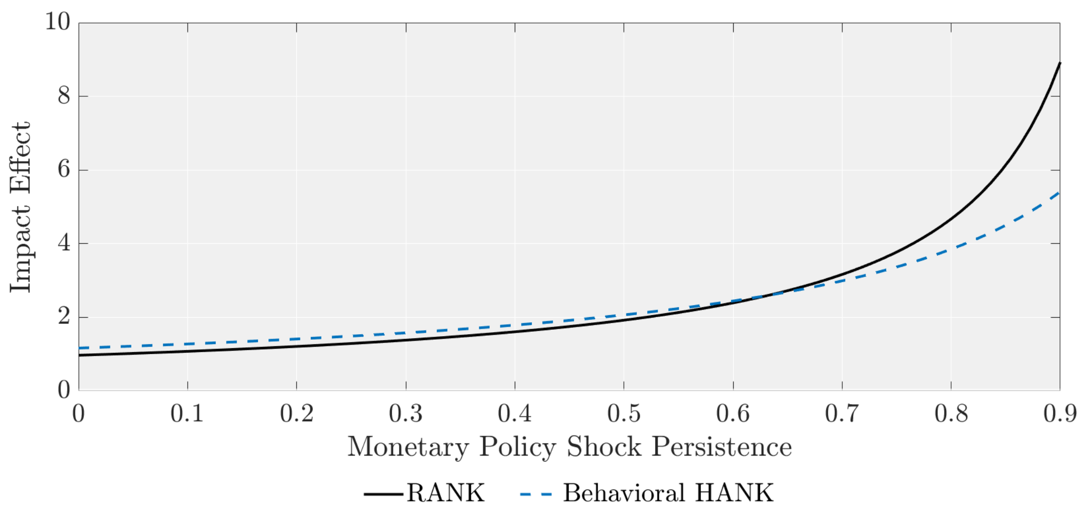

Note: This figure shows the initial output response to monetary policy shocks with different degrees of persistence.

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

$$\widehat{y}_t = -\frac{\frac{\psi_c}{\gamma}}{1 + \frac{\psi_c}{\gamma}\phi\kappa - \left(\psi_f + \psi_c\frac{\kappa}{\gamma}\right)\rho}\varepsilon_t^{MP}. \tag{A.36}$$

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 $\rho = 1$.

A.4.4 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:

$$\begin{pmatrix} \mathbb{E}_t \pi_{t+1} \\ \mathbb{E}_t \widehat{y}_{t+1} \end{pmatrix} = \underbrace{\begin{pmatrix} \frac{1}{\beta M^f} & -\frac{\kappa}{\beta M^f} \\ \frac{\psi_c}{\gamma \psi_f}\left(\phi - \frac{\bar{m}}{\beta M^f}\right) & \frac{1}{\psi_f}\left(1 + \frac{\psi_c \bar{m}\kappa}{\gamma \beta M^f}\right) \end{pmatrix}}_{\equiv A} \begin{pmatrix} \pi_t \\ \widehat{y}_t \end{pmatrix}. \tag{A.37}$$

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).

A.5 Quantitative Behavioral HANK Model

Table A.3 shows how we calibrate the quantitative model introduced in Section 1.4. 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 1.5, we discuss less-progressive tax systems as well.

Patterson (2022) 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 A.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

A.5.1 Sticky Wages

In our baseline model, we assume that prices are sticky and wages fully flexible. We now show that our results are robust if we instead introduce sticky wages and keep prices fully flexible.[6] Given that prices are fully flexible, we also abstract from monopolistic competition of firms, that is, prices are set to marginal costs.

Labor hours N_{it} are determined by union labor demand. Each worker provides N_{ikt} hours of work to a continuum of unions indexed by k. Each union aggregates efficient units of work into a union-specific task. A competitive labor packer then packages these tasks into aggregate employment services according to a CES technology and sells these services to final goods firms at price W_t. We assume that there are quadratic utility costs of adjusting the nominal wage W_{kt}. A union sets a common wage W_{kt} per efficient unit for each of its members. In doing so, the union trades-off the marginal disutility of working given average hours against the marginal utility of consumption given average consumption. The union then calls upon its members to supply hours according to a specific allocation rule: in stationary equilibrium all households' supply the same amount of hours. Outside stationary equilibrium, we assume that each households labor supply is a linear feedback function of changes in aggregate hours according to her productivity:

$$N_{it} = \eta(e)(N_t - \bar{N}) + \bar{N}. \tag{A.38}$$

We assume that $\eta(e_1) = 3, \eta(e_2) = 1, \eta(e_3) = 0.5$. This implies that the labor supply of households with lower idiosyncratic productivity are more exposed to the business cycle. Given that dividends are zero, this also implies fact (ii), namely, that the income of households with higher MPCs are more exposed to changes in aggregate income induced by monetary policy.

All in all, our setup leads to a wage Philips curve given by:

$$\pi_t^W = \kappa\Big(v'(N_t) - (\epsilon_n - 1)/\epsilon_n(1 - \tau_t)\frac{W_t}{P_t}u'(C_t)\Big) + \beta\pi_{t+1}^W, \tag{A.39}$$

where $\epsilon_n = 11$ is the elasticity of substitution between differentiated labor supply and $\kappa = 0.1$ is the slope of the wage Philips Curve.

Figure A.6 shows our main results in our sticky-wage behavioral HANK model. We see that our results are robust. The behavioral HANK model rules out the forward guidance puzzle whereas it is aggravated in the rational HANK model compared to the representative agent model.

[6] Auclert et al. (2021) and Broer et al. (2020) argue in favor of using sticky wages rather than sticky prices in HANK models.

Figure A.6: Sticky Wages

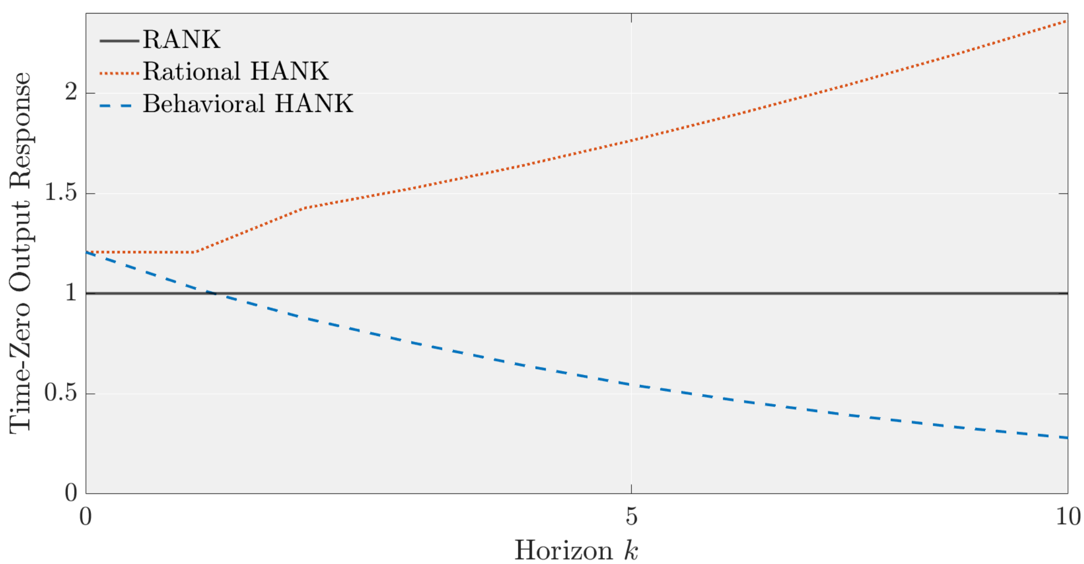

Note: This figure shows the response of total output in period 0 to anticipated i.i.d. monetary policy shocks occurring at different horizons k for the case in which prices are flexible and wages are sticky.

A.5.2 Robustness of Calibration

In our baseline calibration, the unequal dividend shares imply an amplification of contemporaneous monetary policy shocks of 20% compared to the representative agent model. If, instead, we assume a dividend distribution that is even more tilted towards high-productivity households (instead of distributing 40% of dividends to the medium-productivity households and 60% to the high-productivity households, we now assign 20% to the medium and 80% to high-productivity households) we obtain an amplification of 40%. Figure A.7 shows that in this case, the forward guidance puzzle is substantially exacerbated in the rational HANK model while it is still ruled out in the behavioral HANK model.

Figure A.7: Robustness of Calibration

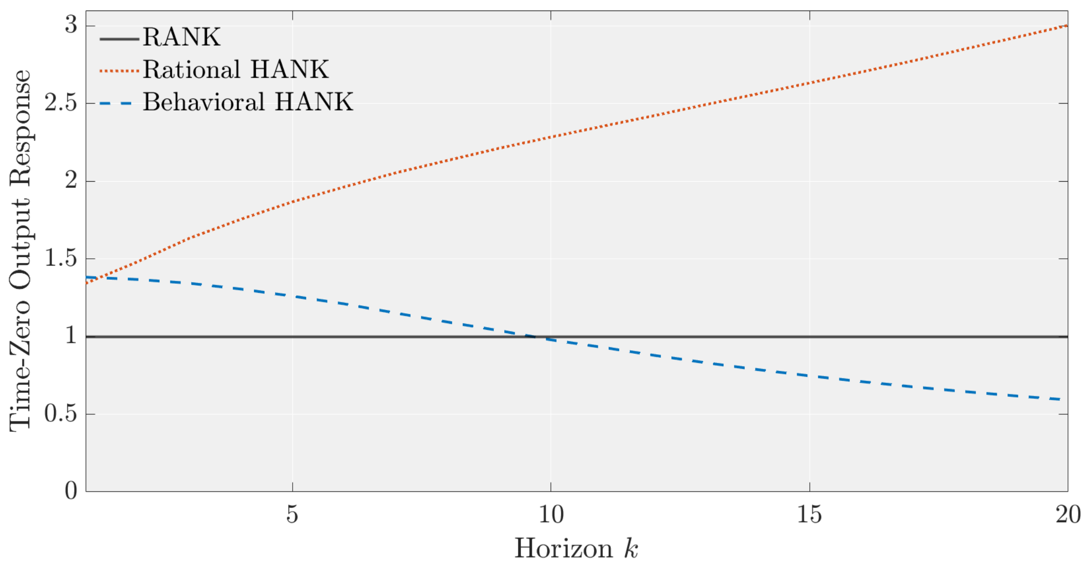

Note: This figure shows the response of total output in period 0 to anticipated i.i.d. monetary policy shocks occurring at different horizons k for the case in which the unequal dividend shares imply an amplification of 40% compared to the representative agent model.

A.5.3 Stability at the ELB and Fiscal Multipliers

Figure A.8 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 A.8: ELB recession in the quantitative behavioral HANK model

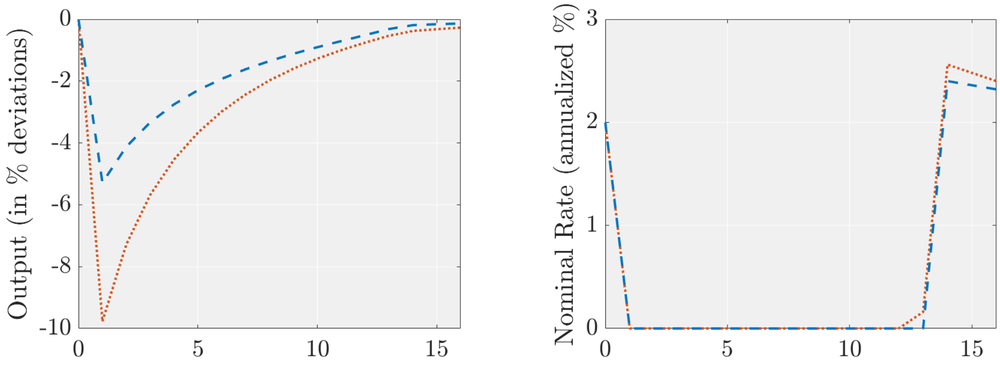

Note: This figure shows the impulse responses of total output and of the nominal interest rate after a discount factor shock that brings the economy to the ELB for 12 quarters.

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 1.3). 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 a positive consumption multiplier under a constant real rate, we redo the experiments from Section A.4.1: 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 A.9 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. (2022)).

Figure A.9: Consumption multiplier in the quantitative behavioral HANK

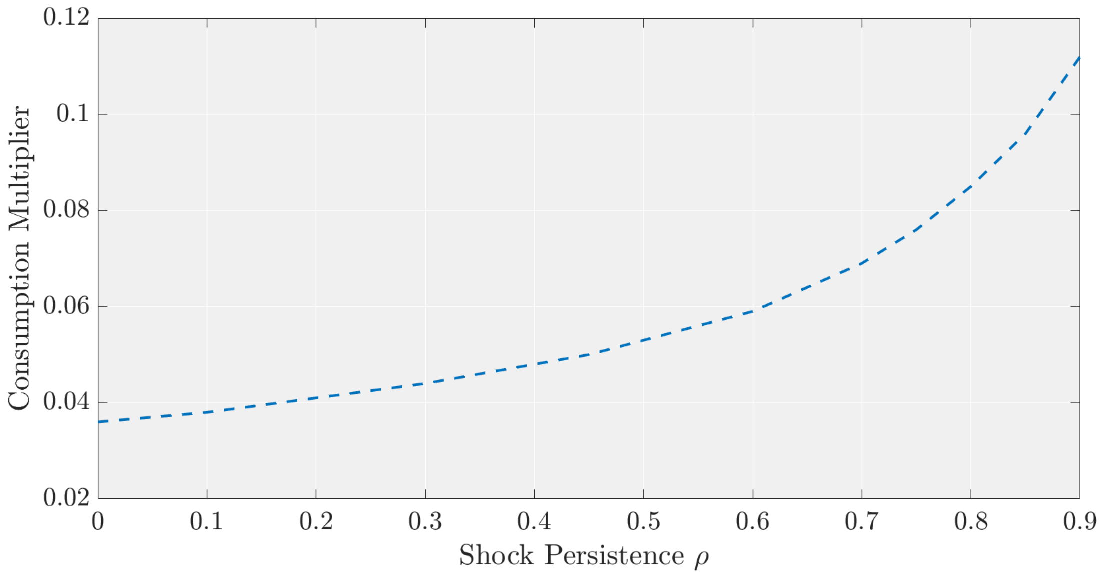

Note: This figure shows the impact consumption multiplier after an exogenous increase in government consumption which is financed by lump-sum taxes for various degrees of persistence.

A.5.4 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 central bank 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 A.10 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 A.10: Heterogeneous $\bar{m}$ and Monetary Policy

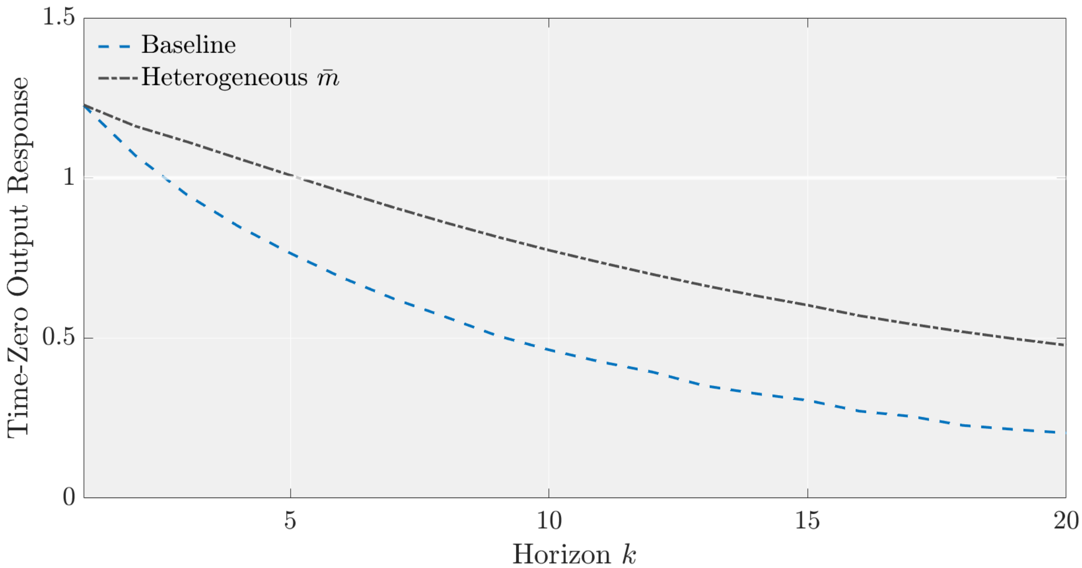

Note: This figure shows the response of total output in period 0 to anticipated i.i.d. monetary policy shocks occurring at different horizons k for the baseline calibration with $\bar{m} = 0.85$ for all households (blue-dashed line), and for the model in which high productivity households have $\bar{m}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

A.6 Additional Results and Figures to Section 1.5

A.6.1 Decomposition of Amplification Channel: Alternative Calibration

As shown in Section 1.5, 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 1.2). 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 A.11 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 than twice as large. Given our discussion in Section 1.5, 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 A.11: Decomposition of the Additional Inflation Increase: Alternative Calibration

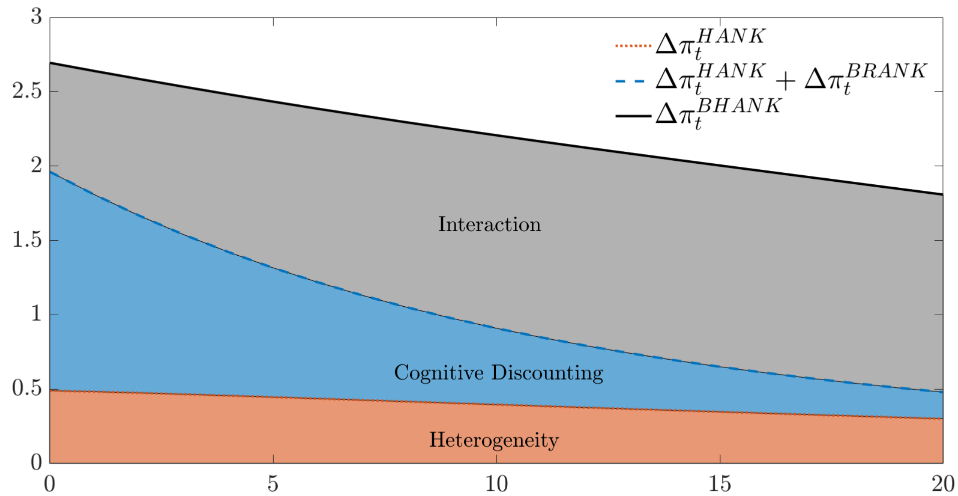

Note: This figure shows the additional inflation increase in the rational HANK model compared to the rational RANK model (orange-dotted line), the sum of the additional increase in the rational HANK and the behavioral RANK compared to the rational RANK model (blue-dashed line) and the additional increase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

A.6.2 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 1.4. 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 A.12 and A.13 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 A.12, 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 1.3). 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 A.13). 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 A.12: Procyclical inequality, strict inflation-targeting

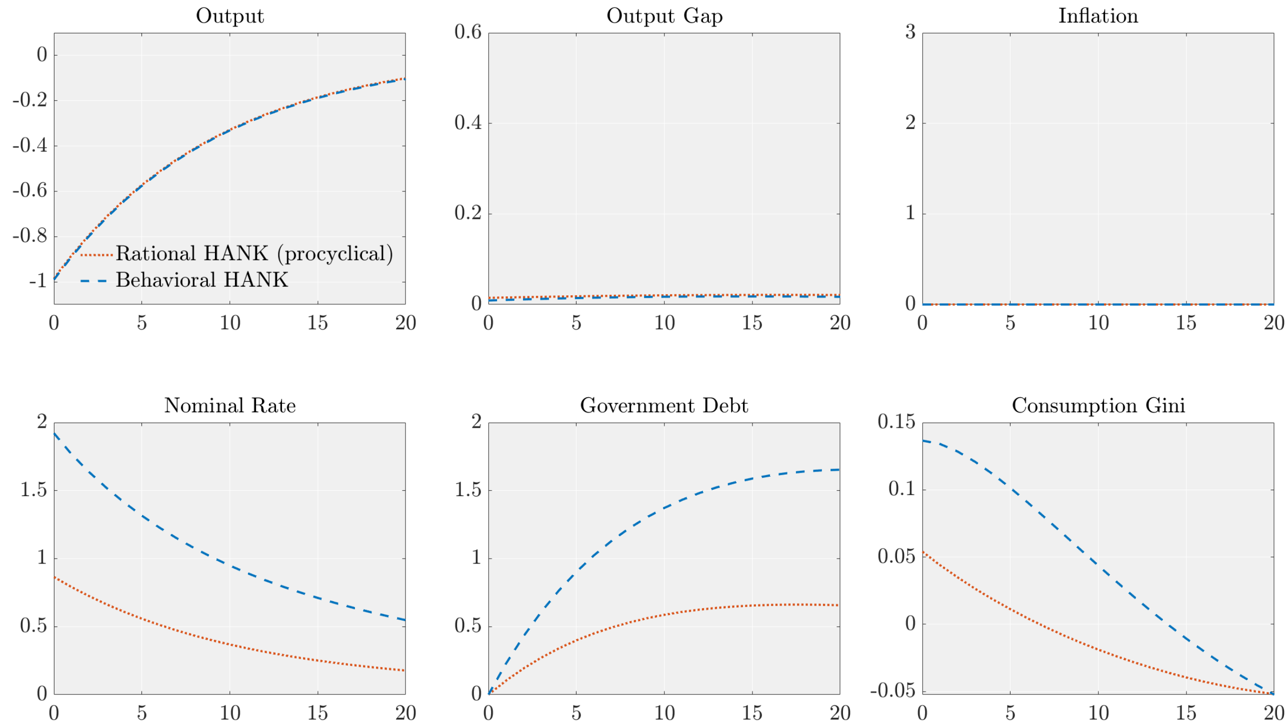

Note: This figure shows the impulse responses after a TFP shock that decreases potential output by 1% in the inflation-stabilizing monetary policy regime for the behavioral HANK model (blue-dashed lines) and for the rational HANK model with procyclical inequality (orange-dotted lines). Output and the 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 A.13: Procyclical inequality, Taylor rule

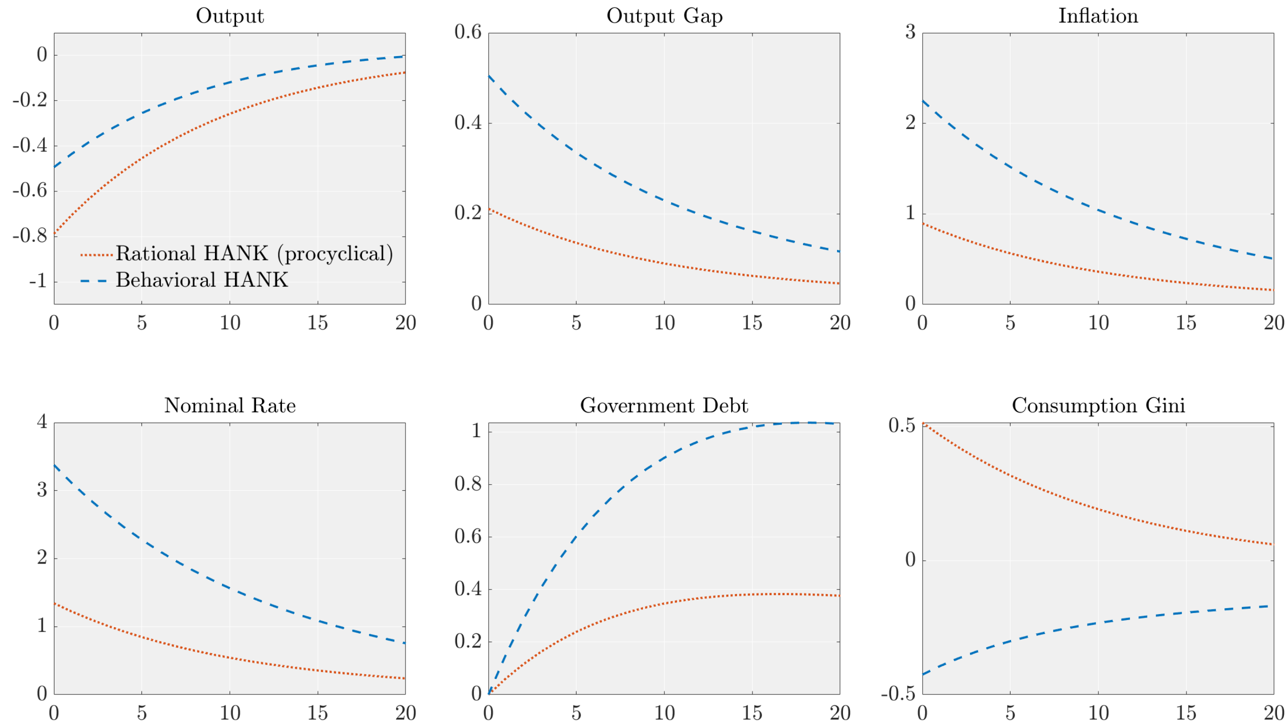

Note: This figure shows the impulse responses after a TFP shock that decreases potential output by 1% when monetary policy follows a Taylor rule for the behavioral HANK model (blue-dashed lines) and for the rational HANK model with procyclical inequality (orange-dotted lines). Output and the output gap 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.

A.6.3 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 A.14 and A.15 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 A.14: Representative Household, strict inflation targeting

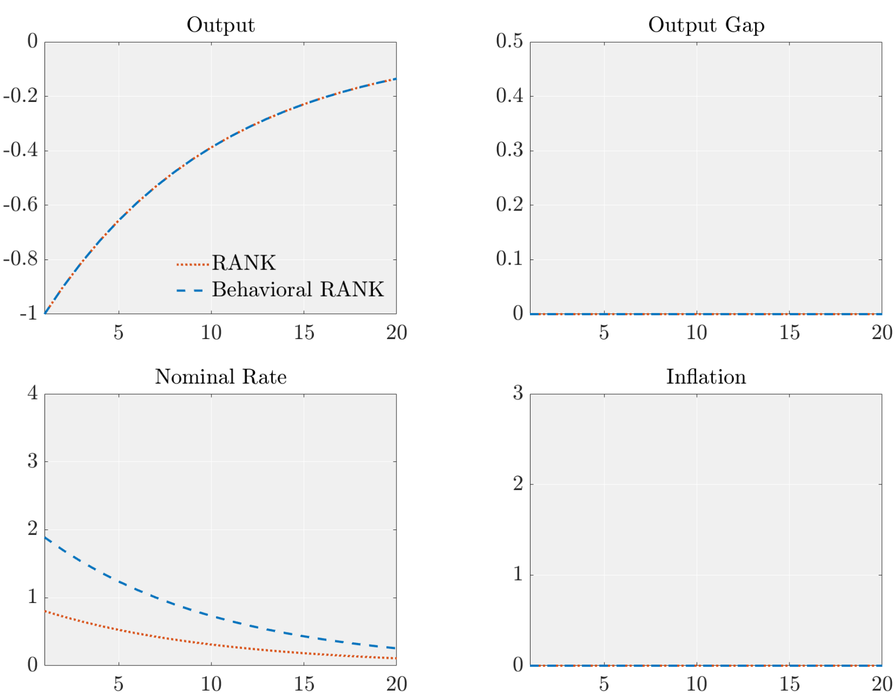

Note: This figure shows the impulse responses after a TFP shock that decreases potential output by 1% in the inflation-stabilizing monetary policy regime for the behavioral RANK model (blue-dashed lines) and for the rational RANK model (orange-dotted lines). Output and the output gap are shown as percentage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

Figure A.15: Representative Household, Taylor rule

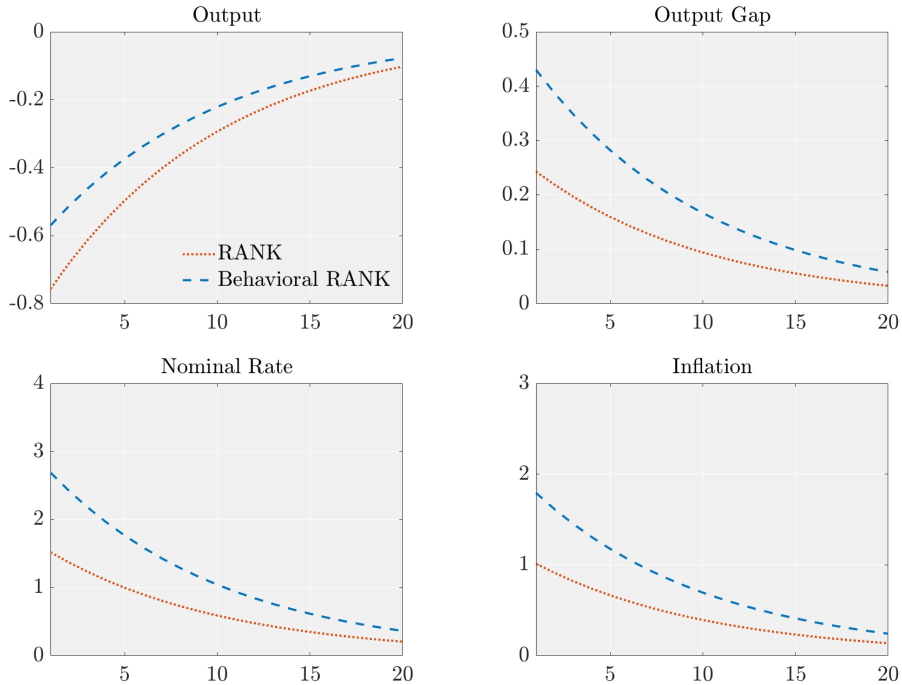

Note: This figure shows the impulse responses after a TFP shock that decreases potential output by 1% when monetary policy follows a Taylor rule for the behavioral RANK model (blue-dashed lines) and for the rational RANK model (orange-dotted lines). Output and the output gap are shown as percentage deviations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

A.6.4 Heterogeneous $\bar{m}$

In the main text, we consider the case of homogeneous cognitive discounting. As discussed in Section 1.4.2, 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 1.4.2, we find that the trade-off between price stability and distributional considerations is slightly weaker.

Figures A.16 and A.17 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 A.16: Heterogeneous $\bar{m}$, strict inflation-targeting

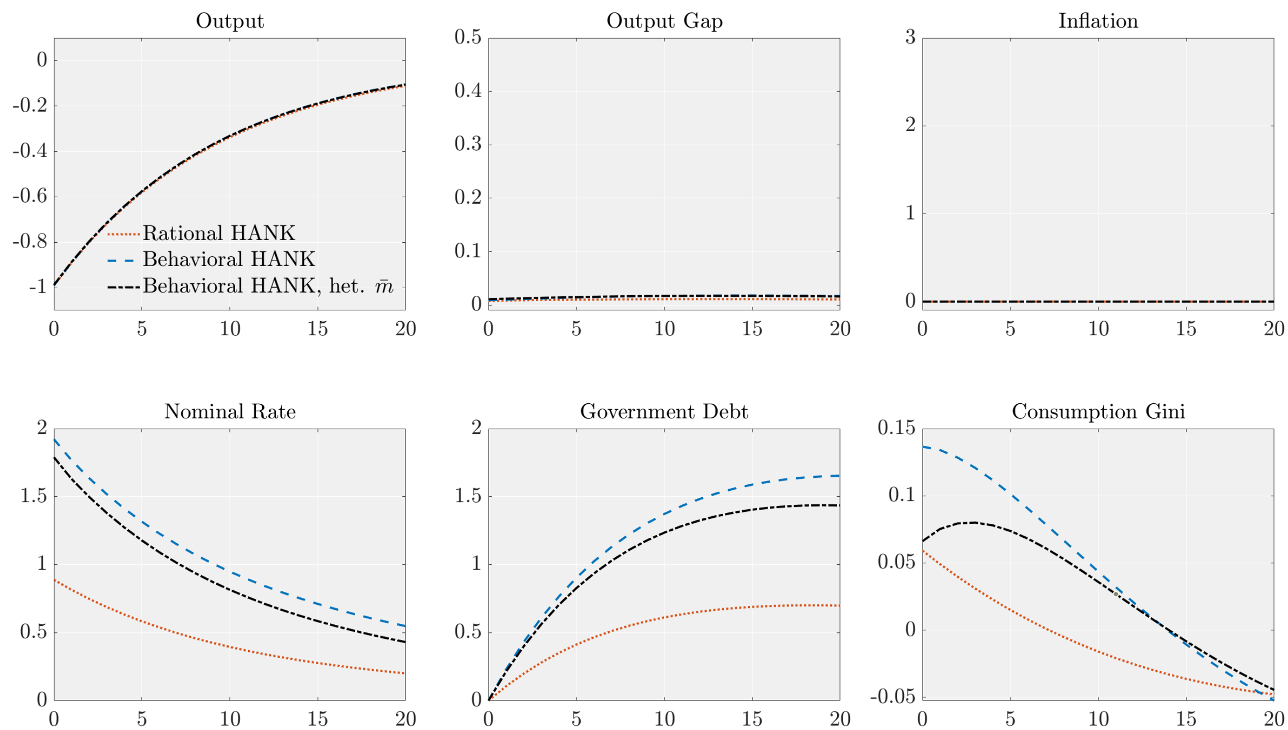

Note: This figure shows the impulse responses after a TFP shock that decreases potential output by 1% in the inflation-stabilizing monetary policy regime and heterogeneous $\bar{m}$ (black-dashed-dotted lines). Output and the output gap are shown as percentage deviations from steady state output, the 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 A.17: Heterogeneous $\bar{m}$, Taylor rule

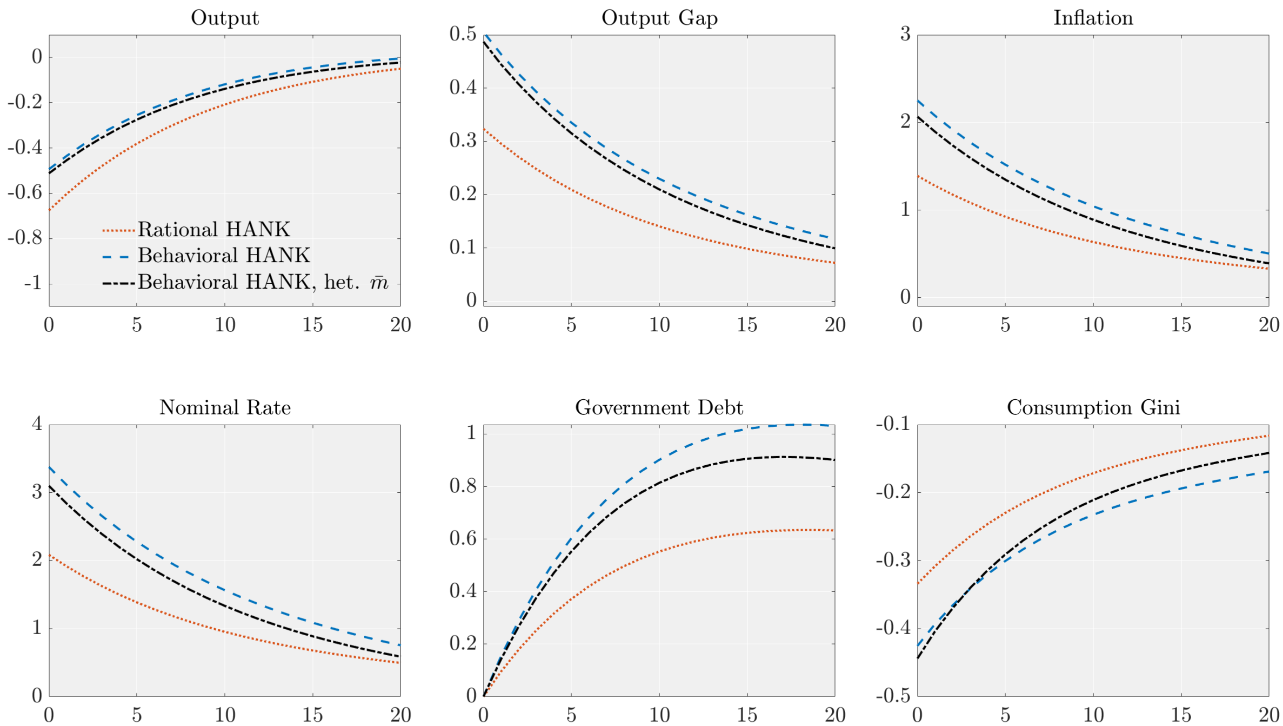

Note: This figure shows the impulse responses after a TFP shock that decreases potential output by 1% when monetary policy follows a Taylor rule and heterogeneous $\bar{m}$ (black-dashed-dotted lines).. Output and the output gap are shown as percentage deviations from steady state output, the nominal 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.

A.6.5 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 1.5 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 A.18 and A.19 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 A.18: Heterogeneous $\bar{m}$, strict inflation-targeting and less progressive taxes

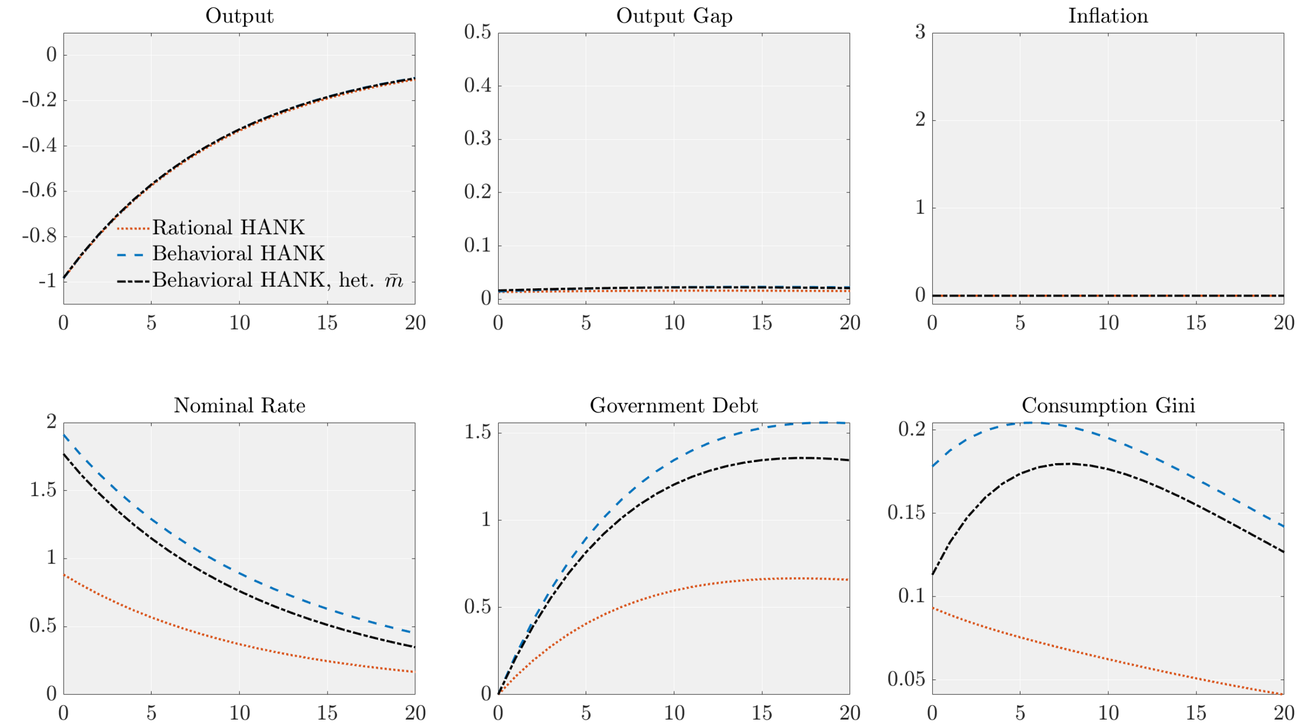

Note: This figure shows the impulse responses after a TFP shock that decreases potential output by 1% in the inflation-stabilizing monetary policy regime with a less-progressive tax system than in the baseline calibration. Output and the output gap are shown as percentage deviations from steady state 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 A.19: Heterogeneous $\bar{m}$, Taylor rule and less progressive taxes

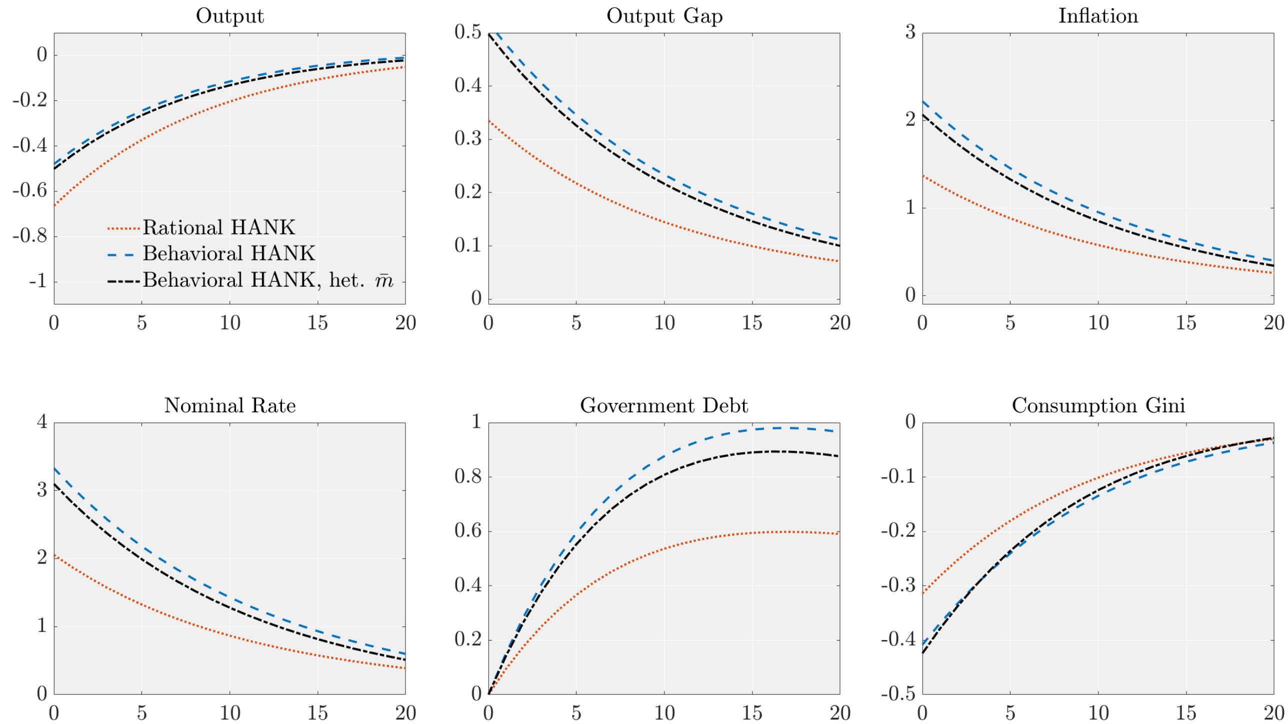

Note: This figure shows the impulse responses after a TFP shock that decreases potential output by 1% when monetary policy follows a Taylor rule with a less-progressive tax system than in the baseline calibration. Output and the output gap are shown as percentage deviations from steady state output, 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.

A.6.6 Productivity Shock with High Initial Debt Levels

Figure A.20 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 1.3, 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 A.20: Inflationary productivity shock with high debt: strict inflation targeting

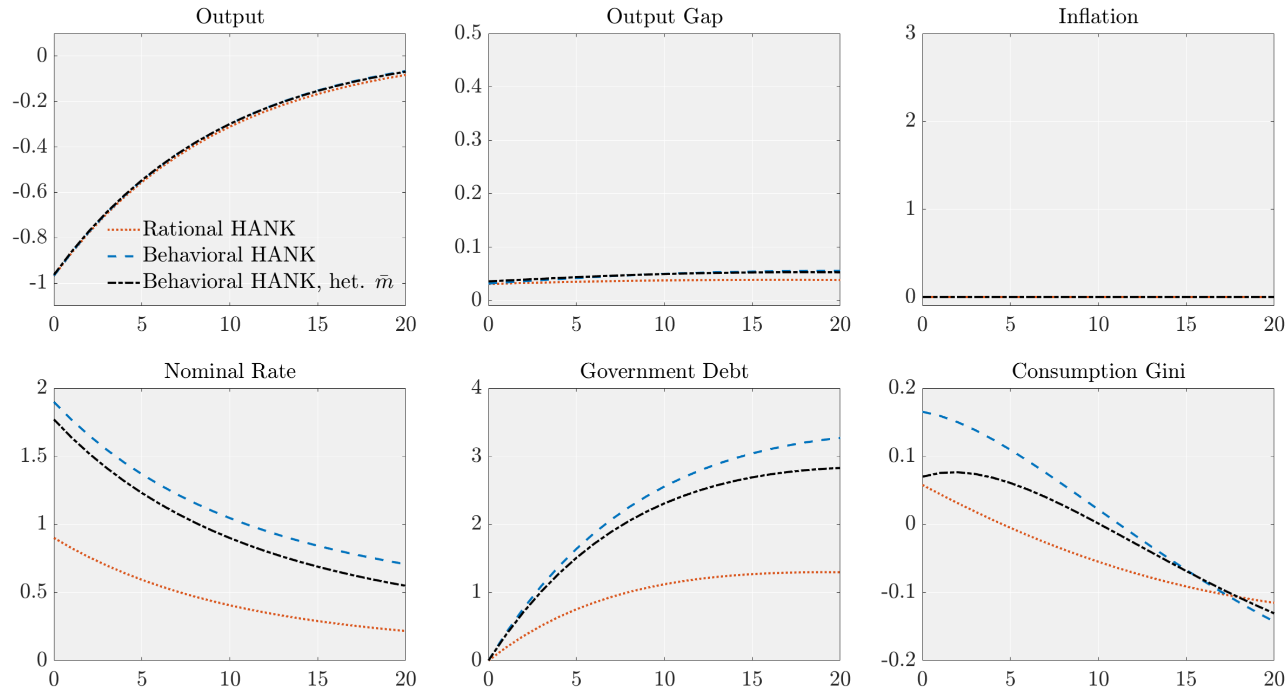

Note: This figure shows the impulse responses after a TFP shock that decreases potential output by 1% in the inflation-stabilizing monetary policy regime when the initial debt level is 90% instead of 50% of annual GDP. Output and the output gap are shown as percentage deviations from steady state output, 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 A.21: Inflationary productivity shock with high debt: Taylor rule

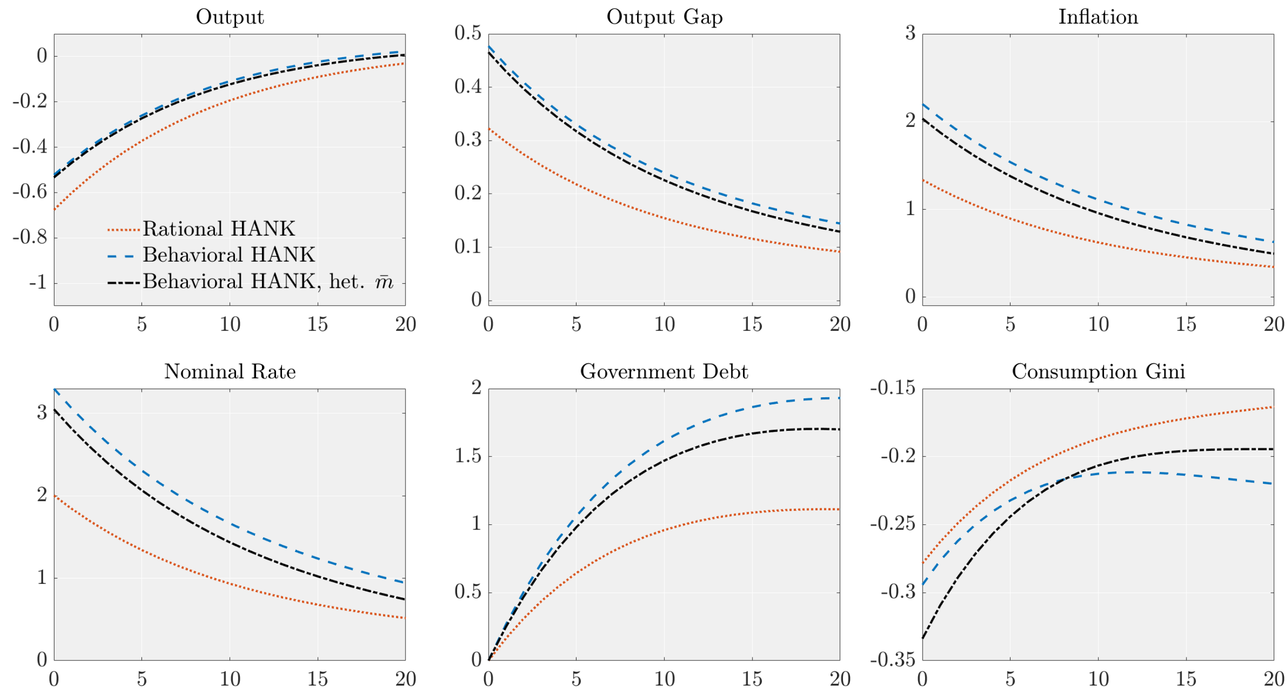

Note: This figure shows the impulse responses after a TFP shock that decreases potential output by 1% in the inflation-stabilizing monetary policy regime when the initial debt level is 90% instead of 50% of annual GDP. Output and the output gap are shown as percentage deviations from steady state output, 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.

A.6.7 Behavioral Firms

Figure A.22 shows the impulse-response functions after a negative productivity shock when monetary policy follows a Taylor rule and in which firms are behavioral (with a cognitive discounting factor of 0.85). We see that the increase in inflation when monetary policy follows a Taylor rule is somewhat muted whereas the increase in the output gap is amplified compared to the case in which firms are rational. The reason is that firms discount the increase in their future marginal costs and thus increase their prices not as strongly. According to the Taylor rule this then leads to a smaller increase in interest rates so that households consume more, leading to an increase in demand and thus, the output gap.

Figure A.22: Inflationary supply shock: behavioral firms

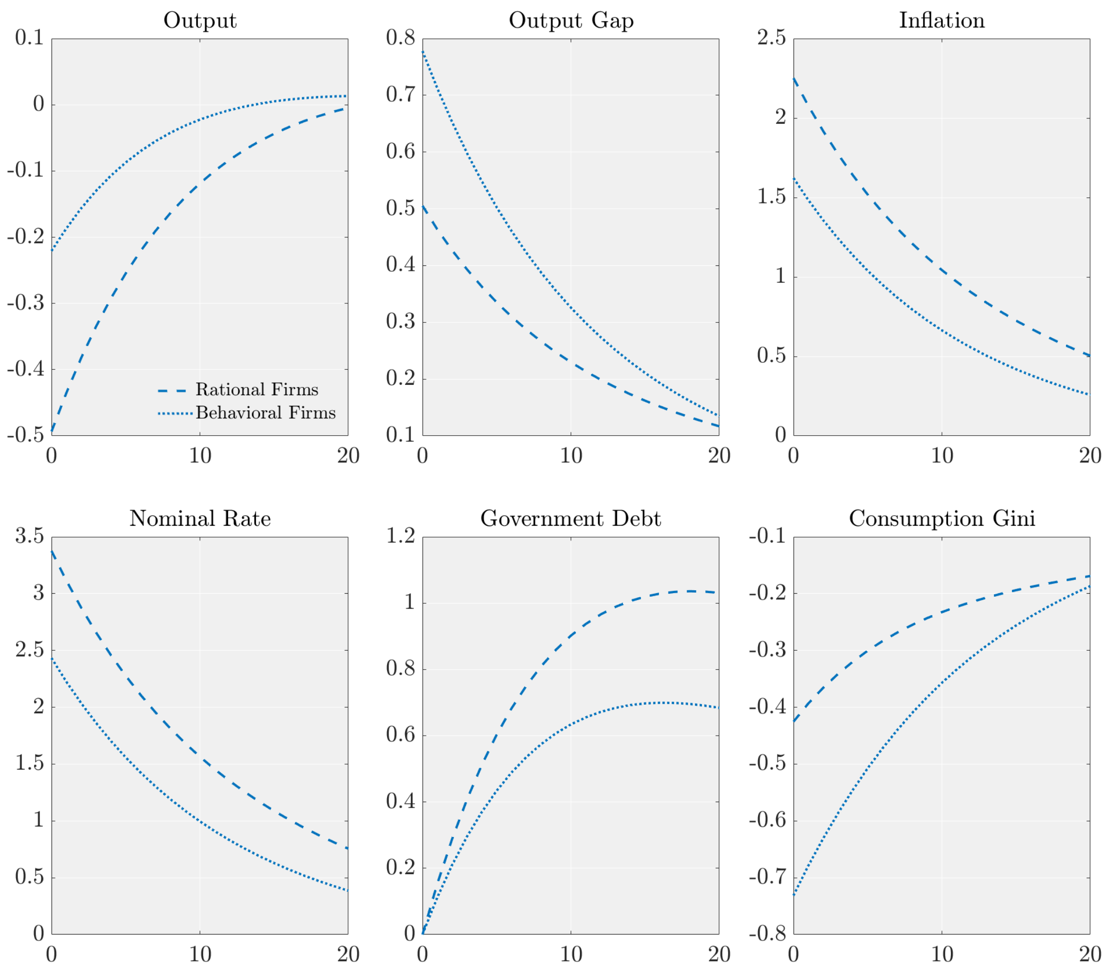

Note: This figure shows the impulse responses after a productivity shock for the case that monetary policy follows a Taylor rule and firms cognitively discount the future with a cognitive discounting parameter of 0.85. Output and the output gap are shown as percentage deviations from steady state output, 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.

A.6.8 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 1.5. Note, that we assume that the shock also applies to the model under flexible prices, thus moves potential output as well.

Figure A.23 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 1.5. 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 A.23: Inflationary cost-push shock: strict inflation targeting

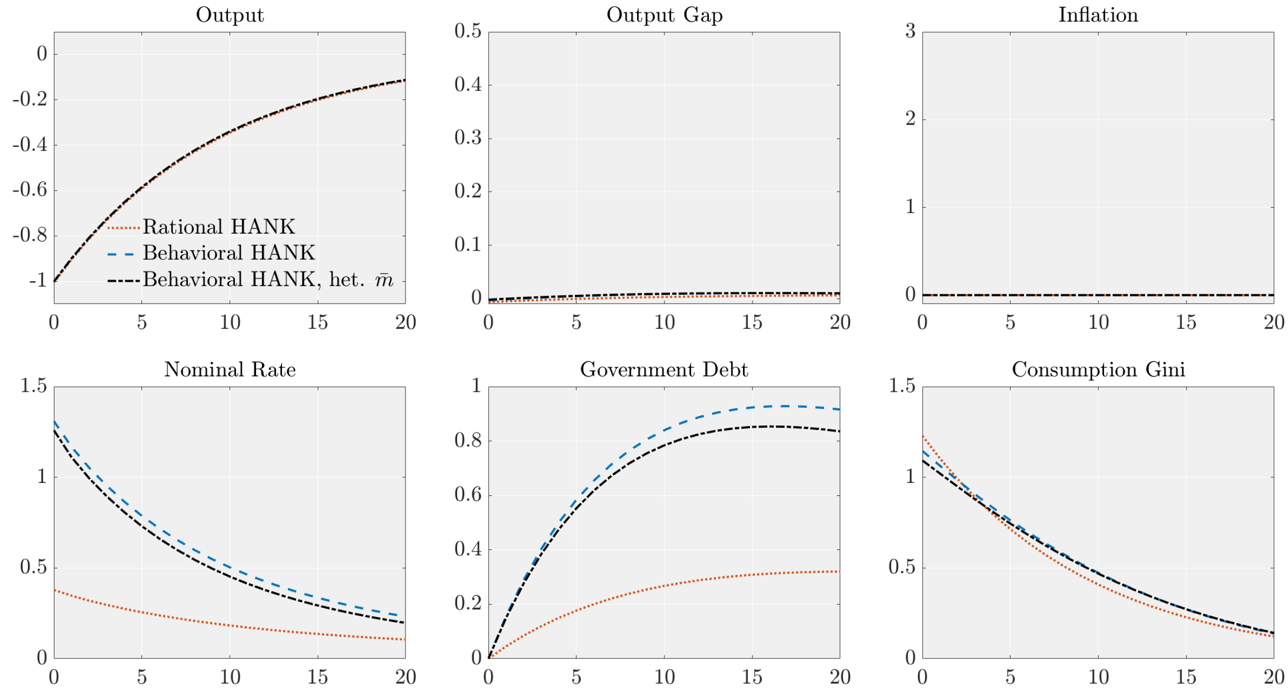

Note: This figure shows the impulse responses after a cost-push shock that decreases potential output by 1% in the inflation-stabilizing monetary policy regime. Output and the output gap are shown as percentage deviations from steady state output, the nominal interest rate and inflation as annualized 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 A.23 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 A.24 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 1.3 and 1.4 do, thereby redistributing towards lower income households which dampens the increase in consumption inequality.

Figure A.24: Inflationary cost-push shock: Taylor rule

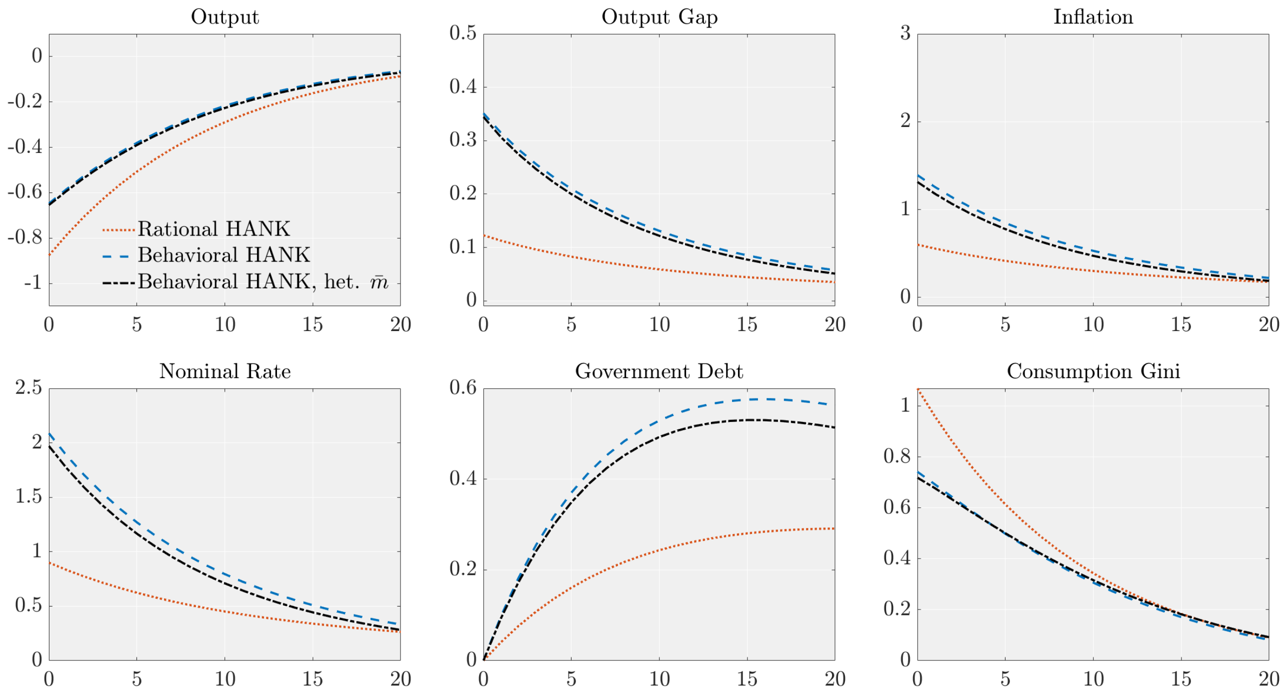

Note: This figure shows the impulse responses after a cost-push shock that decreases potential output by 1% in the Taylor rule monetary policy regime. Output and the output gap are shown as percentage deviations from steady state output, the nominal interest rate and inflation as annualized percentage 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.

A.7 Details on Intertemporal MPCs

In this section, we derive the iMPCs discussed in Section 1.6.1. 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

$$\widehat{c}_t^H = \widehat{y}_t^H + \frac{1-s}{\lambda} \beta^{-1} b_t \tag{A.40}$$
$$\widehat{c}_t^U + \frac{1}{1-\lambda} b_{t+1} = \widehat{y}_t^U + \frac{s}{1-\lambda} \beta^{-1} b_t, \tag{A.41}$$

where b_t denotes real bonds in shares of steady state output. Aggregating (A.40) and (A.41) delivers

$$\widehat{c}_t = \tilde{y}_t + \beta^{-1} b_t - b_{t+1}, \tag{A.42}$$

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

By plugging equations (A.40) and (A.41) into the Euler equation of unconstrained households (1.12), we can derive the dynamics of liquid assets b_t (ignoring changes in the real rate as this is a partial equilibrium exercise):

$$\begin{aligned}\mathbb{E}_t b_{t+2} - b_{t+1}\left[\frac{1}{s\bar{m}} + \beta^{-1}s + \frac{(1-s)^2\beta^{-1}(1-\lambda)}{s\lambda}\right] + \frac{\beta^{-1}}{\bar{m}} b_t = \\ (1-\lambda)\mathbb{E}_t \widehat{y}_{t+1}^U + \frac{1-s}{s}(1-\lambda)\mathbb{E}_t \widehat{y}_{t+1}^H - \frac{1-\lambda}{s\bar{m}} \widehat{y}_t^U.\end{aligned} \tag{A.43}$$

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 (A.43) by $-\mathbb{E}_t \widehat{z}_t$. Factorizing the left-hand side and letting F denote the forward-operator, it follows that

$$(F - \mu_1)(F - \mu_2)\mathbb{E}_t b_t = -\mathbb{E}_t \widehat{z}_t, \tag{A.44}$$

where μ_1 and μ_2 denote the roots of the characteristic equation

$$\mathbb{E}_t b_{t+2} - \phi_1 b_{t+1} - \phi_2 b_t = 0, \tag{A.45}$$

where

$$\phi_1 \equiv \left[\frac{1}{s\bar{m}} + \beta^{-1}s + \frac{(1-s)^2\beta^{-1}(1-\lambda)}{s\lambda}\right] \tag{A.46}$$

and

$$\phi_2 \equiv -\frac{\beta^{-1}}{\bar{m}}. \tag{A.47}$$

Thus, the roots are given by

$$\mu_{1,2} = \frac{\phi_1 \pm \sqrt{\phi_1^2 + 4\phi_2}}{2}. \tag{A.48}$$

It follows that

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

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

$$b_{t+1} = \mu_1 b_t + \frac{1-\lambda\chi}{s}\sum_{l=0}^{\infty}\mu_2^{-(l+1)}\mathbb{E}_t\left(\frac{1}{\bar{m}}\widehat{y}_{t+l} - \delta\widehat{y}_{t+1+l}\right). \tag{A.49}$$

Plugging this in equation (A.42) and taking derivatives with respect to $\widehat{y}_{t+k}$ yields Proposition 4.

iMPCs for more than two periods. Figure A.25 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 1.6.1, 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 A.26 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.

iMPCs and the role of idiosyncratic risk. In Figure A.26, 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 1.6.1. 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

Figure A.25: Intertemporal MPCs

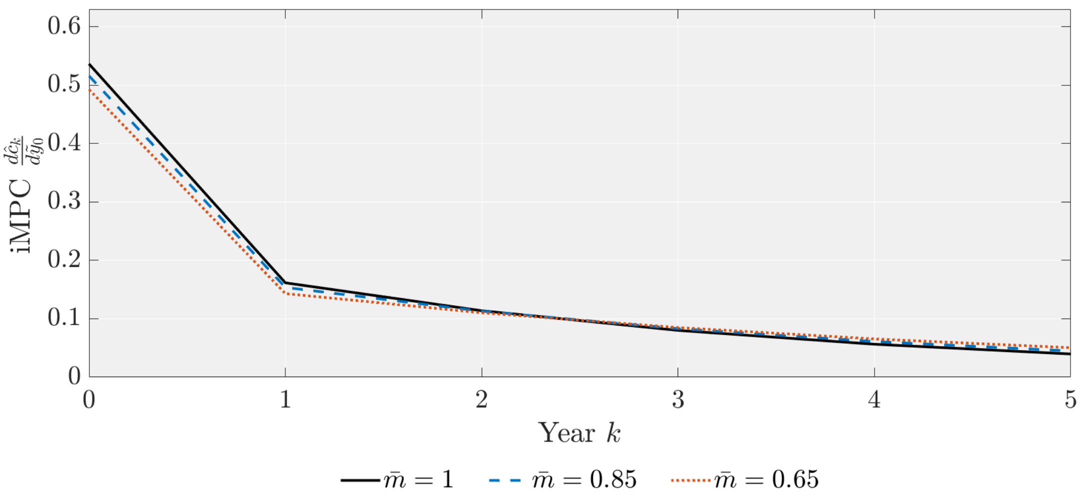

Note: This figure shows the aggregate intertemporal MPCs, i.e., the aggregate consumption response in year k to a change in aggregate disposable income in year 0 for different $\bar{m}$.

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 $1 - s$.

Figure A.26: Intertemporal MPCs, Bounded Rationality and Household Heterogeneity

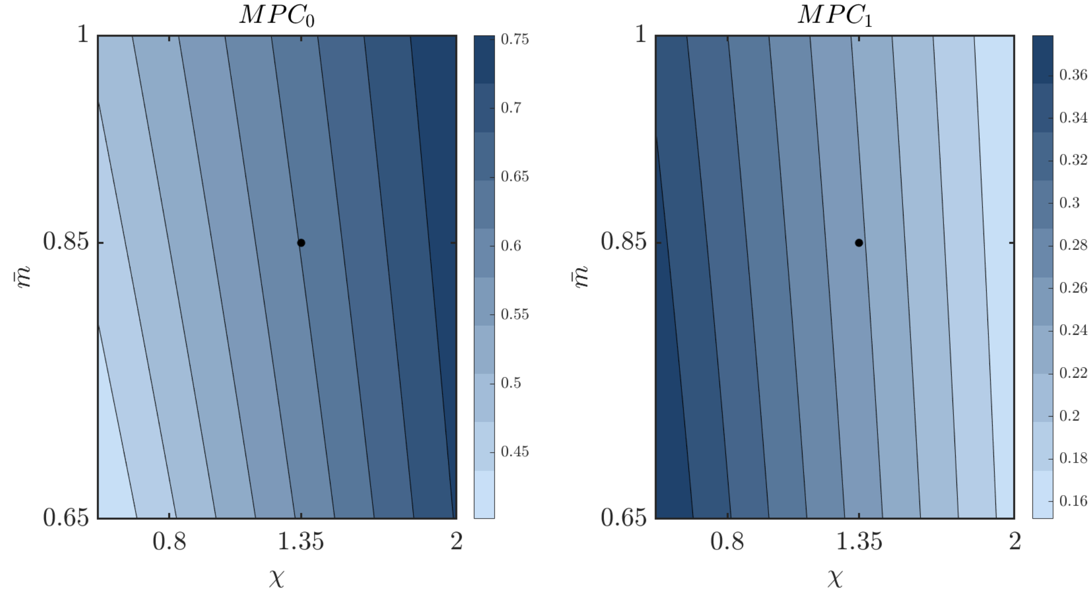

Note: This figure shows the aggregate intertemporal MPCs, i.e., the aggregate consumption response in year 0 (left) and year 1 (right) to a change in aggregate disposable income in year 0 for a transition probability $1 - s = 0.5$.

A.8 Sticky Wages

In this section, we provide details on the sticky-wage extension presented in Section 1.6.2 as well as the calibration used to produce Figures 1.7 and 1.8. The way we introduce sticky wages follows Colciago (2011) and recently adopted by Bilbiie et al. (2022).[7]

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 1.2. 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 A.4 summarizes all equilibrium equations.

The calibration of this extended model is presented in Table A.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. (2022) 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

[7] 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 A.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 A.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

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 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.

Chapter 2

Subjective Housing Price Expectations, Falling Natural Rates and the Optimal Inflation Target

Joint with Klaus Adam and Timo Reinelt

2.1 Introduction

The large and sustained booms and busts in housing prices in advanced economies are often attributed to households holding excessively optimistic or pessimistic beliefs about future housing prices (Piazzesi and Schneider (2006), Kaplan et al. (2020)). This view is supported by a nascent literature that documents puzzling facts about the behavior of housing price expectations. Survey measures of expected future housing prices have been found to be influenced by past housing price changes, but appear to underreact to these changes, and they also miss the tendency of housing prices to mean revert over time (Kuchler and Zafar (2019), Case et al. (2012), Ma (2020) and Armona et al. (2019)).

Documenting in which ways households' housing price expectations deviate from the rational expectations benchmark is an important task but remains in itself uninformative about how important the observed deviations are for economic outcomes in housing markets and for the conduct of monetary policy. Understanding these features requires a structural equilibrium model that quantitatively replicates how households' expectations deviate from rational expectations. Constructing such an equilibrium model, calibrating it to the behavior of household beliefs in survey data, and understanding its implications for the optimal design of monetary policy is the main objective of the present paper. To the best of our knowledge, it is the first paper pursuing this task.

We begin our analysis by comprehensively quantifying the dimensions along which households' housing price expectations deviate from the full-information rational expectations benchmark. To this end, we consider the Michigan Survey of Consumers, which provides the longest available time series of quantitative housing price expectations for the United States, covering the years 2007-2021.

We document three dimensions along which household expectations deviate from rational expectations. First, expectations about future housing prices are revised too sluggishly over time, a feature that housing price expectations share with other household expectations (Coibion and Gorodnichenko (2015a)). Second, households' capital gain expectations covary positively with market valuation, i.e., the price-to-rent ratio, while actual future capital gains covary negatively with market valuation. We show that the difference is striking, highly statistically significant, and in line with findings on investor expectations in stock markets (Adam et al. (2017)).[1] Third, in a dynamic sense, households' capital gain expectations initially underreact to observed capital gains, i.e., households are too pessimistic in the first few quarters following a positive capital gain, but later on overreact, i.e., households hold too optimistic expectations after about twelve quarters. The pattern of initial underreaction and subsequent overreaction is similarly present in other macroeconomic expectations, see Angeletos et al. (2021).

While the first and third deviation from rational housing price expectations have been documented before using different data sets, see Armona et al. (2019), the second deviation from rational expectations is new to the housing literature. We quantify here all three deviations using a single data set, so as to obtain a coherent set of quantitative targets for our structural equilibrium model with subjective housing expectations.

Equipped with these facts, we construct first a simple housing model with optimizing households that hold subjective beliefs about housing price behavior. Bayesian belief updating implies that households weakly extrapolate past capital gains into the future. The model reproduces – as an equilibrium outcome – important patterns of the behavior of U.S. housing prices, in particular, the large and protracted swings in the price-to-rent ratio over time, as well as the three dimensions mentioned above along which household expectations deviate from the rational expectations benchmark. The quantitative fit is surprisingly good, despite the simplicity of the model.

The simple model generates two important insights. First, it shows that the standard deviation for the price-to-rent ratio would be much lower in the presence of rational housing expectations. This suggests that the observed volatility of housing prices is to a significant extent due to the presence of subjective beliefs. This lends credence to the view that the observed deviations from rational expectations substantially contribute to booms and busts in housing markets.

Second, the simple model connects the secular decline in natural rates of interest with higher volatility of housing prices. Specifically, the model predicts that lower real interest rates imply larger effects of belief fluctuations on equilibrium housing prices. This prediction does not emerge

[1]For stock markets, Adam et al. (2021) show that this cannot be explained by investors reporting risk-adjusted expectations.

in the presence of rational housing price expectations, but is consistent with the data. We show that in a number of advanced economies, including the United States, the volatility of housing prices has increased considerably at the same time as the level of the natural rate of interest has fallen.

The most important objective of paper is to understand the monetary policy implications generated by a setting where households (weakly) extrapolate capital gains into the future. We are particularly interested in the optimal policy response to increased housing price volatility that is induced by falling natural rates of interest in a setting where policy rates cannot move into negative territory. To this end, we introduce capital gain extrapolation into an otherwise standard New Keynesian model featuring a housing sector and a lower bound constraint on nominal interest rates.

The sticky price model has a number of attractive features. First, it shares the implications for housing price behavior and household beliefs with the simpler model considered before and thus quantitatively replicates the patterns of belief deviations and housing prices. Second, it is immune to the critique by Barsky et al. (2007) regarding the behavior of sticky price models featuring durable goods. In line with the data, the model implies that housing demand reacts more strongly to monetary disturbances than non-housing demand, despite the fact that housing prices are fully flexible. Third, the model introduces subjective housing beliefs in a way that monetary policy is unable to manipulate household beliefs to its own advantage. This allows for a meaningful discussion of Ramsey optimal monetary policy in the presence of subjective beliefs. Finally, the model makes a minimal departure from rational expectations: expectations about non-housing related variables are rational and all agents maximize given their (subjective) beliefs about the future.

To gain analytic insights, we derive a linear-quadratic approximation to the optimal policy problem and show how it is affected by the presence of subjective housing beliefs. We find that housing price gaps, i.e., deviations of housing prices from their efficient level, affect the economy via two channels. First, inefficiently high housing prices, driven by capital gain optimism, give rise to negative cost-push terms in the Phillips curve.[2] This feature allows the model to potentially generate a non-inflationary housing boom. Yet, a second channel is more important: rising housing price volatility increases the volatility of the natural rate of interest. Since increased housing price volatility is itself triggered by a fall in the average level of the natural rate, this dramatically exacerbates the lower-bound problem for a monetary policy authority confronted with falling natural rates.

The natural rate is affected by housing prices, because higher housing prices make it optimal to allocate more resources to housing investment. This exerts positive pressure on the output gap and counteracting these – so as to keep the output gap stable – requires policy to increase the real interest rate. Under rational expectations, housing prices never deviate from their efficient

[2]Conversely, inefficiently low housing prices, driven by capital gain pessimism, cause positive cost-push terms.

value, so that policy never has to work against inefficient investment pressures. With rational expectations, the volatility of the natural rate is thus independent of the average level of the natural rate.

These contrasting predictions of the model under rational and subjective housing beliefs also lead to rather different policy messages on how the optimal inflation target, i.e., the average inflation rate implied by optimal monetary policy, should respond to a fall in the natural rate of interest. Under rational expectations, the optimal inflation target is nearly invariant to the average level of the natural rate.

In the presence of capital gain extrapolation, the optimal inflation target increases considerably in response to a fall in the average natural rate. This is due to the increased volatility in the natural rate and cost-push shocks, which causes the lower bound on the nominal rate to become more restrictive. A more restrictive lower bound forces monetary policy to rely more strongly on promising future inflation in order to lower the real interest rate. This increases the average inflation rate under optimal policy. For our calibrated model, we find that the optimal inflation target should increase approximately by one third of a percent in response to a one percent fall in the natural rate with the increase becoming non-linear for very low levels of the natural rate.

We also investigate the optimal policy response to housing demand shocks. While inflation and the output gap do not respond to these shocks under rational expectations, capital gains induced by housing demand shocks get amplified by capital gain extrapolation and thereby generate persistent housing price gaps to which monetary policy optimally responds. Housing price gaps, however, generate opposing effects. On the one hand, inefficiently high housing prices generate negative cost-push pressures, which calls for a decrease in the policy rate; on the other hand, inefficiently high housing prices trigger a housing investment boom, which puts upward pressure on the output gap. Counteracting this second effect requires hiking policy rates.

In our calibrated model, the second effect quantitatively dominates. Optimal monetary policy thus 'leans against' housing price movements, but the optimal strength of the reaction depends on the direction of the shock. Following a positive housing preference shock, the increase in the interest rate (nominal and real) is more pronounced than the interest rate decrease following a negative housing demand shock. The presence of the lower-bound constraint thus attenuates the degree to which monetary policy leans against negative housing demand shocks.

We also consider whether macroprudential policies could address the housing market inefficiencies generated by capital gain extrapolation. We do so by considering housing taxes that might be levied on households in order to insulate monetary policymakers from the fluctuations in the housing price gap. We find that the required taxes would have to be large and very volatile. For our calibrated model, taxes must often exceed 20% of the rental value of housing per period and also often require equally sized or even larger housing subsidies. It appears somewhat unlikely that any of the existing macroprudential tools are capable of generating effects of such magnitude. And to the best of our knowledge, none of the available macroprudential tools allows subsidizing private sector behavior. Less aggressive tax policies turn out to be considerably less effective in bringing down the volatility of the housing price gap. This suggests that macroprudential policies are unable to substantially reduce the monetary policy trade-offs arising from subjective housing price expectations.

Related literature. This paper is related to work by Andrade et al. (2019, 2021) who study how the optimal inflation target depends on the natural rate of interest in a setting with a lower bound constraint. In line with our findings, they show that an increase in the inflation target is a promising approach to deal with the lower-bound problem. While their work considers optimized Taylor rules in a medium-scale sticky price model without a housing sector and rational expectations, the present paper studies Ramsey optimal policy in a model featuring a housing sector and subjective housing expectations.

A number of papers consider Ramsey optimal policy in the presence of a lower-bound constraint, but also abstract from housing markets and the presence of subjective beliefs (Eggertsson and Woodford (2003), Adam and Billi (2006a), Coibion et al. (2012)). This literature finds that lower bound episodes tend to be short and infrequent under optimal policy, so that average inflation is very close to zero under optimal policy. The present paper shows that this conclusion is substantially altered in the presence of subjective housing price expectations.

Optimal monetary policy with subjective beliefs has previously been analyzed in Caines and Winkler (2021) and Adam and Woodford (2021). These papers abstract from the lower-bound constraint and consider different belief setups that are not calibrated to replicate patterns of deviations from rational housing price expectations as observed in survey data.[3] We show that taking into account the existence of a lower-bound constraint on nominal rates is quantitatively important for understanding how the optimal inflation target responds to lower natural rates.

Outline. The rest of the paper is structured as follows. Section 2.2 documents how survey expectations about future housing prices deviate from rational expectations. Section 2.3 presents a simple housing model in which households extrapolate capital gains. It shows how this simple model can jointly replicate in equilibrium the behavior of housing prices and the pattern of deviations from rational expectations. Section 2.4 then presents the full housing model with sticky prices, subjective housing beliefs, and a lower-bound constraint on nominal rates. Section 2.5 derives a quadratic approximation to the monetary policy problem, which allows obtaining important analytic insights into the new economic forces arising from the presence of subjective housing price beliefs. We calibrate the model in Section 2.6 and present quantitative results about the optimal inflation target and the optimal policy response to housing shocks in Section 2.7. Section 2.8 discusses macroprudential policies and Section 2.9 concludes.

[3] Adam and Woodford (2021) consider 'worst-case' belief distortions, while Caines and Winkler (2021) consider a setting with 'conditionally model-consistent beliefs'. Both setups generate deviations from rational expectations for variables other than housing prices.

2.2 Cyclical Properties of Housing Price Expectations

This section documents that households' housing price expectations deviate in systematic ways from the full-information rational expectations (RE) benchmark. We consider three rationality tests that have recently been proposed in the literature (Coibion and Gorodnichenko (2015a), Adam et al. (2017) and Angeletos et al. (2021)). These tests cover different dimensions along which subjective expectations deviate from RE.

We measure housing prices using the S&P/Case-Shiller U.S. National Home Price Index and let q_t denote the quarterly average of the monthly housing price index. We consider both nominal and real housing prices with real housing prices being obtained by deflating nominal housing prices with the CPI.[45]

Expectations about housing capital gains are taken from the Michigan household survey. The survey provides subjective expectations about nominal four-quarter-ahead housing price growth, $E_t^{\mathcal{P}}[q_{t+4}/q_t]$, for the period 2007-2021. The survey also provides housing price growth expectations over the next five years. We focus on the shorter horizon because these expectations determine housing prices according to our model. The shorter horizon also allows performing a dynamic decomposition of forecast errors over time in response to realized capital gains.[6]

We consider both mean and median household expectations.[7] When considering real housing price expectations, we deflate the nominal mean (median) capital gain expectations with the mean (median) inflation expectation over the same period, as obtained from the Michigan survey.[8]

Sluggish Updating About the Expected Housing Price Level. We start by documenting that the mean/median household expectation about the future level of housing prices is updated too sluggishly. This can be tested following the approach of Coibion and Gorodnichenko (2015a), which involves considering regressions of the form

$$q_{t+4} - E_t^{\mathcal{P}}\left[q_{t+4}\right] = a^{CG} + b^{CG} \cdot \left(E_t^{\mathcal{P}}\left[q_{t+4}\right] - E_{t-1}^{\mathcal{P}}\left[q_{t+3}\right]\right) + \varepsilon_t. \tag{2.1}$$

[4]The simplified model in the next section makes predictions about real housing prices only, while the survey data contains information about nominal capital gain expectations. This leads us to consider nominal and real housing prices.

[5]We use the "Consumer Price Index for All Urban Consumers: All Items in U.S. City Average"obtained from FRED.

[6]Data limitations make such a decomposition difficult for five-year-ahead forecasts: with only 15 years of data, the dynamic decompositions become largely insignificant. Appendix B.1.1 shows, however, that all other patterns documented below are equally present in five-year-ahead expectations.

[7]Analyzing the dynamics of individual expectations over time is difficult because households in the Michigan survey are sampled at most twice. In general, cross-sectional disagreement between households might partly reflect heterogeneous information on the part of households, see Kohlhas and Walther (2021).

[8]As is well-known, these inflation expectations feature an upward bias relative to actual inflation outcomes. This, however, will not be the source of rejection of the RE hypothesis: all our tests focus on the cyclical properties of capital gain expectations and eliminate mean differences between forecasts and realizations using appropriate regression constants that will not be used in our rationality tests.

Table 2.1: Sluggish adjustment of housing price expectations

	Mean Expectations	Median Expectations
Nominal Housing Prices		
$\widehat{b}^{CG}$	2.22***	2.85***
	(0.507)	(0.513)
Real Housing Prices		
$\widehat{b}^{CG}$	2.00***	2.47***
	(0.332)	(0.366)

Notes: This figure shows the empirical estimates of regression (2.1) for nominal and real housings price and considers mean and median expectations. The reported standard errors are robust with respect to heteroskedasticity and serial correlation (Newey-West with four lags). Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

The regression projects forecast errors about the future housing price level on past forecast revisions. Under the RE hypothesis, information that is contained in agents' information set, i.e., past forecasts and their revisions, should not predict future forecast errors ($H_0 : b^{CG} = 0$).

We estimate equation (2.1) for nominal and real capital gains, using mean and median expectations, respectively. Expectations of the future house price level are computed as $E_t^{\mathcal{P}}[q_{t+4}] = E_t^{\mathcal{P}}[q_{t+4}/q_t]\, q_t$, where $E_t^{\mathcal{P}}[q_{t+4}/q_t]$ denotes the capital gain expectations from the Michigan survey and q_t the S&P/Case-Shiller Index.[9]

Table 2.1 reports the estimated b^{CG} from regression (2.1). We find that $\widehat{b}^{CG} > 0$, which is inconsistent with the RE hypothesis. The regression coefficient is positive and statistically significant at the 1% level in all considered specifications. This implies that the mean/median agent updates beliefs too sluggishly: future realizations move (on average) by more than what is suggested by past forecast revisions. The magnitude of the estimates is also large in economic terms: a coefficient estimate of two suggests that forecast revisions should approximately be three times as strong than they actually are.

Overall, sluggish belief updating is consistent with previous findings on the behavior of survey expectations about output, inflation and unemployment (Coibion and Gorodnichenko (2015a), Angeletos et al. (2021), Kohlhas and Walther (2021)). Furthermore, Bordalo et al. (2020) provide evidence of sluggish belief adjustment in consensus forecasts for other housing variables, such as residential investment and new housing starts.

Appendix B.1.2 shows that our findings are robust to using an instrumental-variable approach for estimating regression (2.1), in which forecast revisions are instrumented with monetary policy shocks obtained via high-frequency identification. Appendix B.1.3 shows that similar results emerge when using capital gains and expected capital gains in equation 2.1 instead of the level and expected level of the housing price.

[9] When considering real housing prices, nominal capital gain expectations from the Michigan survey are deflated using the subjective (mean or median) inflation expectations from the Michigan survey.

Table 2.2: Cyclicality of expected vs. actual capital gains

	$\hat{c}$ (in %)	$\hat{\mathbf{c}}$ (in %)	bias (in %) $-E(\hat{\mathbf{c}} - \hat{c})$	p-value $H_0 : c = \mathbf{c}$
Nominal Housing Prices				
Mean Expectations	0.033	-0.102	0.006	0.000
	(0.008)	(0.007)		
Median Expectations	0.014	-0.102	0.009	0.000
	(0.001)	(0.007)		
Real Housing Prices				
Mean Expectations	0.030	-0.113	-0.003	0.000
	(0.017)	(0.009)		
Median Expectations	0.010	-0.113	0.006	0.000
	(0.004)	(0.009)		

Notes: $\widehat{c}$ is the estimate of c in equation (2.2) and $\widehat{\mathbf{c}}$ the estimate of $\mathbf{c}$ in equation (2.3). The Stambaugh (1999) small sample bias correction is reported in the second-to-last column and the last column reports the p-values for the null hypothesis $c = \mathbf{c}$. Newey-West standard errors using four lags are in parentheses.

Opposing Cyclicality of Actual and Expected Capital Gains. Our second test documents the different cyclicality of actual and expected capital gains in housing markets. Differences between the cyclicality of actual and expected capital gains have previously been documented for stock markets, where actual and expected stock market capital gains covary differently with the price-to-dividend ratio (Greenwood and Shleifer (2014), Adam et al. (2017)). We consider here the cyclicality of expected and actual capital gains in the housing market with the price-to-rent ratio PR:

$$E_t^{\mathcal{P}}\left[\frac{q_{t+4}}{q_t}\right] = a + c \cdot PR_{t-1} + u_t \tag{2.2}$$

$$\frac{q_{t+4}}{q_t} = \mathbf{a} + \mathbf{c} \cdot PR_{t-1} + \mathbf{u}_t. \tag{2.3}$$

The rational expectations hypothesis implies $H_0 : c = \mathbf{c}$, whenever the agents' information set contains the past price-to-rent ratio as an observable.[10] Since the predictor variable used on the right-hand side of the preceding regressions equations is highly persistent, we correct for small sample bias in the coefficient estimates (Stambaugh (1999)).[11]

Table 2.2 reports the regression results. It shows that expected capital gains are positively associated with the PR-ratio, while realized capital gains are negatively associated. Expected capital gains are pro-cyclical, i.e., are high when market valuation is high, while realized capital

[10] In the regressions, we use the lagged PR-ratio, PR_{t-1}, instead of the current value, because the PR-ratio is computed using the average price over a quarter. In Adam et al. (2017) the price-to-dividend ratio was computed using the beginning of quarter stock price, which allowed using the current value in the regression.

[11] The small sample bias correction in Table 2.2 follows the same approach as the one in Table 1A in Adam et al. (2017).

gains are counter-cyclical, i.e., are low when market valuation is high. This pattern of is akin to the one documented in stock markets.

Quantitatively, the results imply that a two standard deviation increase of the PR-ratio by 15.5 units increases the mean household expectations about four-quarter-ahead real capital gains by around 0.5%. Actual four-quarter ahead capital gains, however, fall by around 1.5%, so that the forecast error is approximately 2%.

The last column in Table 2.2 performs a test of the rational expectations hypothesis that the cyclicality of actual and expected returns are equal ($H_0 : c = \mathbf{c}$). The test corrects for small sample bias, which is reported in the second to last column. We find that the difference in the cyclicality of actual and expected capital gains is highly statistically significant in all cases. Appendix B.1.4 shows that similar results are obtained when first subtracting equation (2.2) from (2.3) and estimating the resulting equation with forecast errors on the left-hand side, as in Kohlhas and Walther (2021) who do not consider housing related variables.

Initial under- and subsequent over-reaction of housing price expectations. While the results in Table 2.1 show that households adjust short-term housing price beliefs on average too sluggishly, the results in Table 2.2 indicate over-optimism (over-pessimism) in housing price expectations when the current market valuation is high (low), which points to some form of over-reaction to past housing price increases. It turns out that both patterns can be jointly understood by considering the dynamic response of actual and expected capital gains to housing price changes.

Following the approach in Angeletos et al. (2021), who analyze forecast errors about unemployment and inflation, we investigate how capital gains and forecast errors about these capital gains evolve over time in response to realized capital gains.[12] Provided households observe realized capital gains, the RE hypothesis implies that it should not be possible to predict future forecast errors with current capital gains.

We estimate the dynamic responses using local projections (Jorda (2005)) of the form

$$X_{t+h} = a^h + b^h \frac{q_{t-1}}{q_{t-2}} + u_t^h, \tag{2.4}$$

where the left-hand side variable X_{t+h} is either the cumulative capital gain (q_{t+h+4}/q_t), or the forecast error about the four-quarter-ahead capital gain ($q_{t+h+4}/q_{t+h} - E^{\mathcal{P}}_{t+h}[q_{t+h+4}/q_{t+h}]$), and u_t^h a serially correlated and heteroskedastic error term. Note that forecast errors are positive when households are overly pessimistic about capital gains and negative if households are overly optimistic.

Figure 2.1 reports the estimated coefficients b^h from local projection (2.4). Panel (a) depicts the response of cumulative capital gains. It shows that the initial capital gains is not only persistent, but increases further over time, reaching a plateau after around twelve quarters. Given the high

[12] These dynamic responses are well-defined in econometric terms, even if they cannot be given a structural interpretation, because past capital gains are likely driven by a combination of past shocks.

Figure 2.1: Dynamic responses to a realized capital gain

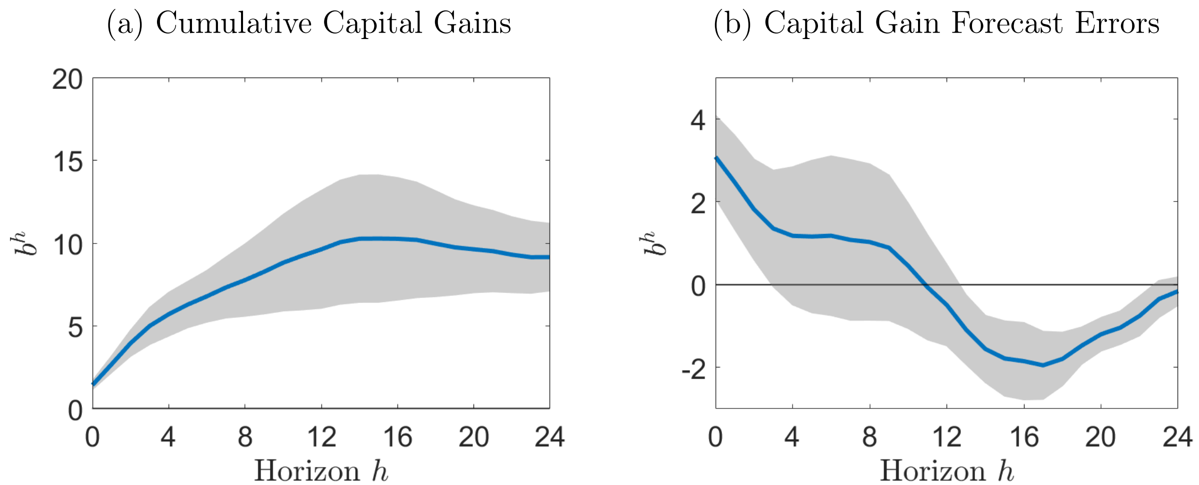

Notes: Panel (a) shows the dynamic response of cumulative real capital gains at horizon h to a one standard deviation innovation in the housing capital gain. Panel (b) reports the dynamic response of housing-price forecast errors at horizon h of one-year ahead expectations to a one standard deviation innovation in the housing capital gain. Positive (negative) values indicate that realized capital gains exceed (fall short of) expected capital gains. The shaded area shows the 90%-confidence intervals, standard errors are robust with respect to serial correlation and heteroskedasticity (Newey-West with $h+1$ lags).

serial correlation displayed by capital gains in housing markets, this feature is perhaps not too surprising.

Panel (b) depicts the dynamic response of forecast errors. Forecast errors are initially positive but later on – once cumulative capital gains reach their plateau – become negative before eventually disappearing. The positive values in initial periods indicates that agents' expectations react too sluggishly: realized capital gains are persistently larger than the expected gains. This also implies an underreaction in terms of the expected level of housing prices. Subsequently, when all actual capital gains have materialized and housing prices start to slightly mean revert, agents are too optimistic about future capital gains. This aligns well with the prior observation that capital gain expectations display the wrong cyclicality with housing market valuation.[13] It also implies that households entirely miss the mean-reversion in capital gains: forecast errors turn negative once housing prices level off and start to slightly mean-revert. This pattern is consistent with the experimental evidence provided in Armona et al. (2019).

In Appendix B.1.5, we show that the nominal forecast error responses look very similar. Likewise, using median expectations instead of mean expectations makes no noticeable difference of the results. In Appendix B.1.6, we show all our results obtained thus far are robust to excluding the Corona Virus period, i.e., to letting the sample period end in the last quarter of 2019.

[13] Since rents move only very slowly over time, changes in housing prices capture changes in the price-to-rent ratio rather well.

Analysis Using Regional Data. As is well known, housing prices often display considerable regional variation across the United States. We thus check whether the three deviations from RE documented above are also present in regional housing prices and housing price beliefs. Appendix B.1.7 uses regional housing price indices and exploits local information contained in the Michigan survey that allows grouping survey respondents into four different U.S. regions (North East, North Central/Midwest, South, and West). Repeating the above analyses at the regional level, it shows that one obtains quantitatively similar results. The next section presents a simple housing model that can quantitatively replicate the forecast error deviations documented in this section.

2.3 Simple Model with Capital Gain Extrapolation

This section presents a bare-bones housing model in which households (weakly) extrapolate past capital gains. The model makes equilibrium predictions for the joint dynamics of housing prices and housing price beliefs. Housing prices in the model depend on households' housing price beliefs, with the latter being influenced by past housing price behavior. We show that equilibrium dynamics of housing prices and housing price beliefs quantitatively replicate key features of housing price behavior in the U.S., as well as the deviations from rational expectations documented in the previous section. The simple model also predicts that low levels of the natural rate of interest give rise to increased housing price volatility. As we show, this prediction is consistent with the evolution of natural rates and housing prices in advanced economies over the past decades.

The full model in Section 2.4 additionally features nominal rigidities, a lower bound constraint on nominal rates, generalized preferences, and endogenous production of consumption goods and housing. The present section abstracts from these features, but nevertheless shares its implications for housing price behavior and housing price beliefs with the full model.

The Household Problem. There is a measure one of identical households.[14] Households are internally rational, as in Adam and Marcet (2011), i.e., maximize utility holding potentially subjective beliefs about variables beyond their control. The representative household chooses consumption C_t, housing units to own D_t, and housing units to rent D_t^R, to maximize

$$\max_{\left\{C_t\geq 0, D_t\in[0,D^{\max}], D_t^R\geq 0\right\}_{t=0}^{\infty}} \quad E_t^{\mathcal{P}} \sum_{t=0}^{\infty} \beta^t \left[C_t + \xi_t^d \left(D_t + D_t^R\right)\right]$$

$$\text{s.t.} \quad : \quad C_t + \left(D_t - (1-\delta)D_{t-1}\right) q_t + R_t D_t^R = Y_t \text{ for all } t \geq 0,$$

where Y_t is an exogenous (and sufficiently large) endowment, q_t the real price of housing, R_t the real rental price and $\delta > 0$ the housing depreciation rate. Rental units and housing units owned are perfect substitutes and $\xi_t^d \geq 0$ denotes a housing preference shock. The household's subjective probability measure $\mathcal{P}$ allows for subjective housing price beliefs. For simplicity, we assume beliefs

[14]The fact that households are identical is not common knowledge among households.

about other variables beyond the household's control, $\{Y_t, \xi_t^d, R_t\}_{t=1}^{\infty}$, to be rational. The latter assumption is not important for the results derived in this section.

Housing choices are subject to a short-selling constraint $D_t \geq 0$, which is standard, and to a long constraint $D_t \leq D^{\max}$. The latter insures existence of optimal plans in the presence of distorted housing beliefs. The long constraint is chosen such that it will never bind in equilibrium, i.e., $D^{\max} > D$, where D denotes the exogenously fixed housing supply. Without loss of generality, rental units are assumed to be in zero net supply.

The household first-order conditions imply that rents are given by

$$R_t = \xi_t^d \tag{2.5}$$

and that equilibrium housing prices satisfy[15]

$$q_t = \xi_t^d + \beta(1-\delta)E_t^{\mathcal{P}} q_{t+1}. \tag{2.6}$$

Capital Gain Extrapolation. We now introduce subjective price beliefs that give rise to capital gain extrapolation, using the setup in Adam et al. (2016). Importantly, the precise details generating capital gain extrapolation are not essential for the results in this section and we could have used alternative belief assumptions, e.g., learning from life-time experience as in Nagel and Xu (2022) and Malmendier and Nagel (2011, 2015), or could have directly assumed extrapolative behavior as in Barberis et al. (2015).

Households perceive capital gains to evolve according to

$$\frac{q_t}{q_{t-1}} = b_t + \varepsilon_t, \tag{2.7}$$

where $\varepsilon_t \sim iiN(0, \sigma_\varepsilon^2)$ is a transitory component of capital gains and b_t a persistent component, which itself evolves according to $b_t = b_{t-1} + \nu_t$, with $\nu_t \sim iiN(0, \sigma_\nu^2)$.[16] Households observe the realized capital gains (q_t/q_{t-1}) and use Bayesian belief updating to decompose observed capital gains into their persistent and transitory components. With conjugate prior beliefs, the subjective conditional one-step-ahead capital gain expectations

$$\beta_t \equiv E_t^{\mathcal{P}} \left(q_{t+1}/q_t\right) \tag{2.8}$$

[15]This holds true in equilibrium because $0 < D < D^{\max}$. For the household, however, first-order conditions may hold only with inequality under the subjectively optimal plans, due to the presence of short and long constraints. The latter explains why rational households can hold price expectations that differ from the discounted sum of future rents, see Adam and Marcet (2011) for details and Adam and Nagel (2022) for related arguments.

[16]In the full model in Section 2.4, we will assume the same beliefs for risk-adjusted house price growth. With risk neutrality, the two coincide.

Figure 2.2: Capital gain surprises and revisions

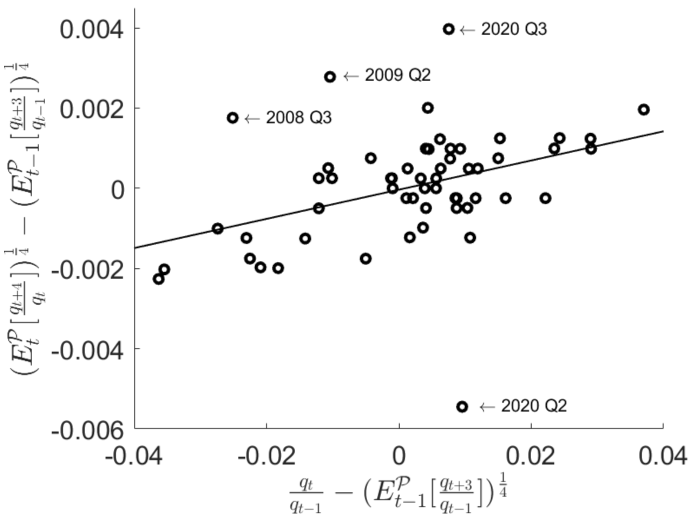

Notes: This figure plots the capital gain surprises against capital gain revisions in the Michigan survey (2007-2021), along with a linear regression line.

evolve according to

$$\beta_t = \min\left\{\beta_{t-1} + \frac{1}{\alpha}\left(\frac{q_{t-1}}{q_{t-2}} - \beta_{t-1}\right), \beta^U\right\}, \tag{2.9}$$

where $1/\alpha$ is the Kalman gain determining how strongly households' capital gain expectations respond to past capital gain surprises.[17] The Kalman gain thus captures the degree to which past capital gain surprises are extrapolated into the future. The upper bound β^U on the beliefs in equation (2.9) is there to insure that capital gain optimism is bounded from above, so as to keep subjectively expected utility finite.[18]

Figure 2.2 illustrates the relationship between belief revisions and forecast errors implied by equation (2.9) using the Michigan survey data. The figure plots on the vertical axis a measure of the quarterly revision in capital gain expectations, $\left(E_t^{\mathcal{P}}\left(q_{t+4}/q_t\right)\right)^{\frac{1}{4}} - \left(E_{t-1}^{\mathcal{P}}\left(q_{t+3}/q_{t-1}\right)\right)^{\frac{1}{4}}$, and on the horizontal axis a measure of the forecast error in quarterly capital gains, $\frac{q_t}{q_{t-1}} - \left(E_{t-1}^{\mathcal{P}}\left(q_{t+3}/q_{t-1}\right)\right)^{\frac{1}{4}}$, for all quarters in the Michigan survey. Consistent with equation (2.9), there is a clear positive and approximately linear relationship between capital gain surprises and belief revisions in Figure 2.2. The most notable deviations from the regression line are the ones around the Great Recession (2008Q3 and 2009Q2) and the Covid Recession (2020Q2 and 2020Q3).

Equilibrium Dynamics of Housing Prices and Capital Gain Expectations. From equation (2.6) and the definition of subjective beliefs β_t it follows that the equilibrium housing

[17]The (steady-state) Kalman gain depends on the subjectively perceived values for $(\sigma_\varepsilon^2, \sigma_\nu^2)$.

[18]The bound can be interpreted as a short-cut for a truncated prior support or b_t. The bounding function in (2.9) is a special case of the bounding function used in Adam et al. (2016), obtained by setting $\beta^L = \beta^U$.

price is given by

$$q_t = \frac{1}{1-\beta(1-\delta)\beta_t}\xi_t^d, \tag{2.10}$$

where β_t evolves according to (2.9). Equations (2.9) and (2.10) thus jointly characterize the equilibrium dynamics of housing prices and subjective beliefs. From equations (2.5) and (2.10) follows that the equilibrium price-to-rent ratio is given by

$$PR_t \equiv \frac{q_t}{R_t} = \frac{1}{1-\beta(1-\delta)\beta_t}. \tag{2.11}$$

Calibration. The simple model just described can generate empirically plausible housing price behavior and the resulting housing price beliefs quantitatively match the deviations from rational expectations presented in the previous section. The calibration in this section is identical to the one for the full model, with the exception for the standard deviation of the innovations to housing preferences.[19] We consider housing preference shocks evolving according to

$$\log \xi_t^d = (1-\rho_\xi)\log \underline{\xi}^d + \rho_\xi \log \xi_{t-1}^d + \varepsilon_t^d, \tag{2.12}$$

where $\varepsilon_t^d \sim iiN$ satisfies $E[e^{\varepsilon_t^d}] = 1$. Following Adam and Woodford (2021), we set $\rho_\xi = 0.99$ and $\delta = 0.03/4$. The standard deviation of ε_t^d is set to 0.0067, so that the model replicates the empirical standard deviation of the price-to-rent ratio, expressed in percent deviation from its mean, over the period for which we have survey data on housing expectations (2007-2021). The average value of the housing preference $\underline{\xi}^d > 0$ is irrelevant, as we are only concerned with moments characterizing cyclical properties (deviations from average values).

For the subjective belief process, we completely tie our hands and set $1/\alpha = 0.007$, which is the value estimated in Adam et al. (2016) using stock market expectations. The low value for the Kalman gain implies that agents extrapolate only weakly, as they believe most of the realized capital gains being due to transitory components. The value for the upper belief bound β^U is set as in the full model, where it matches the maximum observed deviation of the price-to-rent ratio from its mean. Finally, the quarterly discount factor β is set such that the real interest rate is equal to 0.75%, which is the average value of the estimated U.S. natural rate over the period 2007-2021, according to estimates using the approach of Holston et al. (2017).

Housing Price Behavior. Table 2.3 illustrates that the subjective belief model replicates surprisingly well the empirical behavior of the price-to-rent ratio and of capital gains. While the standard deviation of the price-to-rent ratio is a targeted moment, all other moments are untargeted. The model matches very well the high quarterly autocorrelation of the price-to-rent ratio and the fairly high quarterly autocorrelation of capital gains. It undershoots somewhat

[19] This is so because the present section matches moments for a different time period than the full model, i.e., the period for which we have subjective expectations data (2007-2021).

Table 2.3: Housing price moments: data versus model

	Data	Subjective Belief Model	RE Model
$std\,(PR_t)$	8.6	8.6	2.69
$corr(PR_t, PR_{t-1})$	0.99	0.99	0.99
$std(q_t/q_{t-1})$	0.06	0.04	0.003
$corr(q_t/q_{t-1}, q_{t-1}/q_{t-2})$	0.97	0.94	-0.01

Notes: The table reports the standard deviation and first-order autocorrelation of price-to-rent ratios and capital gains in the data, for the model under subjective housing beliefs and the model under rational expectations.

Table 2.4: Patterns of deviations from rational expectations: data versus model

	Subjective Belief Model	Data	
		Mean Expectations	Median Expectations
b^{CG} from (2.1)	2.09	1.68	2.12
		(0.355)	(0.394)
c (in %) from (2.2)	0.03	0.030	0.010
		(0.172)	(0.043)
$\mathbf{c}$ (in %) from (2.3)	-0.063	-0.113	-0.113
		(0.009)	(0.009)

Notes: This table shows the model-implied regression coefficients of regressions (2.1), (2.2) and (2.3) for a natural rate of 0.75% (annualized) in the first column and the empirical results (for real housing prices) in the second and third column.

the standard deviation of quarterly capital gains, illustrating that it features perhaps too little high-frequency variation in prices.[20]

Table 2.3 also reports the rational expectations (RE) outcome using the same calibration as for the subjective belief model. It shows that about 70% of the fluctuations in the price-dividend ratio in the subjective belief model is due to capital gain extrapolation. Adam et al. (2016) explain how capital gain extrapolation generates momentum and mean reversion in prices and thus contributes to asset price volatility.

While the ability of capital gain extrapolation to increase the price volatility is well-known, we now turn to the new question of whether the model with capital gain extrapolation matches the structure of forecast errors documented in Section 2.2.

Belief revisions and forecast errors. The simple model quantitatively matches the three deviations from rational housing price expectations documented in Section 2.2.

Table 2.4 reports the outcomes of population regressions of equations (2.1), (2.2) and (2.3) for the calibrated subjective belief model. The results shows the model matches sluggish updating about expected housing prices ($b^{CG} > 0$) and the opposing cyclicality of actual and expected capital

[20]This could easily be remedied by adding some iid shocks, say iid shocks to the discount factor β.

Figure 2.3: Dynamic forecast error response: data versus model

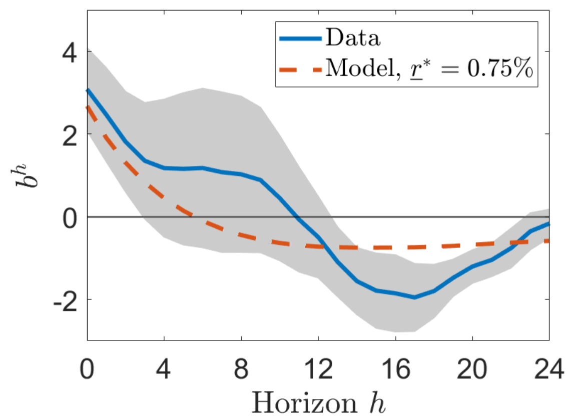

Notes: The figure shows impulse-response functions of housing-price forecast errors of one-year ahead expectations to a one standard deviation innovation in the housing capital gain from the model and the data. The shaded area shows the 90%-confidence intervals of the empirical estimates, standard errors are robust with respect to serial correlation and heteroskedasticity (Newey-West with $h+1$ lags).

gains ($c > 0$ and $\mathbf{c} < 0$). For better comparison, Table 2.4 also reports also the empirical estimates of the corresponding coefficients from Tables 2.1 and 2.2. The magnitude of the coefficients generated by the model closely match the ones obtained using survey data, with the exception that the model underpredicts the counter-cyclicality of actual capital gains.

Figure 2.3 shows that the simple model is able to match the dynamic response of forecast errors documented empirically in Figure 2.1(b). We compute model-implied forecast errors as $FE^{model}_{t+h} = \frac{q_{t+4+h}}{q_{t+h}} - (\beta_{t+h})^4$ and compute the population local projections (2.4). Consistent with the data, the model generates initial underprediction of capital gains (over-pessimism) and subsequently overprediction (over-optimism).

Appendix B.2.1 reports the dynamic forecast error responses for the model and in the data about the expected housing price level (rather than the expected capital gain). It shows that the model matches equally well the patterns of forecasts errors about the future housing price level.

Falling Natural Rates and Rising Housing Price Volatility. The simple housing model also predicts that falling natural rates of interest will give rise to higher volatility for housing prices. Such a relationship between natural rates and housing volatility is present in the data. It can be seen by considering regressions of the form

$$Std(PR_{t-\frac{h}{2}}, ..., PR_{t+\frac{h}{2}}) = a^*_h - b^*_h \cdot r^*_t + u_{t,h}, \tag{2.13}$$

where r^*_t denotes the natural rate of interest from Holston et al. (2017) and $Std(PR_{t-\frac{h}{2}}, ..., PR_{t+\frac{h}{2}})$ the standard deviation of the price rent ratio using a window of $h+1$ quarters centered around period t. Under the standard assumption that the natural rate of interest is only a function of

Figure 2.4: The natural rate and housing price fluctuations

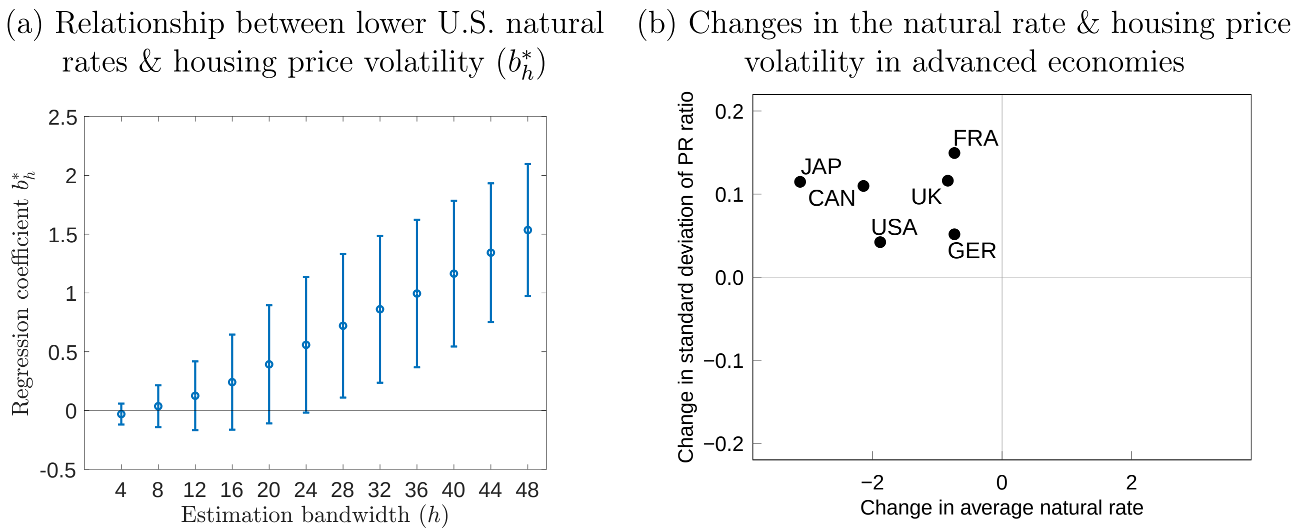

Notes: Panel (a) reports the regression coefficient b_h^* from equation (2.13) together with 68% Newey-West error bands using h lags. Panel (b) plots the pre-/post-1990 changes in the average natural rate against the changes in the volatility of the price-to-rent ratio for different advanced economies. The volatilities of the price-to-rent ratios in the pre-/post-1990 periods are the standard deviations relative to the period-specific mean values.

exogenous fundamentals, the regression coefficients b_h^* can be interpreted as capturing a causal relationship.

Panel (a) in Figure 2.4 reports the coefficients b_h^* for the United States using various estimation bandwidths h. While narrow bandwidths generate insignificant outcomes, most likely due to the difficulty associated with reliably estimating the standard deviation of the PR-ratio, the coefficient becomes positive and significant for larger bandwidth and is quite large when using a bandwidth of 48 quarters (+/- 2 years). We can thus conclude that the standard deviation of U.S. housing prices is rising as the natural rate falls.

Panel (b) in Figure 2.4 shows that this relationship is also present in other countries: it plots the change in the average level of the natural rate from the period before 1990 to the period after 1990 for the U.S., Canada, France, Germany, and the United Kingdom, against the change in the standard deviation of the price-to-rent ratio over the same periods. To take possible shifts in the mean of the PR-ratio over time into account, e.g., due to falling real interest rates, the standard deviation of the PR-ratio is computed in each of the two sub-periods for the *percent* deviation of the PR-ratio from its period-specific mean.[21] In all six advanced economies, the PR-ratio has become more volatile as the average level of the natural rate has declined.

Equations (2.9) and (2.10) reveal how housing prices in our simple model are affected by the level of the natural rate of interest. The natural rate of interest is given by $\underline{r}^* = 1/\beta - 1$ and only depends on the discount factor $\beta \in (0, 1)$. A discount factor closer to one thus lowers the natural

[21]The empirical results become even stronger if one considers instead the absolute standard deviation of the PR-ratio.

Figure 2.5: Model-implied housing price response for different natural rates

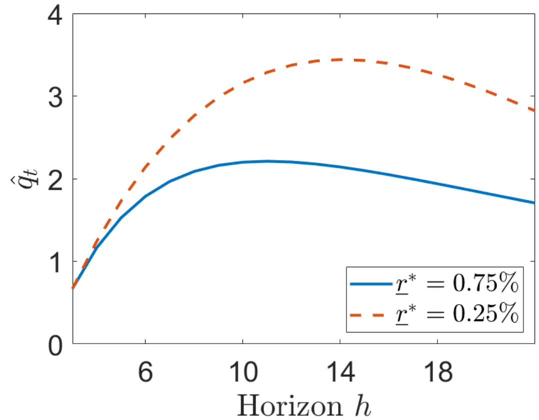

rate of interest.

Figure 2.5 illustrates that the model in fact generates a negative relationship between the level of the natural rate and housing price volatility.[22] It presents the impulse response of real housing prices to a positive housing preference shock ξ_t^d, which is the only shock driving housing prices in the model. It considers this response for the calibrated level of the natural rate of 0.75% and for a lower natural rate level equal to 0.25%.[23] The key message of Figure 2.5 is that housing prices respond stronger to housing demand shocks when natural rates are lower: the same shock gives rise to an approximately 75% stronger housing price response when the natural rate is at its lower level.

This surprising model outcome can be explained as follows. The capital gain increase triggered by the fundamental shock in the initial period leads to an upward revision of capital gain expectations. Equation (2.10) implies, however, that these higher capital gain expectations produce larger realized capital gains, the higher is the value for β, i.e., the lower is the natural rate of interest. Higher realized capital gains produce stronger upward revisions in beliefs in the future and thus feed stronger capital gains in the subsequent period. Through this feedback loop, low natural rates generate more momentum in housing price changes following fundamental shocks, allowing the model to replicate the relationship between natural rates and the volatility of housing prices.

2.4 Full Model with Capital Gain Extrapolation

This section studies the monetary policy implications of falling natural rates of interest and rising housing price volatility. To this end, we embed capital gain extrapolation into a sticky price model with a housing sector. The model features endogenous production of consumption goods and

[22] The full model presented later on will also be able to quantiatively replicate this relationship, see section 2.6.

[23] To account for the higher housing price levels associated with lower natural rates, we show impulse responses in terms of percent deviations from their respective steady state values. The model-implied response for the PR-ratio to a housing preference shock looks very similar and is shown in Appendix B.2.2.

housing and generalizes the setup in Adam and Woodford (2021) by allowing for belief distortions that are not absolutely continuous with respect to the beliefs held by the policymaker. This permits analyzing the subjective housing beliefs as in equation (2.7), which give rise to capital gain extrapolation and deviations from rational expectations matching patterns in the survey data. In addition, we consider a lower-bound constraint on nominal rates, which we show to be quantitatively important for understanding how the optimal inflation target responds to lower natural rates in the presence of subjective housing beliefs.

We consider an economy populated by internally rational decision makers (Adam and Marcet (2011)): households maximize utility and firms maximize profits, but both do so using a potentially subjective probability measure $\mathcal{P}$, which assigns probabilities to all external variables, i.e., to all variables that are beyond agents' control. These variables include fundamental shocks, as well as competitive market prices (wages, goods prices, housing prices and rents). The setup delivers rational expectations in the special case where $\mathcal{P}$ is the objective probability measure.

The economy is made up of identical infinitely-lived households, each of which maximizes the following objective function[24]

$$U \equiv E_0^{\mathcal{P}} \sum_{t=0}^{\infty} \beta^t \left[\tilde{u}(C_t;\xi_t) - \int_0^1 \tilde{v}(H_t(j);\xi_t)dj + \tilde{\omega}(D_t + D_t^R;\xi_t) \right], \tag{2.14}$$

subject to the sequence of flow budget constraints

$$\begin{aligned} & C_t + B_t + (D_t - (1-\delta)D_{t-1}) \frac{q_t^u}{\tilde{u}_C(C_t;\xi_t)} + k_t + R_t D_t^R = \\ & \tilde{d}(k_t;\xi_t) \frac{q_t^u}{\tilde{u}_C(C_t;\xi_t)} + \int_0^1 w_t(j) H_t(j) dj + \frac{B_{t-1}}{\Pi_t}(1+i_{t-1}) + \frac{\Sigma_t}{P_t} + \frac{T_t}{P_t}, \end{aligned} \tag{2.15}$$

where C_t is an aggregate consumption good, $H_t(j)$ is the quantity supplied of labor of type j and $w_t(j)$ the associated real wage, D_t the stock of owned houses, D_t^R the units of rented houses, $\delta \in [0,1]$ the housing depreciation rate, and q_t^u the real price of houses in marginal utility units, defined as

$$q_t^u \equiv q_t \tilde{u}_C(C_t;\xi_t),$$

where q_t is the real house price in units of consumption.[25] The variable q_t^u provides a measure of whether housing is currently expensive or inexpensive, in units that are particularly relevant for determining housing demand. The variable k_t denotes investment in new houses and $\tilde{d}(k_t;\xi_t)$ the resulting production of new houses.[26] $B_t \equiv \widetilde{B}_t/P_t$ denotes the real value of nominal government

[24]It cannot be common knowledge to households that they are representative whenever $\mathcal{P}$ deviates from the rational measure.

[25]In Section 2.3, q_t^u and q_t coincide due to risk-neutrality.

[26]We consolidate housing production into the household budget constraint. It would be equivalent to have instead a separate housing production sector that is owned by households.

bond holdings $\widetilde{B}_t$ and P_t the nominal price of consumption. $\Pi_t = P_t/P_{t-1}$ is the inflation rate, i_t the nominal interest rate, R_t the real rental rate for housing units, and ξ_t is a vector of exogenous disturbances, which may induce random shifts in the functions $\tilde{u}$, $\tilde{v}$, $\tilde{\omega}$ and $\tilde{d}$. T_t denotes nominal lump sum transfers (taxes if negative) from the government and Σ_t nominal profits accruing to households from the ownership of firms.

Households discount future payoffs at the rate $\beta \in (0,1)$. Since our model is formulated in terms of growth-detrended variables, the discount rate β jointly captures the time preference rate $\widetilde{\beta} \in (0,1)$ and the steady-state growth rate of marginal utility. Letting $g_c \geq 0$ denote the steady-state growth rate of consumption in non-detrended terms, we have

$$\beta \equiv \widetilde{\beta}\frac{\tilde{u}_C(C(1+g_c);\underline{\xi})}{\tilde{u}_C(C;\underline{\xi})}, \tag{2.16}$$

where $\underline{\xi}$ denotes the steady state value of the disturbance ξ_t. When the growth rate g_c of the economy falls, the discount rate β increases because marginal utility falls less strongly. We can thus capture a fall in the trend growth rate of the economy simply via an increase in the time discount rate β. Declining trend growth causes the steady-state real interest rate and thus the average natural rate of interest to fall, which is in line with the estimates provided in Holston et al. (2017) (see Appendix B.7).

The aggregate consumption good is a Dixit-Stiglitz aggregate of each of a continuum of differentiated goods,

$$C_t \equiv \left[\int_0^1 c_t(i)^{\frac{\eta-1}{\eta}} di\right]^{\frac{\eta}{\eta-1}}, \tag{2.17}$$

with an elasticity of substitution $\eta > 1$. We further assume isoelastic functional forms

$$\begin{aligned}
\tilde{u}(C_t;\xi_t) &\equiv \frac{C_t^{1-\tilde{\sigma}^{-1}}\bar{C}_t^{\tilde{\sigma}^{-1}}}{1-\tilde{\sigma}^{-1}}, \\
\tilde{v}(H_t(j);\xi_t) &\equiv \frac{\lambda}{1+\nu}\left(H_t(j)\right)^{1+\nu}\bar{H}_t^{-\nu}, \\
\tilde{\omega}(D_t + D_t^R;\xi_t) &\equiv \xi_t^d\left(D_t + D_t^R\right), \\
\tilde{d}(k_t;\xi_t) &\equiv \frac{A_t^d}{\tilde{\alpha}}k_t^{\tilde{\alpha}},
\end{aligned} \tag{2.18}$$

where $\tilde{\sigma}, \nu > 0$, $\tilde{\alpha} \in (0,1)$ and $\{\bar{C}_t, \bar{H}_t, \xi_t^d, A_t^d\}$ are bounded, exogenous and positive disturbance processes which are among the exogenous disturbances included in the vector ξ_t.

Our specification includes two housing-related disturbances, namely ξ_t^d, which captures shocks to housing preferences, and A_t^d, which captures shocks to the productivity in the construction of new houses. We impose linearity in the utility function (2.18), because it greatly facilitates the characterization of optimal policy, with rented and owned housing units being perfect substitutes. Introducing a weight on rental units relative to housing units would allow us to perfectly match the average price-to-rent ratio we observe in the data. However, since this does not change any other results, we abstract from such a scaling parameter and assign equal weight to housing and renting in the utility.

Each differentiated good is supplied by a single monopolistically competitive producer; there is a common technology for the production of all goods, in which (industry-specific) labor is the only variable input,

$$y_t(i) = A_t f(h_t(i)) = A_t h_t(i)^{1/\phi}, \tag{2.19}$$

where A_t is an exogenously varying technology factor, and $\phi > 1$. The Dixit-Stiglitz preferences (2.17) imply that the quantity demanded of each individual good i will equal[27]

$$y_t(i) = Y_t \left(\frac{p_t(i)}{P_t} \right)^{-\eta}, \tag{2.20}$$

where Y_t is the total demand for the composite good defined in (2.17), $p_t(i)$ is the price of the individual good, and P_t is the price index,

$$P_t \equiv \left[\int_0^1 p_t(i)^{1-\eta} di \right]^{\frac{1}{1-\eta}}, \tag{2.21}$$

corresponding to the minimum cost for which a unit of the composite good can be purchased in period t. Total demand is given by

$$Y_t = C_t + k_t + g_t Y_t, \tag{2.22}$$

where g_t is the share of the total amount of composite goods purchased by the government, treated here as an exogenous disturbance process.

2.4.1 Household Optimality Conditions

Internally rational households choose state-contingent sequences for the choice variables $\left\{C_t, H_t(j), D_t, D_t^R, k_t, B_t\right\}$ so as to maximize (2.14), subject to the budget constraints (2.15), taking as given their beliefs about the processes $\{P_t, w_t(j), q_t^u, R_t, i_t, \Sigma_t/P_t, T_t/P_t\}$, as determined by the (subjective) measure $\mathcal{P}$.

We shall be particularly interested in the policy implications generated by subjective housing price beliefs. To insure that an optimum exists in the presence of potentially distorted beliefs about the housing price q_t^u, we require housing choices to lie in some compact choice set $D_t \in [0, D^{\max}]$, as discussed in Section 2.3, where the upper bound can be arbitrarily large.

[27] In addition to assuming that household utility depends only on the quantity obtained of C_t, we assume that the government also cares only about the quantity obtained of the composite good defined by (2.17), and that it seeks to obtain this good through a minimum-cost combination of purchases of individual goods.

The first order conditions give rise to an optimal labor supply relation

$$w_t(j) = \frac{\tilde{v}_H(H_t(j);\xi_t)}{\tilde{u}_C(C_t;\xi_t)}, \tag{2.23}$$

a consumption Euler equation

$$\tilde{u}_C(C_t;\xi_t) = \beta E_t^{\mathcal{P}}\left[\tilde{u}_C(C_{t+1};\xi_{t+1})\frac{1+i_t}{P_{t+1}/P_t}\right], \tag{2.24}$$

an equation characterizing optimal investment in new houses

$$k_t = \left(A_t^d q_t^u \frac{C_t^{\tilde{\sigma}^{-1}}}{\bar{C}_t^{\tilde{\sigma}^{-1}}}\right)^{\frac{1}{1-\tilde{\alpha}}}, \tag{2.25}$$

an optimality condition for rental units

$$\xi_t^d = R_t \tilde{u}_C(C_t, \xi_t), \tag{2.26}$$

and a set of conditions determining the optimal housing demand D_t:

$$\begin{array}{ll} q_t^u < \xi_t^d + \beta(1-\delta)E_t^{\mathcal{P}} q_{t+1}^u & \text{if } D_t = D^{\max} \\ q_t^u = \xi_t^d + \beta(1-\delta)E_t^{\mathcal{P}} q_{t+1}^u & \text{if } D_t \in (0, D^{\max}) \\ q_t^u > \xi_t^d + \beta(1-\delta)E_t^{\mathcal{P}} q_{t+1}^u & \text{if } D_t = 0. \end{array} \tag{2.27}$$

With rational expectations, the upper and lower holding bounds never bind.[28] Since we are interested in how the presence of belief distortions about future housing values affect equilibrium outcomes, the bounds in equation (2.27) can potentially bind under the *subjectively* optimal plans. This explains why an internally rational household can hold subjective housing price expectations, even if she holds rational expectations about the preference shocks ξ_t^d in equation (2.27).

Forward-iterating on equation (2.24), which holds with equality under all belief-specifications, delivers a present-value formulation of the consumption Euler equation

$$\tilde{u}_C(C_t;\xi_t) = \lim_{T\to\infty} E_t^{\mathcal{P}}\left[\tilde{u}_C(C_T;\xi_T)\beta^T \prod_{k=0}^{T-t} \frac{1+i_{t+k}}{P_{t+k+1}/P_{t+k}}\right], \tag{2.28}$$

which will be convenient to work with, especially under subjective belief specifications. Household choices must also satisfy the transversality constraint

$$\lim_{T\to\infty} \beta^T E_t^{\mathcal{P}}\left[\tilde{u}_C(C_T;\xi_T)B_T + D_T q_T^u\right] = 0. \tag{2.29}$$

[28] The upper bound $D^{\max}$ has been chosen sufficiently large for this to be true. The lower bound is never reached because the housing production function satisfies Inada conditions.

Optimal household behavior under potentially distorted beliefs is jointly characterized by equations (2.23) and (2.25)-(2.29).

2.4.2 Optimal Price Setting by Firms

The producers in each industry fix the prices of their goods in monetary units for a random interval of time, as in the model of staggered pricing introduced by Calvo (1983) and Yun (1996). Producers use the representative households' subjectively optimal consumption plans to discount profits and are assumed to know the product demand function (2.20). They need to formulate beliefs about the future price levels P_T, industry-specific wages $w_T(j)$, aggregate demand Y_T, and productivity A_T.

Let $0 \leq \alpha < 1$ be the fraction of prices that remain unchanged in any period. A supplier i in industry j that changes its price in period t chooses its new price $p_t(i)$ to maximize

$$E_t^{\mathcal{P}} \sum_{T=t}^{\infty} \alpha^{T-t} Q_{t,T} \Pi\left(p_t(i), P_T, w_T(j), Y_T, A_T\right), \tag{2.30}$$

where $E_t^{\mathcal{P}}$ denotes the expectations of price setters conditional on time t information, which are identical to the expectations held by consumers. Firms discount random nominal income in period T using households' subjective stochastic discount factor $Q_{t,T}$, which is given by

$$Q_{t,T} = \beta^{T-t} \frac{\tilde{u}_C\left(C_T, \xi_T\right)}{\tilde{u}_C\left(C_t, \xi_t\right)} \frac{P_t}{P_T}.$$

The term α^{T-t} in equation (2.30) captures the probability that a price chosen in period t will not have been revised by period T, and the function $\Pi\left(p_t(i), ...\right)$ indicates the nominal profits of the firm in period t, as discussed next.

Profits are equal to after-tax sales revenues net of the wage bill. Sales revenues are determined by the demand function (2.20), so that (nominal) after-tax revenue equals

$$\left(1 - \tau_t\right) p_t(i) Y_t \left(\frac{p_t(i)}{P_t}\right)^{-\eta}.$$

Here τ_t is a proportional tax on sales revenues in period t, $\{\tau_t\}$ is treated as an exogenous disturbance process, taken as given by the monetary policymaker. We assume that τ_t fluctuates over a small interval around a non-zero steady state level $\underline{\tau}$. We allow for exogenous variations in the tax rate in order to include the possibility of "pure cost-push shocks" that affect the equilibrium pricing behavior while implying no change in the efficient allocation of resources.

The labor demand of firm i at a given industry-specific wage $w_t(j)$ can be written as

$$h_t(i) = \left(\frac{Y_t}{A_t}\right)^{\phi} p_t(i)^{-\eta\phi} P_t^{\eta\phi}, \tag{2.31}$$

which follows from (2.19) and (2.20). Using this, the nominal wage bill is given by

$$P_t w_t(j) h_t(i) = P_t w_t(j) \left(\frac{Y_t}{A_t}\right)^{\phi} p_t(i)^{-\eta\phi} P_t^{\eta\phi}.$$

Subtracting the nominal wage bill from the above expression for nominal after tax revenue, we obtain the function $\Pi\left(p_t(i), P_T, w_T(j), Y_T, A_T\right)$ used in (2.30).

Each of the suppliers that revise their prices in period t chooses the same new price p_t^*, that maximizes (2.30). The first-order condition with respect to $p_t(i)$ is given by[29]

$$E_t^{\mathcal{P}} \sum_{T=t}^{\infty} \alpha^{T-t} Q_{t,T} \Pi_1\left(p_t(i), P_T, w_T(j), Y_T, A_T\right) = 0.$$

The equilibrium choice p_t^*, which is the same for each firm i in industry j, is the solution to this equation. Letting p_t^j denote the price charged by firms in industry j at time t, we have $p_t^j = p_t^*$ in periods in which industry j resets its prices and $p_t^j = p_{t-1}^j$ otherwise.

Under the assumed isoelastic functional forms, the optimal choice has a closed-form solution

$$\left(\frac{p_t^*}{P_t}\right)^{1+\eta(\phi-1)} = \frac{E_t^{\mathcal{P}} \sum_{T=t}^{\infty} \alpha^{T-t} Q_{t,T} \frac{\eta}{\eta-1} \phi w_T(j) \left(\frac{Y_T}{A_T}\right)^{\phi} \left(\frac{P_T}{P_t}\right)^{\eta\phi+1}}{E_t^{\mathcal{P}} \sum_{T=t}^{\infty} \alpha^{T-t} Q_{t,T} (1-\tau_T) Y_T \left(\frac{P_T}{P_t}\right)^{\eta}}. \tag{2.32}$$

The price index evolves according to a law of motion

$$P_t = \left[(1-\alpha) p_t^{*1-\eta} + \alpha P_{t-1}^{1-\eta}\right]^{\frac{1}{1-\eta}}, \tag{2.33}$$

as a consequence of (2.21). The equilibrium inflation in any period is characterized by

$$\left(\frac{P_t}{P_{t-1}}\right)^{\eta-1} = \frac{1-(1-\alpha)\left(\frac{p_t^*}{P_t}\right)^{1-\eta}}{\alpha}. \tag{2.34}$$

The welfare loss from price adjustment frictions can be captured by price dispersion, which is

[29]Note that supplier i's profits in (2.30) are a concave function of the quantity sold $y_t(i)$, since revenues are proportional to $y_t(i)^{\frac{\eta-1}{\eta}}$ and hence concave in $y_t(i)$, while costs are convex in $y_t(i)$. Moreover, since $y_t(i)$ is proportional to $p_t(i)^{-\eta}$, the profit function is also concave in $p_t(i)^{-\eta}$. The first-order condition for the optimal choice of the price $p_t(i)$ is the same as the one with respect to $p_t(i)^{-\eta}$; hence the first-order condition with respect to $p_t(i)$ is both necessary and sufficient for an optimum.

defined as

$$\Delta_t \equiv \int_0^1 \left(\frac{p_t^j}{P_t}\right)^{-\eta(1+\omega)} dj \geq 1, \tag{2.35}$$

where

$$\omega \equiv \phi(1+\nu) - 1 > 0$$

is the elasticity of real marginal cost in an industry with respect to industry output.

Using equation (2.33) together with the fact that the relative prices of the industries that do not change their prices in period t remain the same, one can derive a law of motion for the price dispersion term Δ_t of the form

$$\Delta_t = h(\Delta_{t-1}, P_t/P_{t-1}), \tag{2.36}$$

with

$$h(\Delta_t, P_t/P_{t-1}) \equiv \alpha\Delta_t \left(\frac{P_t}{P_{t-1}}\right)^{\eta(1+\omega)} + (1-\alpha)\left(\frac{1-\alpha\left(\frac{P_t}{P_{t-1}}\right)^{\eta-1}}{1-\alpha}\right)^{\frac{\eta(1+\omega)}{\eta-1}}.$$

As is commonly done, we assume that the initial degree of price dispersion is small ($\Delta_{-1} \sim O(2)$).

Equations (2.32), (2.34), and (2.36) jointly define a short-run aggregate supply relation between inflation, output and house prices (via the aggregate demand equation (2.22) and (2.25)), given the current disturbances ξ_t, and expectations regarding future wages, prices, output, consumption and disturbances. Equation (2.36) describes the evolution of the costs of price dispersion over time.

For future reference, we remark that all firms together make total profits equal to

$$\frac{\Sigma_t}{P_t} = (1-\tau_t)Y_t - w_t H_t, \tag{2.37}$$

where $w_t H_t = \int_0^1 w_t(j) H_t(j) dj$.

2.4.3 Government Budget Constraint and Market Clearing Conditions

The government consumes goods $g_t Y_t$, imposes a sales tax τ_t, issues nominal bonds $\widetilde{B}_t \equiv P_t B_t$, and pays lump sum transfers T_t to households. The government budget constraint is given by

$$B_t = B_{t-1}\frac{1+i_{t-1}}{P_t/P_{t-1}} + \frac{T_t}{P_t} + (g_t - \tau_t)\,Y_t.$$

For simplicity, we assume that lump sum transfers (taxes if negative) are set such that they keep real government debt constant at some initial level B_{-1}. This implies that government transfers are given by

$$\frac{T_t}{P_t} = -(g_t - \tau_t)Y_t + B_{t-1}\left(1 - \frac{1+i_{t-1}}{P_t/P_{t-1}}\right). \tag{2.38}$$

Using (2.22) and (2.25), one can express the market clearing condition for the consumption/investment good as

$$Y_t = \frac{C_t + \Omega_t C_t^{\frac{\tilde{\sigma}-1}{1-\tilde{\alpha}}}}{1-g_t}, \tag{2.39}$$

where

$$\Omega_t \equiv \left(A_t^d \bar{C}_t^{-\tilde{\sigma}-1} q_t^u \right)^{\frac{1}{1-\tilde{\alpha}}} > 0 \tag{2.40}$$

is a term that depends on exogenous shocks and belief distortions in the housing market only, see equation (2.27). The previous two equations implicitly define a function

$$C_t = C(Y_t, q_t^u, \xi_t), \tag{2.41}$$

which delivers the market clearing consumption level, for a given output level Y_t, given housing prices q_t^u and given exogenous disturbances ξ_t.

The market clearing condition for housing is

$$D_t = (1-\delta)D_{t-1} + \tilde{d}(k_t; \xi_t), \tag{2.42}$$

and rental market clearing requires

$$D_t^R = 0. \tag{2.43}$$

Labor market clearing requires that the supply of labor of type j in (2.23) is equal to labor demand of industry j, which is given by (2.31), as all firms in the industry charge the same price. This delivers

$$w_t(j) = \frac{\tilde{v}_H(H_t(j);\xi_t)}{\tilde{u}_C(C_t;\xi_t)} = \frac{\lambda\,(H_t(j))^{\nu}\,\bar{H}_t^{-\nu}}{C_t^{-\tilde{\sigma}^{-1}}\bar{C}_t^{\tilde{\sigma}^{-1}}} = \lambda \frac{\bar{H}_t^{-\nu}}{\bar{C}_t^{\tilde{\sigma}^{-1}}} \left(\frac{Y_t}{A_t}\right)^{\nu\phi} C_t^{\tilde{\sigma}^{-1}} \left(\frac{p_t^j}{P_t}\right)^{-\nu\eta\phi}, \tag{2.44}$$

where $p_t^j = p_t^*$ in periods where industry j can adjust prices and $p_t^j = p_{t-1}^j$ otherwise.

2.4.4 Equilibrium and Ramsey Problem with Subjective Beliefs

We now define the equilibrium in the presence of subjective beliefs, as well as the nonlinear Ramsey problem characterizing the monetary policymaker's optimization problem in the presence of subjective beliefs.

We start by defining an *Internally Rational Expectations Equilibrium (IREE)*, which is a generalization of the notion of a Rational Expectations Equilibrium (REE) to settings with subjective private sector beliefs:

An internally rational expectations equilibrium (IREE) is a bounded stochastic process for $\{Y_t, C_t, k_t, D_t, \{w_t(j)\}, p_t^*, P_t, \Delta_t, q_t^u, i_t\}_{t=0}^{\infty}$ satisfying the aggregate supply equations (2.32), and

(2.34), the law of motion for the evolution of price distortions (2.36), the household optimality conditions (2.25), (2.27), (2.28), and the market clearing conditions (2.39), (2.42) and (2.44) for all j.

The equilibrium features ten variables (counting the continuum of wages as a single variable) that must satisfy nine conditions, leaving one degree of freedom to be determined by monetary policy.[30] In the special case with rational beliefs ($E_t^{\mathcal{P}}[\cdot] = E_t[\cdot]$), the IREE is a Rational Expectations Equilibrium (REE).

Given the equilibrium outcome, the remaining model variables can be determined as follows. Equilibrium profits are given by equation (2.37) and equilibrium taxes by equation (2.38). Equilibrium labor supply $H_t(j)$ follows from equation (2.23) for each labor type j. Equilibrium bond holdings satisfy $B_t = B_{-1}$ and equilibrium inflation is $\Pi_t \equiv P_t/P_{t-1}$. Equilibrium rental units are given by equation (2.43) and equilibrium rental prices by equation (2.26).

The Ramsey problem allows the policymaker to choose the sequence of policy rates, prices and allocations to maximize household utility, subject to the constraint that prices and allocations constitute an IREE. The policymaker thereby maximizes household utility under rational expectations, i.e., under a probability measure that is different from the one entertained by households, whenever the latter hold distorted beliefs. Benigno and Paciello (2014) refer to such a policymaker as a 'paternalistic' policymaker. The non-linear Ramsey problem is spelled out in Appendix B.3. To gain economic insights into the forces shaping the policy problem, the next section considers a quadratic approximation to the nonlinear problem.

2.5 The Monetary Policy Problem: Analytic Insights

This section derives analytic insights into the monetary policy problem. In particular, it presents a quadratic approximation to the policymaker's Ramsey problem that highlights the new economic forces arising from the presence of capital gain extrapolation.[31] It shows how subjective capital gain expectations shift the Phillips curve and affect the natural rate of interest in the IS equation.

The quadratic approximation derived below is valid for two alternative belief settings.[32] The first setting is standard and assumes rational expectations. While constituting a useful benchmark, the assumption of rational housing price expectations is strongly rejected by the survey evidence in Section 2.2.

[30] The transversality condition (2.29) must also be satisfied in equilibrium, but is not imposed as an equilibrium condition, as it will hold for all belief specifications considered below.

[31] The nonlinear problem can be found in Appendix B.3. The quadratic approximation delivers a valid second-order approximation to the problem, whenever (i) the steady-state Lagrange multipliers associated with the non-linear constraints are of order $O(1)$, which is the case when the steady state output distortion $\Theta \equiv \log\left(\frac{\eta}{\eta-1}\frac{1-g}{1-\tau}\right)$ is of order $O(1)$, and (ii) the gap between the steady-state interest rate and the lower bound, i.e., $\frac{1}{\beta} - 1$, is also of $O(1)$, i.e., when steady state real interest rates/natural rates are low.

[32] Recall from our earlier discussion that firms must hold beliefs about future values of P_t, $w_t(j)$, Y_t and that households must hold beliefs about future values of $(P_t, w_t(j), q_t^u, R_t, i_t, \Sigma_t/P_t, T_t/P_t)$. Both actors must additionally hold beliefs about the fundamental shocks entering their decision problem.

The second setting considers subjective housing beliefs. In particular, it considers capital gain extrapolation according to equations (2.7)-(2.9) introduced in the simple model in Section 2.3, but with the variable q_t being replaced by q_t^u. The latter implies that households extrapolate capital gains in units of marginal utility rather than in units of consumption. Specifying subjective beliefs in units of marginal utility leaves the ability of the learning rule to replicate the survey evidence unchanged, but has three advantages.[33]

First, the dynamics of housing prices in units of marginal utility is unaffected by monetary policy, even if housing prices in units of consumption do depend on policy. As a result, the object about which agents learn does not depend on policy. The policymaker thus cannot 'manipulate' households' subjective housing price beliefs in a way to achieve outcomes that are potentially better than under rational expectations.[34] In addition, it allows side-stepping the otherwise thorny issue of how the learning rule should respond to the conduct of monetary policy.

Second, the belief setup allows replicating the fact that housing demand/investment responds more strongly to monetary policy disturbances than non-housing demand, thereby avoiding the pitfalls described in Barsky et al. (2007). Appendix B.4 shows that in response to an exogenous shift in the path of nominal interest rates i, the change in housing investment and consumption satisfies at all times

$$\frac{d\log k_t}{d\boldsymbol{i}} = \frac{1}{1-\tilde{\alpha}}\frac{1}{\tilde{\sigma}} \cdot \frac{d\log C_t}{d\boldsymbol{i}}, \tag{2.45}$$

where $1/(1-\tilde{\alpha})$ is the price elasticity of housing supply and $1/\tilde{\sigma}$ the coefficient of relative risk aversion in consumption. The calibrated model considered later on features $1/((1-\tilde{\alpha})\tilde{\sigma}) > 1$.[35]

Third, the belief specification greatly simplifies the algebra involved in deriving the second-order approximation to the Ramsey problem, because it allows for a relatively straightforward determination of the equilibrium path of subjectively optimal consumption choices.

Overall, we wish to consider a minimal deviation from rational expectations, therefore keep expectations about all other variables rational to the extent possible.[36] Finally, to insure that households' subjectively optimal plans satisfy the transversality condition, we assume that households hold rational capital gain expectations in the very long run, i.e., after some arbitrarily large but finite period $\bar{T} < \infty$.[37] We then consider the policy problem with subjective beliefs in periods

[33]This is so because we consider log consumption preferences which imply that contributions from fluctuations in marginal utility are orders of magnitude smaller than those generated by subjective beliefs.

[34]This is a key distinction to the setups analyzed in Molnar and Santoro (2014), Mele et al. (2020), and Caines and Winkler (2021).

[35]The calibration use log utility in consumption ($1/\tilde{\sigma} = 1$) and a supply elasticity of $1/(1-\tilde{\alpha}) = 5$.

[36]In particular, household continue to hold rational expectations about all other prices, i.e., about $\{P_t, w_t(j), i_t\}$ and firms hold rational expectations about $\{P_t, w_t(j), Y_t\}$. Furthermore, all actors continue to hold rational expectations about the exogenous fundamentals. Beliefs about profits and lump sum taxes, $\{\Sigma_t/P_t, T_t/P_t\}$ continue to be determined by equations (2.37) and (2.38), evaluated with rational output expectations and the state-contingent optimal choices for $\{H_t, k_t, B_t\}$. Rental price expectations, however, cannot be kept rational: they need to satisfy equation (2.26), which shows that they are influenced by the subjectively optimal consumption plans implied by equation (2.28).

[37]Appendix B.5 shows that this is sufficient to insure that subjectively optimal plans satisfy the transversality

$t \ll \overline{T}$.

For the two belief settings just described, the quadratic approximation of the Ramsey problem is given by[38]

$$\max_{\{\pi_t, y_t^{gap}, \widehat{q}_t^u, i_t \geq \underline{i}\}} -E_0 \sum_{t=0}^{\infty} \beta^t \frac{1}{2} \left(\Lambda_\pi \pi_t^2 + \Lambda_y \left(y_t^{gap}\right)^2 + \Lambda_q \left(\widehat{q}_t^u - \widehat{q}_t^{u*}\right)^2 \right) \tag{2.46}$$

s.t.:

$$\pi_t = \beta E_t \pi_{t+1} + \kappa_y y_t^{gap} + \kappa_q \left(\widehat{q}_t^u - \widehat{q}_t^{u*}\right) + u_t \text{ for } t \geq 0 \tag{2.47}$$

$$y_t^{gap} = \lim_{T \to \infty} E_t y_T^{gap} - \varphi E_t \sum_{k=0}^{\infty} \left(i_{t+k} - \pi_{t+1+k} - r_{t+k}^{*,RE} \right) - \frac{C_q}{C_Y} \left(\widehat{q}_t^u - \widehat{q}_t^{u*}\right) \tag{2.48}$$

as well as equations determining $\left(\widehat{q}_t^u - \widehat{q}_t^{u*}\right)$ and initial pre-commitments,

where $\pi_t = \log \Pi_t$ denotes inflation and y_t^{gap} the output gap, which is defined as $y_t^{gap} = \log Y_t - \log Y_t^*$, with Y_t^* denoting the efficient level of output, as defined in equation (B.15) in appendix B.6. The housing price gap $\widehat{q}_t^u - \widehat{q}_t^{u*}$ is the difference between the housing price $\widehat{q}_t^u = \log q_t^u$ and its efficient welfare-maximizing level $\widehat{q}_t^{u*}$, which is given by[39]

$$q_t^{u*} = \overline{\xi}_t^d, \tag{2.49}$$

where $\overline{\xi}_t^d \equiv \sum_{T=t}^{\infty} E_t[(1-\delta)^{T-t} \beta^{T-t} \xi_T^d]$.

The policymaker's objective (2.46) involves the standard terms of squared inflation and the squared output gap, but also depends on the squared housing price gap. The latter arises because any deviation of housing prices from their efficient level distorts – for a given level of the output gap – housing investment, as we explain below. The equilibrium value of the housing price gap will depend on the belief specification and will be discussed in detail in the next two sections.

Constraint (2.47) is the New Keynesian Phillips Curve and depends on the housing price gap. The coefficients $\kappa_q < 0$ and $\kappa_y > 0$ are defined in Appendix B.6.3 and imply that positive housing price gaps exert *negative* cost-push effects: high housing prices increase housing investment and – for a given output gap – decrease non-housing consumption. The latter raises the marginal utility of non-housing consumption and thereby depresses wages and marginal costs. This allows the model to potentially produce a non-inflationary boom in housing prices and housing investment. The mark-up disturbance u_t is a function of exogenous disturbances only.

Constraint (2.48) is the linearized and forward-iterated IS equation. A key new insight here is that the IS equation also depends on the housing price gap. This implies that the housing

constraint (2.29).

[38]See Appendix B.6 for a derivation.

[39]See the derivation in Appendix B.8.4. All variables in the approximation are expressed in terms of log deviations from the efficient steady state.

price gap affects the natural rate of interest, as discussed in detail below. The coefficients $C_q < 0$ and $C_Y > 0$ are the derivatives of the function $C(\cdot)$ defined in (2.41) with respect to q^u and Y, respectively, evaluated at the efficient steady state. The variable $r_t^{*,RE}$ in equation (2.48) denotes the natural interest rate under RE and is a function of exogenous disturbances only.[40] The long-run output gap expectations $\lim_T E_t y_T^{gap}$ in equation (2.48) are the ones associated with a setting in which agents hold rational housing expectations.[41]

Note, that the policymaker's choice of the nominal interest rate i_t is subject to an effective lower bound $i_t \geq \underline{i}$, where the bound $\underline{i} < 0$ is expressed in terms of deviations from the interest rate in a zero-inflation steady state. For the special case with a zero lower bound, we have $\underline{i} = -(1-\beta)/\beta$. In the absence of a lower bound constraint or when economic shocks never cause the bound to become binding, the IS equation (2.48) can be dropped from the policy problem.

Interestingly, the expectations showing up in the monetary policy problem (2.46)-(2.48) are all rational. The way subjective housing price expectations affect the monetary policy problem are thus fully captured through their effects on the housing price gap. The next two sections determine the housing price gaps under rational and subjective beliefs and what they imply for optimal policy.

2.5.1 Rational Housing Price Expectations

With fully rational expectations we have

$$\widehat{q}_t^{u,RE} = \widehat{q}_t^{u*}, \tag{2.50}$$

which shows that the housing price *gap* is zero at all times, independently of monetary policy and independently of the economic disturbances hitting the economy.[42] Under RE, the Ramsey problem with a lower bound constraint (2.46) is thus isomorphic to the Ramsey problem in a standard New Keynesian model without a housing sector, as considered for instance in Adam and Billi (2006a). This result may appear surprising because monetary policy decisions do affect the housing price in units of consumption $\widehat{q}_t$. Yet, as the policy problem (2.46) makes clear, it is only the housing price gap in units of marginal utility, $\widehat{q}_t^u - \widehat{q}_t^{u*}$, that is relevant from a welfare perspective. Under RE, the presence of a housing sector thus generates no fundamentally new economic insights into the monetary policy problem.[43]

The RE setup also has difficulties in making a connection between the average natural rate of

[40] More precisely, $r_t^{*,RE}$ is the real interest rate consistent with the optimal consumption level in a setting with flexible prices and fully rational expectations, see Appendix B.6.4 for details.

[41] Recall that housing expectations are assumed rational in the long-run in both belief settings.

[42] See Appendix B.8.1 for proofs on the results about housing prices and the price-rent ratio presented in this section.

[43] The inclusion of a housing sector only affects the definition of the output gap, which now also depends on housing sector disturbances.

interest and the volatility of the price-to-rent (PR) ratio. Under RE, the equilibrium PR-ratio is

$$PR_t^{RE} = \frac{q_t^u}{\xi_t^d}, \tag{2.51}$$

which to a first-order approximation is given by

$$\widehat{PR}_t^{RE} = Z \cdot \widehat{\xi}_t^d, \tag{2.52}$$

with $Z \equiv \beta(1-\delta)\left(\rho_\xi - 1\right)/(1-\beta(1-\delta)\rho_\xi)$. Equation (2.52) shows that the PR-ratio displays persistent variation under RE, if and only if housing demand shocks $\widehat{\xi}_t^d$ are persistent. In fact, replicating the high quarterly auto-correlation of the PR-ratio in Table 2.3 requires choosing a shock persistence ρ_ξ very close to one. Yet, in the limit $\rho_\xi \to 1$, the derivative $\partial Z/\partial\beta$ uniformly converges to zero for all $\beta \in [0,1]$. This implies that the volatility of the PR ratio will be largely independent of the natural rate of interest when housing demand shocks are sufficiently persistent. Under RE, there is thus no quantitatively important relationship between the average natural rate of interest and the volatility of the PR ratio, unlike in the case with capital gain extrapolation.

Given equation (2.50), the IS equation (2.48) implies that setting

$$i_t - E_t\pi_{t+1} = r_t^{*,RE} \quad \text{for all } t \geq 0 \tag{2.53}$$

is consistent with a constant output gap, i.e.,

$$y_t^{gap} = \lim_T E_t y_T^{gap} \quad \text{for all } t \geq 0.$$

This justifies our interpretation of $r_t^{*,RE}$ as the natural rate of interest under RE.[44] It also shows that the volatility of the natural rate of interest is independent of the average value of the natural rate under RE. This will cease to be the case under subjective housing beliefs.

2.5.2 Subjective Housing Price Expectations

This section discusses three new economic forces showing up in the monetary policy problem in the presence of subjective housing price beliefs. It shows (i) how housing price fluctuations are affected by the average level of the natural rate of interest, (ii) how these fluctuations affect the volatility of the natural rate of interest, and (iii) how these fluctuations distort the allocation of output.

Housing prices under subjective beliefs are jointly determined by equations (2.9) and (2.10), where q_t should again be replaced by q_t^u. Since these equations do not depend on policy, the

[44] In the presence of a lower bound constraint on nominal rates, it might not be feasible to implement (2.53) at all times.

policymaker can treat the housing price gap as exogenous, as is the case with RE.[45] Yet, the housing price gap will now generally differ from zero, as the housing price gap can become positive or negative depending on the degree of capital gain optimism/pessimism.

The average natural rate and housing price volatility. With subjective housing price expectations, the equilibrium housing price is given by[46]

$$q_t^{u,\mathcal{P}} = \frac{1}{1-\beta(1-\delta)\beta_t}\xi_t^d \tag{2.54}$$

and the price-to-rent ratio by

$$PR_t^{\mathcal{P}} = \frac{q_t^{u,\mathcal{P}}}{\xi_t^d}. \tag{2.55}$$

For the limit with persistent housing demand shocks ($\rho_\xi \to 1$), we can derive the first-order approximation

$$\widehat{q}_t^{u,\mathcal{P}} = \widehat{q}_t^{u,RE} + (\beta_t - 1)\frac{\beta(1-\delta)}{1-\beta(1-\delta)\beta_t}\left(1+\widehat{\xi}_t^d\right), \tag{2.56}$$

which decomposes the equilibrium housing price into its RE value plus a contribution coming from the presence of subjective beliefs. We then also have

$$E_t^{\mathcal{P}}\left[\widehat{q}_{t+1}^{u,\mathcal{P}}\right] = E_t\left[\widehat{q}_{t+1}^{u,RE}\right] + (\beta_t - 1)\left[1+\frac{\beta(1-\delta)}{1-\beta(1-\delta)\beta_t}\left(1+\widehat{\xi}_t^d\right)\right], \tag{2.57}$$

which shows that subjective housing price expectations are equal to their RE equilibrium value whenever expected capital gains are equal to one ($\beta_t = 1$). Capital gain extrapolation, however, will induce fluctuations of β_t around one and thus drive a wedge between the housing price under learning and RE.[47]

As explained for the simple model in Section 2.3, lower values for the average natural rate (discount factors β closer to one), will induce stronger fluctuations in capital gain expectations (β_t), because housing prices are more sensitive to belief revisions, see equation (2.54). Lower average values for the natural rates will thus be associated with increased fluctuations in housing prices and the PR-ratio, in line with empirical evidence presented in Section 2.3.

Housing price fluctuations and the natural rate of interest. The presence of non-zero housing price gaps also affects the natural rate of interest. This can be seen by considering a policy that sets real interest rates equal to the RE natural real rate ($r_t^{*,RE}$). Such a policy now ceases to deliver a constant output gap, instead implies

[45] This does not imply that the housing price q_t is invariant to monetary policy: monetary can determine how variations in q_t^u get split up into variations of the housing price q_t and variations in marginal utility $\widetilde{u}_C(C_t;\xi_t)$.

[46] See Appendix B.8.1 for a derivation of this and subsequent results, including the generalized expressions for the case with $\rho_\xi < 1$.

[47] In the limit where the Kalman gain ($1/\alpha$) in the updating equation (2.9) approaches zero, the model with capital gain extrapolation converges to the RE model.

$$y_t^{gap} = \lim_T E_t y_T^{gap} - \frac{C_q}{C_Y}\left(\widehat{q}_t^u - \widehat{q}_t^{u*}\right). \tag{2.58}$$

Since $C_q/C_Y < 0$, a positive (negative) housing price gap is then associated with a positive (negative) output gap: high housing prices stimulate housing investment and thereby output. Since the output expansion is inefficient, the policymaker might find it optimal to *lean against housing prices*. The extent to which this is optimal will be explored quantitatively in Section 2.7 below.

The following lemma derives the natural rate $r_t^{*,\mathcal{P}}$ for our setting with subjective housing price beliefs:[48]

Lemma 2. *Let the natural rate of interest under subjective beliefs be given by*

$$r_t^{*,\mathcal{P}} \equiv r_t^{*,RE} - \frac{1}{\varphi}\frac{C_q}{C_Y}\left(\left(\widehat{q}_t^u - \widehat{q}_t^{u*}\right) - E_t\left(\widehat{q}_{t+1}^u - \widehat{q}_{t+1}^{u*}\right)\right) \qquad \textit{for all } t. \tag{2.59}$$

When real interest rates are equal to $r_t^{*,\mathcal{P}}$ for all $t \geq 0$, then the IS equation (2.48) is consistent with

$$y_t^{gap} = \lim_T E_t y_T^{gap} \qquad \textit{for all } t. \tag{2.60}$$

The proof can be found in Appendix B.8.1. Equation (2.59) generalizes the natural interest rate definition under RE to a setting with potentially subjective beliefs. In the special case with a constant housing price gap, we have $r_t^{*,\mathcal{P}} = r_t^{*,RE}$, even when the constant housing price gap differs from zero. This shows that the natural rate under subjective beliefs differs from its RE value if and only if the housing price gap is expected to go up or down. Since $C_q/C_Y < 0$, the natural rate will exceed (fall short of) its RE level, when the current housing price gap is higher (lower) than tomorrow's (expected) gap.

Since fluctuations in housing prices become larger when the average natural rate falls, the expected changes in the housing price gap will also become more volatile. A lower average level of the natural rate is thus not only associated with more volatile housing prices but also with more volatile natural rates of interest.

Figure 2.6 shows that this model prediction is consistent with the data. The figure plots the changes in the average level of the natural rate from the period before 1990 to the period after 1990 on the horizontal axis and the corresponding change in the natural rate *volatility* on the vertical axis. The volatilities of the natural rates in the pre-/post-1990 periods are the standard deviations of the linearly detrended series. The figure is again based on the estimates in Holston et al. (2017). While the level of the natural rate decreased, the volatility of it increased in four out of the five advanced economies. Appendix B.7 discusses the robustness of these results.

[48] As is the case with RE, it will generally not be optimal (or not even feasible) to set interest rates equal to the natural rate at all times due to the presence of a lower bound constraint on nominal rates

Figure 2.6: Changes in the average natural rate vs. changes in the volatility of the natural rate

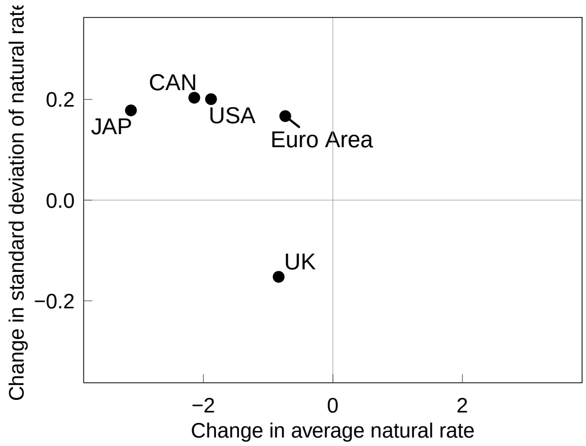

Notes: This figure plots the pre-/post-1990 changes in the average natural rates against the changes in the natural rate volatility for several advanced economies. The volatilities of the natural rates in the pre-/post-1990 periods are the standard deviations of the linearly detrended series.

Housing price fluctuations and the misallocation of output. We now show that fluctuations in the housing price gap distort the allocation of output between its alternative uses, i.e., between housing investment and non-housing consumption. The housing investment gap, i.e., the difference between actual investment $\widehat{k}_t$ and its efficient level $\widehat{k}_t^*$, is – to a first-order approximation – given by

$$\widehat{k}_t - \widehat{k}_t^* = \frac{\tilde{\sigma}^{-1} C_Y}{1-\tilde{\alpha}} y_t^{gap} + \frac{1+\tilde{\sigma}^{-1} C_q}{1-\tilde{\alpha}} \left(\widehat{q}_t^u - \widehat{q}_t^{u*}\right). \tag{2.61}$$

Under rational expectations, the housing price gap is zero and the investment gap is only distorted to the extent that the output gap is not closed. Additional output then gets allocated in constant proportions to housing investment and non-housing consumption, as $(\tilde{\sigma}^{-1} C_Y)/(1-\tilde{\alpha}) > 0$. In the presence of subjective beliefs, however, an additional distortion arises: the housing investment gap is then also driven by the housing price gap. Given the calibration considered later on, we have $(1+\tilde{\sigma}^{-1} C_q)/(1-\tilde{\alpha}) > 0$, so that a positive housing price gap ($\widehat{q}_t^u - \widehat{q}_t^{u*} > 0$) reinforces the investment distortions generated by a positive output gap.[49] This explains why the squared housing price gap shows up in the policymaker's objective function (2.46). While monetary policy cannot affect the housing price gap within our belief setup, it is the case that larger housing price gap fluctuations, as induced by lower natural rates, contribute to increased welfare losses.

2.6 Model Calibration

To explore the quantitative implications for monetary policy arising from the presence of capital gain extrapolation, we consider a calibrated model. The calibration strategy consists of choosing

[49]This distortion in the allocation of output between housing investment and non-housing consumption is present independently of other frictions such as sticky prices or the lower-bound constraint on nominal interest rates.

a set of standard parameter values previously considered in the literature and of matching salient features of the behavior of natural interest rates and housing prices in the United States in the pre-1990 period. We then test the model by considering its predictions for the lower natural rate levels observed in the post-1990 period up to 2021. We compare across long time spans of 30 years each to obtain more reliable estimates of housing price volatility, which is difficult to estimate given the high degree of serial correlation of housing prices.

Table 2.5: Model calibration

Parameter	Value	Source/Target
Preferences and technology		
β	0.9917	Average U.S. natural rate pre 1990
φ	1	Adam and Billi (2006a)
κ_y	0.057	Adam and Billi (2006a)
$\frac{\Lambda_y}{\Lambda_\pi}$	0.007	Adam and Billi (2006a)
κ_q	-0.0023	Adam and Woodford (2021)
$\frac{C_q}{C_Y}$	-0.29633	Adam and Woodford (2021)
δ	0.03/4	Adam and Woodford (2021)
Exogenous shock processes		
ρ_{r^*}	0.8	Adam and Billi (2006a)
σ_{r^*}	0.2940% (RE)	Adam and Billi (2006a)
	0.1394% (subj. beliefs)	
ρ_ξ	0.99	Quarterly autocorrel. of the PR-ratio of 0.99
σ_{ξ^d}	0.0233 (RE)	Std. dev. of price-to-rent ratio pre 1990
	0.0165 (subj. beliefs)	
Subjective belief parameters		
α	1/0.007	Adam et al. (2016)
β^U	1.0031	Max. percentage deviation of PR-ratio from mean

Calibration to the pre-1990 period. Table 3.4 summarizes the model parameterization. The quarterly discount factor β is chosen such that the steady-state natural rate equals the pre-1990 average of the U.S. natural rate of 3.34%, as estimated by Holston et al. (2017). The interest rate elasticity of output φ, the slope of the Phillips curve κ_y, and the welfare weight $\frac{\Lambda_y}{\Lambda_\pi}$ are taken from Table 2 in Adam and Billi (2006a). The Phillips Curve coefficient κ_q and the ratio C_q/C_y are set as in Adam and Woodford (2021).[50]

We now discuss the parameterization of the exogenous shock processes. The persistence of the housing preference shock ρ_ξ is set such that the RE model captures the high serial autocorrelation of the PR ratio in the data. The standard deviation of the innovations to the housing preferences

[50]The calibration target for the ratio C_q/C_y is the ratio of residential fixed investment over the sum of nonresidential fixed investment and personal consumption expenditure, which is on average approximately equal to 6.3% in the US. This and the remaining parameters then imply $\kappa_q = -0.0023$, see Appendix B.8.8 for details.

Figure 2.7: Standard deviation of price-to-rent ratio and natural rate

(a) Standard deviation of price-to-rent ratio (relative to corresponding mean)

(b) Standard deviation of the natural rate relative to case with $\underline{r}^{*,RE} = 3.34\%$

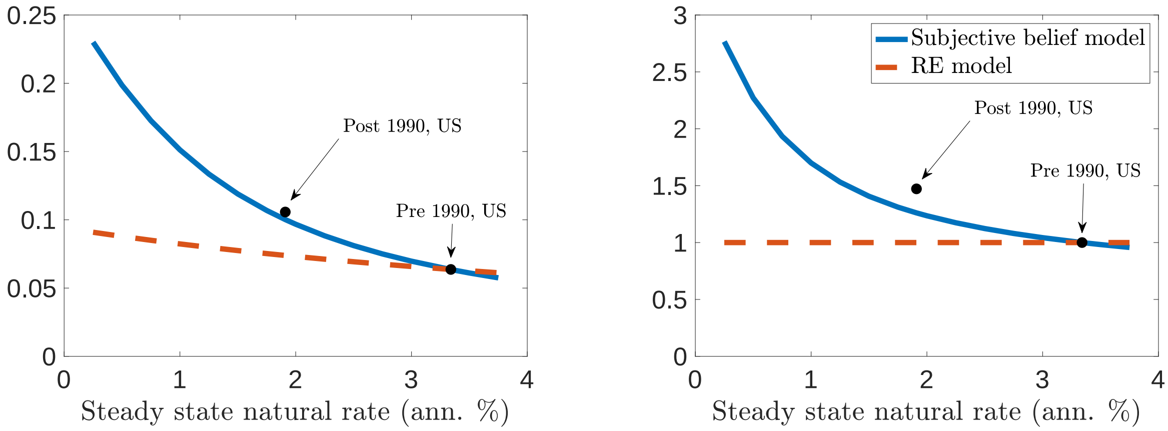

Notes: This figure plots, for different steady state levels of the natural rate, the standard deviation of the price-to-rent ratio (relative to its mean) and the standard deviation of the natural rate.

σ_ξ are set such that the rational expectations and subjective belief models both replicate the pre-1990 standard deviation of the PR-ratio. For the subjective belief model, this is achieved by simulating equations (2.9) and (2.11), which requires specifying the belief updating parameters α and β^U. We set $\alpha = 1/0.007$ following Adam et al. (2016) and determine σ_{ξ^d} and β^U jointly such that (i) we match the volatility of the price-to-rent ratio and (ii) the simulated data matches the maximum deviation of the price-to-rent ratio in the data from its sample mean. The latter statistic identifies β^u. This procedure yields $\beta^U = 1.0031$ and $\sigma_{\xi^d} = 0.0165$. Housing demand disturbances are less volatile than under RE because fluctuations in subjective beliefs contribute to the fluctuations in housing prices. In fact, the calibration implies that about 50% of housing price fluctuations are due to subjective beliefs.

We consider the natural rate process

$$r_t^{*,RE} = \rho_{r*} r_{t-1}^{*,RE} + \varepsilon_t^r, \tag{2.62}$$

where $\varepsilon_t^r \sim iiN(0, \sigma_{r^*}^2)$. For the RE model, we set ρ_{r^*} and σ_{r^*} equal to the values in Adam and Billi (2006a). For the subjective believe model, we use the same value for ρ_{r^*} but choose σ_{r^*} such that the generalized natural rate for the subjective belief model, defined in equation (2.59), has the same volatility as the natural rate in the RE model. This yields $\sigma_{r^*,RE} = 0.1393\%$, which is lower than under RE, because fluctuations in the housing price gap contribute to fluctuations in the natural rate in the presence of subjective beliefs. To economize on the number of state variables in the policy problem, we abstract from the presence of mark-up shocks.[51]

[51] Adam and Billi (2006a) show that mark-up shocks are too small and display too little persistence to cause the

Evaluation of the model in the post-1990 period. Figure 2.7 illustrates the predictions of the RE model (dashed line) and subjective belief model (solid line) for the standard deviation of the price-to-rent ratio (panel a) and the standard deviation of the natural rate of interest (panel b). The panels depict these outcomes, which are independent of monetary policy, on the vertical axis for various levels of the steady-state natural rate on the horizontal axis. Variations in the steady-state level of the natural rate are achieved via appropriate variations in the discount factor.[52] The dots in Figure 2.7 report the average values for the pre- and post-1990 U.S. sample periods, where the average natural rate was equal to 3.34% and 1.91%, respectively.[53]

Since the model has been calibrated to the pre-1990 period, the RE and subjective belief model both match the pre-1990 data points in Figure 2.7. The subjective belief model also performs quite well in matching the post-1990 outcomes, despite the fact that these outcomes are not calibration targets. In particular, the standard deviation of the price-to-rent ratio and the standard deviation of the natural rate endogenously increase as the natural rate falls, with the magnitudes roughly matching the increase observed in the data. In contrast, the RE model produces no increase in the volatility of the natural rate and only a weak increase in the volatility of the price-to-rent ratio, for reasons discussed in Section 2.5.1. Matching the increase in housing price volatility under RE requires increasing the volatility of housing demand shocks. Since such an increase is irrelevant for monetary policy under RE, we leave the volatility of housing preference shocks unchanged. Similarly, matching the increase in the natural rate volatility under RE would require increasing σ_{r^*}. We will consider such increases when discussing our quantitative results.

2.7 Quantitative Implications for Monetary Policy

This section illustrates the quantitative implications of falling natural rates and rising housing price volatility for the conduct of optimal monetary policy. It starts by determining the implications of falling natural rates for the optimal inflation target, i.e., for the average inflation rate implied by optimal monetary policy. It then illustrates the dramatically different optimal response to housing demand shocks under subjective and objective housing beliefs. Details of the nonlinear numerical solution procedure underlying the results in this section can be found in Appendix B.8.6.

2.7.1 The Optimal Inflation Target

Figure 2.8 depicts the optimal inflation target for different steady-state levels of the natural rate of interest, i.e., the average inflation rate implied by optimal monetary policy. It shows the optimal target for the setup with subjective housing beliefs (upper line), for the case with rational housing

lower-bound constraint to become binding.

[52] As discussed before, variations in the discount factor may be driven by variations in the long-term growth rate and/or by variations in time-preferences.

[53] The reported increase in the standard deviation of the natural rate is again based on the estimates in Holston et al. (2017).

Figure 2.8: Average inflation under optimal monetary policy

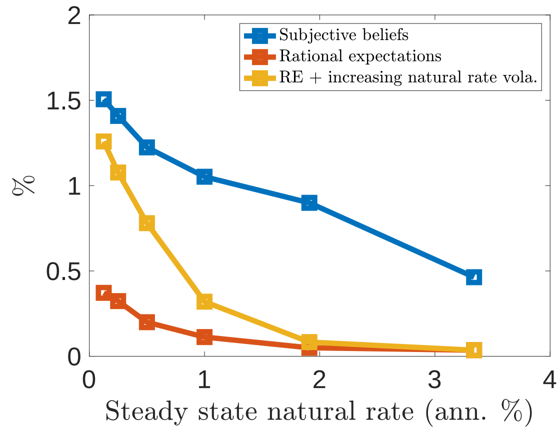

Notes: The figure reports the optimal inflation target for different average levels of the natural rate in the presence of a zero lower bound constraint. The red line depicts the optimal target for the case with rational housing price beliefs and the blue line the one with subjective housing price beliefs. The yellow line shows the optimal average inflation under RE where the exogenous volatility of the natural rate is adjusted such that it matches the endogenous volatility increase under subjective beliefs.

price beliefs (lower line), and for a third case that we discuss below.

We find that the optimal target is close to zero, whenever housing expectations are rational. This holds quite independently of the average level of the natural rate, confirming earlier findings in Adam and Billi (2006a) who considered the value for the average natural rate at the upper end of the range shown in Figure 2.8. This may appear surprising given that it is optimal for monetary policy to promise future inflation, so as to lower real interest rates, whenever adverse natural rate shocks cause the lower-bound constraint on nominal rates to bind. While the lower bound is reached more often when the average natural rate is low, inflation promises still have to be made relatively infrequently and can be quite modest. Hence, they do not significantly affect the average rate of inflation.

This result differs quite substantially from the ones reported in Andrade et al. (2019), who find that the optimal target should move up approximately one-to-one with a fall in the natural rate under rational expectations. Besides that Andrade et al. (2019) consider a medium-scale sticky price model without housing, the main difference to our approach is that they study Taylor rules with optimized intercepts rather than optimal monetary policy. As shown in Coibion et al. (2012) it makes a big difference for the optimal inflation target whether the monetary policy maker follows a Taylor rule or Ramsey optimal policy.

While lower natural rates trigger (slightly) larger housing price fluctuations under rational expectations, increased volatility is fully efficient and does not affect the natural rate of interest. Under rational expectations, the optimal inflation target is thus unaffected by housing price fluctuations, including for very low levels of the natural rate.

The upper line in Figure 2.8 shows that the situation is quite different with subjective housing beliefs. The optimal inflation target is overall substantially higher and also reacts more strongly to a fall in the average natural rate of interest. In fact, a fall in the steady-state natural rate from its pre-1990 average (3.34%) to its post-1990 average (1.9%) causes the optimal inflation target to increase by almost 0.5%. The corresponding increase under rational expectations is less than 0.05%. This difference is due to the fact that the endogenous volatility component of the natural rate increases once the natural rate drops. This reinforces the stringency of the zero lower bound, but is an effect that is absent under RE. It requires that the central bank engages more often in inflation promises, as it faces the lower bound constraint.

The optimal inflation target with subjective housing beliefs is substantially higher than the optimal target with RE, even at the pre-1990 average level of the natural rate. This is the case although the volatility of the natural rate is calibrated at this point to be equal across both models. This is due to two reasons: First, fluctuations in the housing price gap also generate cost-push term in the Phillips curve. Second, belief fluctuations induce more persistent variations in the natural rate than the exogenous natural rate shocks. This puts further upward pressure on the optimal inflation rate, as it requires larger and more persistent inflation promises by the central bank.

To illustrate this last point, the middle line in Figure 2.8 depicts the optimal inflation rate under rational expectations, when we set the volatility of the (exogenous) natural rate in the RE model such that it matches the volatility of the natural rate in the subjective belief model, for each considered level of the natural rate. While the optimal inflation rate increases relative to the benchmark RE setting, the level of the optimal inflation target still falls short of the one implied by subjective beliefs.

2.7.2 Leaning Against Housing Demand Shocks

We now examine the optimal monetary policy response to housing demand shocks. Under RE, housing demand shocks affect the housing price and the efficient housing price identically, so that the housing price gap remains at zero. As a result, neither the output gap nor inflation respond to housing demand shocks. In contrast, it becomes optimal to "lean against" housing demand shocks in the presence of subjective beliefs. Yet, due to the lower bound constraint, the optimal response to positive and negative housing demand shocks displays considerable asymmetry.

The top row in Figure 2.9 shows the response of housing-related variables to a persistent positive/negative housing demand shock of 5%.[54] On impact, the shock triggers capital gains of an equal amount, which then trigger belief revisions that fuel further upward movements of the housing price in the same direction. The positive shock, for instance, pushes housing prices up

[54]We initialize the economy at its ergodic mean and then hit the economy with a one-time shock of three standard deviations. We then average the subsequent response over the possible future shock realizations. We assume a steady-state natural rate equal to its post-1990 mean (1.91%).

Figure 2.9: Impulse responses to a housing preference shock

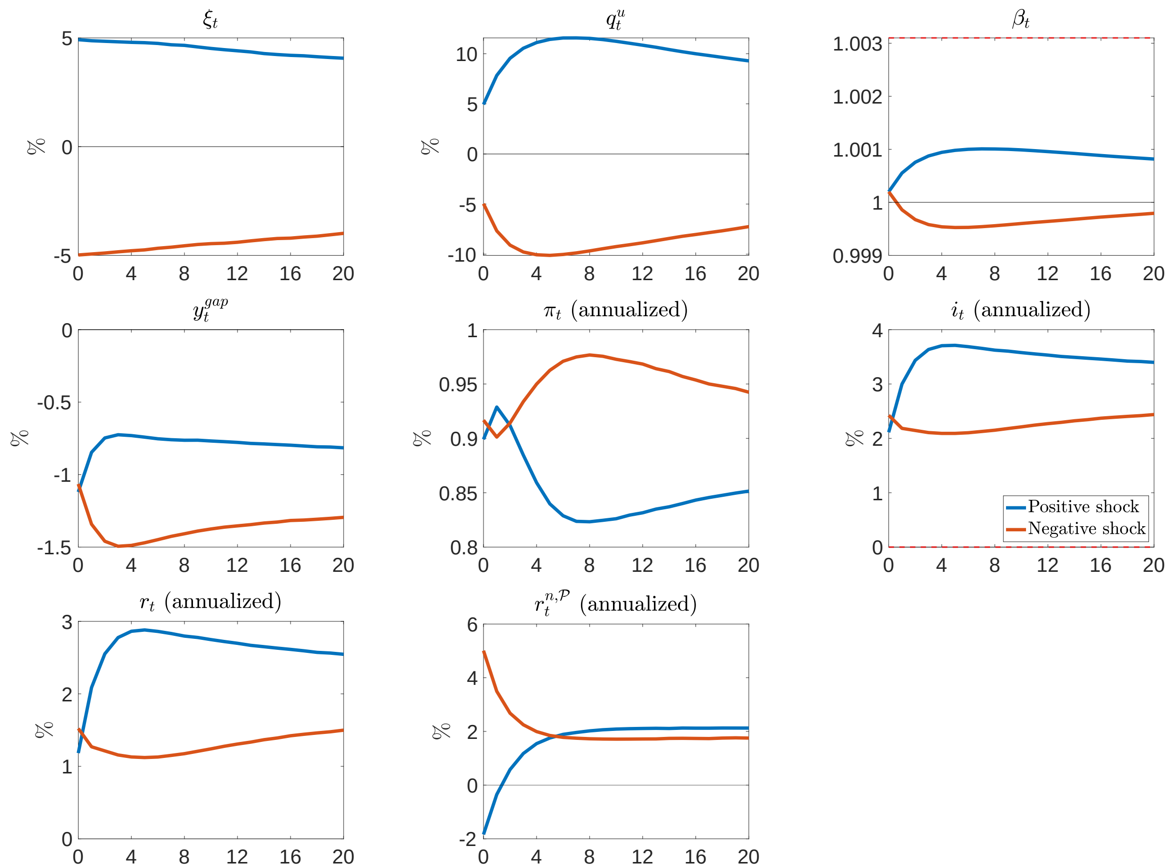

Notes: The figure reports the average impulse responses of the economy under subjective beliefs (at $\underline{r}^* = 1.91\%$) after a three-standard-deviation housing demand shock. The blue lines show the responses after a positive shock and the red lines after a negative shock.

by about 5% on impact, with belief momentum generating approximately another 5% in the first six quarters after the shock. This causes the housing price gap to become significantly positive (not shown in the figure). Once actual housing price increases start to fall short of the expected housing price increases, the housing boom reverts direction.

Higher housing prices push up housing investment, which causes upward pressure on the output gap. Optimal monetary policy leans strongly against the housing price and increases nominal and real interest rates. It does so despite the fact that the natural rate of interest falls in response to the shock. The policy response causes a fall in inflation, which is amplified by the fact that the increase in housing prices and investment increases the marginal utility of consumption, hence, dampens wages and marginal costs. A positive housing demand shock thus results – in the presence of subjective housing beliefs – in a disinflationary housing boom episode under optimal monetary policy.

The policy response to a positive housing demand shock is much stronger than that to a negative housing demand shock. In particular, nominal and real interest rates fall considerably less following a negative shock realization. This is so because a negative housing price gap is inflationary and inflation is already high to start with. Negative housing demand shocks thus move inflation further away from its optimal level of zero.[55] Yet, policy still "leans against" the housing price decrease: real interest rates fall despite the fact that the natural rate increases.

The fact that leaning against housing prices can be optimal in the presence of capital gain extrapolation is in line with results in Caines and Winkler (2021), who consider a setting with 'conditionally model consistent beliefs' in which expectations differ for many variables from rational expectations, and with results in Adam and Woodford (2021), who consider a setting where the policymaker fears 'worst-case' belief distortions about inflation and housing price expectations. As none of these papers consider a lower-bound constraint, the policy response to positive and negative shocks is symmetric in their settings.

2.8 The Role of Macroprudential Policy

It is often argued that macroprudential policies can be used to stabilize financial markets and that this would allow monetary policy to ignore disturbances coming from the housing sector, see Svensson (2018) for a prominent exposition of this view. In this section, we evaluate the quantitative plausibility of this view within our setup with subjective housing beliefs.

We show below that fully eliminating fluctuations in the housing price gap requires imposing large and volatile macroprudential taxes. None of the macroprudential instruments thus far available in advanced economies appear suited to achieve economic effects anywhere near the required size. In addition, it is often necessary for macroprudential policy to pay substantial subsidies. To the best of our knowledge, none of the available macroprudential instruments acts in a way that subsidizes actions by economic actors. Less aggressive policies, that aim at only partly eliminating the housing price gap, still require considerable tax volatility, because fluctuations in subjective beliefs turn out not to be independent of tax policy pursued.

We analyze the issue by considering a setup in which the policymaker can tax or subsidize the ownership of housing. While actual macroprudential policies often operate via constraints imposed on the banking sector, their ultimate effect is to make housing more or less expensive to households. For this reason, we consider taxes and subsidies at the household level.

Specifically, we analyze a proportional and time-varying tax τ_t^D that is applied to the rental value of housing in every period t. A household owning D_t units of houses, then has to pay taxes of

$$\tau_t^D D_t R_t \tag{2.63}$$

[55] While the output gap is moved closer to its optimal level, the weight on the output gap in the welfare function is two orders so magnitude smaller than that on inflation, see Table 3.4.

units of consumption.[56] We find this specification more plausible than a policy that taxes the market value of housing, as it is difficult to determine market values in real time. A setup that taxes the physical housing units, i.e., where taxes are equal to $\tau_t^D D_t$, delivers very similar results, but is analytically more cumbersome. Furthermore, the tax setup in equation (2.63) is equivalent to a setup where taxes directly affect household utility, i.e., where the utility contribution from owning houses would instead be given by $\xi_t^d \left(1 - \tau_t^D\right) D_t$ and no monetary taxes would have to be paid. We prefer the formulation in equation (2.63) because it allows expressing taxes in monetary units.

In the presence of these taxes, housing prices under subjective beliefs are given by

$$q_t^u = \frac{\left(1 - \tau_t^D\right) \xi_t^d}{1 - \beta(1-\delta)\beta_t}, \tag{2.64}$$

and the housing-price *gap* in percentage deviations from the steady state (where $\underline{\tau}^D = 0$) is

$$\widehat{q}_t^u - \widehat{q}_t^{u*} = \frac{(1 - \beta(1-\delta))\left(1 - \tau_t^D\right) \widehat{\xi}_t^d}{1 - \beta(1-\delta)\beta_t} + \frac{\beta(1-\delta)(\beta_t - 1)}{1 - \beta(1-\delta)\beta_t} - \frac{1 - \beta(1-\delta)}{1 - \beta(1-\delta)\beta_t}\tau_t^D - \frac{1 - \beta(1-\delta)}{1 - \beta(1-\delta)\rho_\xi}\widehat{\xi}_t^d.$$

The previous equation shows that macroprudential policy must eliminate housing price gap fluctuations that are due to housing demand shocks ($\widehat{\xi}_t^d$) and due to fluctuations in subjective capital gain expectations (β_t). Doing so requires setting the tax according to

$$\tau_t^{D*} = \frac{\beta(1-\delta)}{1 + \widehat{\xi}_t^d}\left[\frac{(\beta_t - \rho_\xi)}{1 - \beta(1-\delta)\rho_\xi}\widehat{\xi}_t^d + \frac{1}{1 - \beta(1-\delta)}(\beta_t - 1)\right]. \tag{2.65}$$

To understand what the preceding equation implies for the behavior of taxes, one has to take into account that the fluctuations in subjective beliefs (β_t) depend themselves on the tax: the tax influences housing prices, see equation (2.64), and thus – via capital gain extrapolation – the evolution of subjective beliefs.

To analyze the behavior of taxes, we consider the calibrated subjective belief model from Section 2.6 for the case where the average natural rate is equal to its post-1990 average (1.9%). We consider also intermediate forms of taxation that do not aim at fully eliminating the housing gap, by specifying taxes as

$$\tau_t^D = \lambda^D \tau_t^{D*},$$

where $\lambda^D \in [0, 1]$ is a sensitivity parameter. Our prior setup assumed $\lambda^D = 0$, while fully eliminating the housing price gap using macroprudential policy requires setting $\lambda^D = 1$. We then simulate the dynamics of housing prices, beliefs and taxes for alternative values of λ^D.

Table 2.6 reports the main outcomes. It shows that a higher tax sensitivity (λ^D) steadily

[56]To keep the rest of the model unchanged, the household also needs to expect lump sum tax rebates that are equal to the amount of subjectively expected tax payments.

Table 2.6: Taxes and housing price gap fluctuations for alternative tax sensitivities λ^D

Tax sensitivity λ^D	Housing Price Gap $\widehat{q}_t^u - \widehat{q}_t^{u*}$	Housing Taxes τ_t^D		
Value	Std. dev.	Std. dev.	Maximum	Minimum
0.0	14.2%	0.0%	0.0%	0.0%
0.2	9.8%	2.4%	7.0%	-12.1%
0.4	6.4%	4.2%	13.8%	-21.7%
0.6	3.7%	5.7%	18.0%	-30.0%
0.8	1.7%	7.0%	21.3%	-36.2%
1.0	0.0%	8.0%	23.9%	-41.8%

Notes: The table reports the standard deviation of the housing gap, $\widehat{q}_t^u - \widehat{q}_t^{u*}$, as well as the standard deviation, minimum value and maximum value of the macroprudential tax τ^D, for different tax sensitivities λ^D.

reduces the standard deviation of the housing price gap (second column). However, the standard deviation of taxes has to steadily increase. For a policy that fully eliminates the housing price gap ($\lambda^D = 1$), the standard deviation of taxes is a staggering 8% of the rental value of housing. Taxes reach maximum values up to 24% and minimum values deeply in negative territory, with subsidies above 40% of the rental value. These taxes fully stabilize the housing price gap but still induce substantial variation in subjective beliefs. The latter explains why taxes have to remain rather volatile. Intermediate policies, say ones that set $\lambda^D = 0.4$, substantially reduce the volatility of the housing gap, but still require rather volatile taxes and often very large subsidies.

Given the outcomes in Table 2.6, we conclude that the currently available macroprudential instruments will unlikely be able to insulate the monetary authority from disturbances in the housing sector arising from capital gain extrapolation.

2.9 Conclusion

This paper documents systematic deviations from rational housing price expectations and constructs a structural equilibrium model that jointly replicates the behavior of housing prices and the patterns of deviations from rational expectations. The model shows that subjective housing price beliefs significantly contribute to housing price fluctuations and that lower natural rates of interest generate increased volatility for housing prices and the natural rate.

Optimal monetary policy responds to falling and more volatile natural rates by implementing higher average rates of inflation. Monetary policy should also lean against housing price fluctuations induced by housing demand shocks, with reactions to housing price increases being more forceful than the reaction to housing price downturns. None of these features is optimal if households hold rational housing price expectations. This highlights the importance of basing policy advice on economic models featuring empirically plausible specifications for household beliefs.

Appendix B

Appendix to Chapter 2

B.1 Additional Results for Section 2.2

B.1.1 Five-Year-Ahead Capital Gain Expectations

While for our baseline results in Section 2.2 we focus on short-term housing price expectations, our findings equally hold for medium-term five-year-ahead expectations. We estimate the five-year analogue of regression (2.1) as follows:

$$q_{t+20} - E_t^{\mathcal{P}}\left[q_{t+20}\right] = a^{CG} + b^{CG} \cdot \left(E_t^{\mathcal{P}}\left[q_{t+20}\right] - E_{t-1}^{\mathcal{P}}\left[q_{t+19}\right]\right) + \varepsilon_t. \tag{B.1}$$

Table B.1 reports the estimates of b^{CG} showing that five-year expectations are updated sluggishly.

Table B.1: Sluggish adjustment of five-year-ahead housing price expectations

	Mean Expectations	Median Expectations
$\widehat{b^{CG}}$	6.95***	6.89***
	(1.703)	(1.680)

Notes: This table reports the empirical estimates of regression (B.1) using nominal housing-price expectations. The reported standard errors are robust with respect to heteroskedasticity and serial correlation (Newey-West with four lags). Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

We also run five-year-ahead versions of the regressions (2.2) and (2.3):

$$E_t^{\mathcal{P}}\left[\frac{q_{t+20}}{q_t}\right] = a + c \cdot PR_{t-1} + u_t \tag{B.2}$$

$$\frac{q_{t+20}}{q_t} = \mathbf{a} + \mathbf{c} \cdot PR_{t-1} + \mathbf{u}_t. \tag{B.3}$$

Table B.2 shows that five-year-ahead capital gain expectations covary positively with the price-to-rent ratio, whereas actual capital gains covary negatively.

Table B.2: Expected vs. actual capital gains using five-year-ahead housing price expectations

	$\hat{c}$ (in %)	$\hat{\mathbf{c}}$ (in %)	bias (in %) $-E(\hat{\mathbf{c}} - \hat{c})$	p-value $H_0 : c = \mathbf{c}$
Mean Expectations	0.045	-1.889	0.0159	0.000
	(0.0001)	(0.01997)		
Median Expectations	0.044	-1.889	0.0155	0.000
	(0.00024)	(0.01997)		

Notes: $\widehat{c}$ is the estimate of c in equation (B.2) and $\widehat{\mathbf{c}}$ the estimate of $\mathbf{c}$ in equation (B.3). The Stambaugh (1999) small sample bias correction is reported in the second-to-last column and the last column reports the p-values for the null hypothesis $c = \mathbf{c}$. Newey-West standard errors using four lags in parentheses.

B.1.2 IV Estimation of Sluggish Belief Updating

To insure that the results obtained from regression (2.1) in Section 2.2 are not driven by forecast revisions being correlated with the error term, we follow Coibion and Gorodnichenko (2015a) by adopting an Instrumental Variable approach. Specifically, we consider monetary policy shocks as an instrument for forecast revisions. We identify daily monetary policy shocks as changes of the current-month federal funds future in a 30-minute window around scheduled FOMC announcements (following the approach in Gürkaynak et al. (2005) and Gorodnichenko and Weber (2016)). We then aggregate shocks to quarterly frequency by assigning daily shocks partly to the current quarter and partly to the consecutive quarter, based on the number of remaining days in the current quarter. Table B.3 reports the results of the IV regression. The coefficients are positive and statistically significant, with point estimates that are even larger than the ones reported in Section 2.2.

Table B.3: Instrumental variable regression

	Mean Expectations	Median Expectations
Nominal Housing Prices		
$\widehat{b}^{CG}$	2.85^{**}	3.84^{***}
	(1.259)	(1.497)
First-stage F-statistic	21.88	17.78
Real Housing Prices		
$\widehat{b}^{CG}$	2.62^{***}	3.45^{***}
	(0.745)	(0.649)
First-stage F-statistic	44.49	34.13

Notes: $\widehat{b}^{CG}$ report the results from regression (2.1), instrumenting forecast revisions using monetary policy shocks, obtained via high-frequency identification. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

B.1.3 Sluggish Adjustment of Capital Gain Expectations

Regression (2.1) in Section 2.2 studies sluggish adjustment of expectations about the housing price level. Similar results can be obtained when considering expectations about capital gains. *Specification 1* in Table B.4 reports the regression coefficient when one replaces actual and expected housing price levels on the left-hand side of equation (2.1) with actual and expected capital gains. The coefficient estimates remain positive and highly statistically significant. *Specification 2* in Table B.4 reports results when replacing expectations about housing price levels with expectations about capital gains on the right-hand side of equation (2.1) and *Specification 3* reports results when replacing levels by (actual and expected) capital gains on both sides of equation (2.1). The coefficient estimates remain positive, but the significance levels are lower for Specifications 2 and 3.

Table B.4: Sluggish adjustment of housing price growth expectations

	Mean Expectations	Median Expectations
Specification 1		
Nominal Housing Prices		
$\widehat{b}^{CG}$	0.023***	0.030***
	(0.005)	(0.005)
Real Housing Prices		
$\widehat{b}^{CG}$	0.024***	0.031***
	(0.004)	(0.004)
Specification 2		
Nominal Housing Prices		
$\widehat{b}^{CG}$	492*	182
	(279)	(210)
Real Housing Prices		
$\widehat{b}^{CG}$	302*	158
	(164)	(168)
Specification 3		
Nominal Housing Prices		
$\widehat{b}^{CG}$	5.20*	2.16
	(2.896)	(2.06)
Real Housing Prices		
$\widehat{b}^{CG}$	3.23*	2.06
	(1.678)	(1.835)

Notes: This table shows the results of regression (2.1) in terms of house-price growth rates instead of house-price levels. *Specification 1* denotes the case in which we replace housing-price levels with capital gains on the left-hand side of regression (2.1), *Specification 2* the case in which we replace the right-hand side and *Specification 3* denotes the case in which we replace levels with capital gains on both sides of regression (2.1). The reported standard errors are robust with respect to heteroskedasticity and serial correlation (Newey-West with four lags). Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

B.1.4 Cyclicality of Housing Price Forecast Errors

A similar version of the test from Adam et al. (2017) presented in Section 2.2, which considers the cyclicality of expected gains, is proposed by Kohlhas and Walther (2021). In this case, we regress forecast errors about housing prices on the price-to-rent ratio. Formally, we estimate

$$\frac{q_{t+4}}{q_t} - E_t^{\mathcal{P}}\left[\frac{q_{t+4}}{q_t}\right] = \alpha + \gamma \cdot PR_{t-1} + \varepsilon_t. \tag{B.4}$$

Table B.5 shows the results. We find a negative and statistically significant coefficient in all cases. Thus, consumers tend to become too optimistic (pessimistic) when they observe high (low) housing valuations, inconsistent with rational expectations.

Table B.5: Forecast errors and price-to-rent ratios

	Mean Expectations	Median Expectations
Nominal Housing Prices		
$\widehat{\gamma}$	-0.5^{***}	-0.5^{***}
	(0.09)	(0.10)
Real Housing Prices		
$\widehat{\gamma}$	-0.5^{***}	-0.5^{***}
	(0.08)	(0.10)

Notes: This table shows the results of regression (B.4), whereas the estimated regression coefficients (and standard errors) are multiplied by one hundred for better readability. The reported standard errors are robust with respect to heteroskedasticity and serial correlation (Newey-West with four lags). Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

B.1.5 Dynamics of forecast errors with median and nominal housing price expectations

Figure B.1 shows alternative specifications of the dynamic forecast error responses presented in Section 2.2. Panel (a) presents the response of forecast errors for nominal housing prices. Panel (b) shows the response of forecast errors for real housing prices (as in Section 2.2) but considering median expectations. The figure shows that these responses are very close to the baseline specification shown in Section 2.2.

Figure B.1: Dynamic Forecast error response to realized capital gains

(a) Nominal (Mean) Capital Gain Expectations

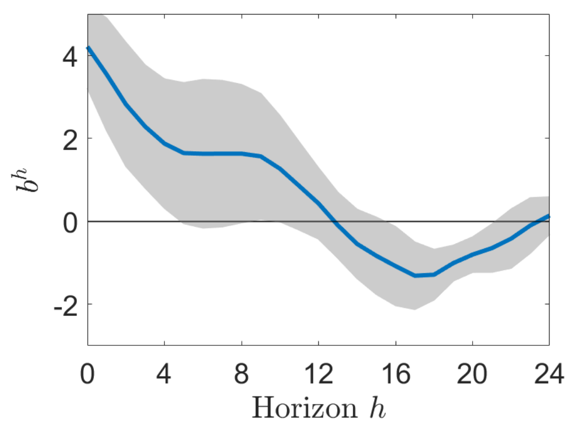

(b) Median (Real) Capital Gain Expectations

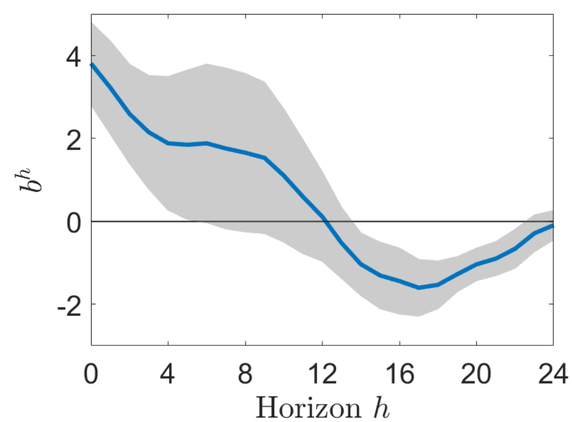

Notes: Panel (a) shows impulse-response functions of nominal capital gain forecast errors to a one standard deviation innovation in the housing capital gain. Panel (b) shows the impulse-response functions of median (real) capital gain forecast errors of one-year ahead expectations to a one standard deviation innovation in the housing capital gain. The shaded area shows the 90%-confidence intervals, standard errors are robust with respect to serial correlation and heteroskedasticity (Newey-West with $h+1$ lags).

B.1.6 Results when Excluding the Corona Virus Period

The empirical results reported in Section 2.2 are based on the entire period for which household-survey expectations are available, i.e., 2007-2021. This section reports results obtained when ending the sample in 2019, thereby excluding the recent Corona Virus crisis period. This is motivated by the fact that the two largest outliers in Figure 2.2 fall into the period after 2019. Tables B.6 and B.7 show, however, that our results are qualitatively and quantitatively robust to excluding observations from the years 2020 and 2021.

Table B.6: Sluggish adjustment of housing price expectations: excluding coronavirus crisis

	Mean Expectations	Median Expectations
Nominal Housing Prices		
$\widehat{b}^{CG}$	2.18***	2.80***
	(0.503)	(0.502)
Real Housing Prices		
$\widehat{b}^{CG}$	1.97***	2.43***
	(0.332)	(0.360)

Notes: This table shows the results of regression (2.1) excluding the coronavirus crisis, i.e., we exclude the years 2020 and 2021. The reported standard errors are robust with respect to heteroskedasticity and serial correlation (Newey-West with four lags). Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table B.7: Expected vs. actual capital gains: excluding coronavirus crisis

	$\hat{c}$ (in %)	$\hat{\mathbf{c}}$ (in %)	bias (in %) $-E(\hat{\mathbf{c}} - \hat{c})$	p-value $H_0 : c = \mathbf{c}$
Nominal Housing Prices				
Mean Expectations	0.058	-0.065	0.0036	0.000
	(0.0066)	(0.0126)		
Median Expectations	0.018	-0.065	0.0118	0.042
	(0.0010)	(0.0126)		
Real Housing Prices				
Mean Expectations	0.0614	-0.0483	-0.0009	0.000
	(0.0136)	(0.0090)		
Median Expectations	0.196	-0.483	0.076	0.017
	(0.0034)	(0.0090)		

Notes: This table shows the results of regressions (2.2) and (2.3) excluding the coronavirus period, i.e., we exclude the years 2020 and 2021. $\widehat{c}$ is the estimate of c in equation (2.2) and $\widehat{\mathbf{c}}$ the estimate of $\mathbf{c}$ in equation (2.3). The small sample bias correction is reported in the second to last column and the last column reports the p-values for the null hypothesis $c = \mathbf{c}$ in the fifth column. Newey-West standard errors using four lags in parentheses.

B.1.7 Regional Housing Prices and Expectations

This appendix considers regional variation in housing prices and housing price expectations. This is possible because the Michigan Survey reports the location of respondents using four different regions: West, North East, North Central (or Midwest) and South. While the Case-Shiller Price Index is not available at this level of regional disaggregation, we can construct a regional housing price index using the Case-Shiller Index that is available for twenty large U.S. cities. Following the definition of the regions in the Michigan Survey, we assign the twenty cities to the four regions and then aggregate city price indices to a regional index using two alternative approaches. The first approach weighs cities by population (as of 2019) within each region, while the second approach uses equal weights for all cities within a region.

Table B.8 lists all twenty cities, the region to which we allocate them and their regional population weights.[1] As in our baseline approach using aggregate data, we deflate nominal housing price indices by the aggregate CPI to obtain a real housing price index. We obtain real housing price expectations by deflating nominal (mean) expectations with region-specific (mean) inflation expectations.

Table B.8: Regions, cities and their weights

City	Region	Weight	City	Region	Weight
Denver	West	$\frac{0.705}{10.595}$	Chicago	North Central	$\frac{2.71}{4.189}$
Las Vegas	West	$\frac{0.634}{10.595}$	Cleveland	North Central	$\frac{0.385}{4.189}$
Los Angeles	West	$\frac{3.97}{10.595}$	Detroit	North Central	$\frac{0.674}{4.189}$
Phoenix	West	$\frac{1.633}{10.595}$	Minneapolis	North Central	$\frac{0.42}{4.189}$
Portland	West	$\frac{0.645}{10.595}$	Atlanta	South	$\frac{0.488}{4.209}$
San Diego	West	$\frac{1.41}{10.595}$	Charlotte	South	$\frac{0.857}{4.209}$
San Francisco	West	$\frac{0.874}{10.595}$	Dallas	South	$\frac{1.331}{4.209}$
Seattle	West	$\frac{0.724}{10.595}$	Miami	South	$\frac{0.454}{4.209}$
Boston	North East	$\frac{0.68}{9.1}$	Tampa	South	$\frac{0.387}{4.209}$
New York	North East	$\frac{8.42}{9.1}$	Washington DC	South	$\frac{0.692}{4.209}$

Notes: This table lists the twenty cities for which the Case-Shiller Home Price Index is available, the region to which the cities are allocated based on the Michigan Survey and their respective weights within region.

Table B.9 reports the region-specific estimates of b^{CG} in regression equation (2.1). All point estimates are positive with magnitudes that are broadly in line with the estimates at the national level. Furthermore, all regional estimates are significant at the 1% level. This shows that households update expectations sluggishly in all regions, consistent with the findings reported for the national level reported in the main text.

[1] The weights are calculated as the ratio of the population in the considered city, divided by the sum of populations in all cities in the respective region.

Table B.9: Sluggish adjustment of housing price expectations across regions

	Weighted	Unweighted
$\widehat{b}^{CG,W}$	2.00***	1.95***
	(0.411)	(0.374)
$\widehat{b}^{CG,NE}$	1.24***	1.15***
	(0.385)	(0.441)
$\widehat{b}^{CG,NC}$	1.97***	1.95***
	(0.461)	(0.459)
$\widehat{b}^{CG,S}$	1.74***	1.94***
	(0.385)	(0.393)

Notes: This table shows the results of regression (2.1) using regional housing prices and expectations. The superscripts W, NE, NC and S denote the regions West, North East, North Central (or Midwest) and South, respectively. The reported standard errors are robust with respect to heteroskedasticity and serial correlation (Newey-West with four lags). Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table B.10 reports the region-specific estimates of c and $\mathbf{c}$ from regressions (2.2) and (2.3). Since regional price-to-rent ratios are not available, the regression uses real housing prices on the right-hand side. In line with our findings at the aggregate level, we find $c > 0$ and $\mathbf{c} < \mathbf{0}$ in all the regions with the differences being largely highly statistically significant.

Figure B.2 shows the dynamic forecast errors responses to a one standard deviation innovation in the real housing capital gain in each of the four regions. Figure B.3 shows the results for the case in which the cities within regions are equally weighted. In line with the findings reported in the main text, households' housing capital gain expectations initially underreact but overshoot after some time.

Table B.10: Expected vs. actual capital gains across regions

	$\hat{c}$ (in %)	$\hat{\mathbf{c}}$ (in %)	bias (in %) $-E(\hat{\mathbf{c}} - \hat{c})$	p-value $H_0 : c = \mathbf{c}$
West				
Population-weighted	0.109	-0.216	0.090	0.083
	(0.0036)	(0.1360)		
Equally weighted	0.132	-0132	0.137	0.183
	(0.0034)	(0.1197)		
North Central				
Population-weighted	0.045	-0.544	0.008	0.000
	(0.0089)	(0.0256)		
Equally weighted	0.088	-0.458	0.0191	0.000
	(0.0118)	(0.0769)		
North East				
Population-weighted	0.013	-0.474	0.001	0.000
	(0.0089)	(0.0072)		
Equally weighted	0.126	-0.315	0.023	0.000
	(0.0187)	(0.0838)		
South				
Population-weighted	0.210	-0.008	0.137	0.144
	(0.0023)	(0.1067)		
Equally weighted	0.163	-0.238	0.055	0.014
	(0.0044)	(0.1250)		

Notes: This table shows the results of regressions (2.2) and (2.3) for different regions. $\widehat{c}$ is the estimate of c in equation (2.2) and $\widehat{\mathbf{c}}$ the estimate of $\mathbf{c}$ in equation (2.3). The small sample bias correction is reported in the second to last column and the last column reports the p-values for the null hypothesis $c = \mathbf{c}$ in the fifth column. Newey-West standard errors using four lags in parentheses.

Figure B.2: Regional dynamic forecast error responses to realized capital gains (population-weighted city housing price index)

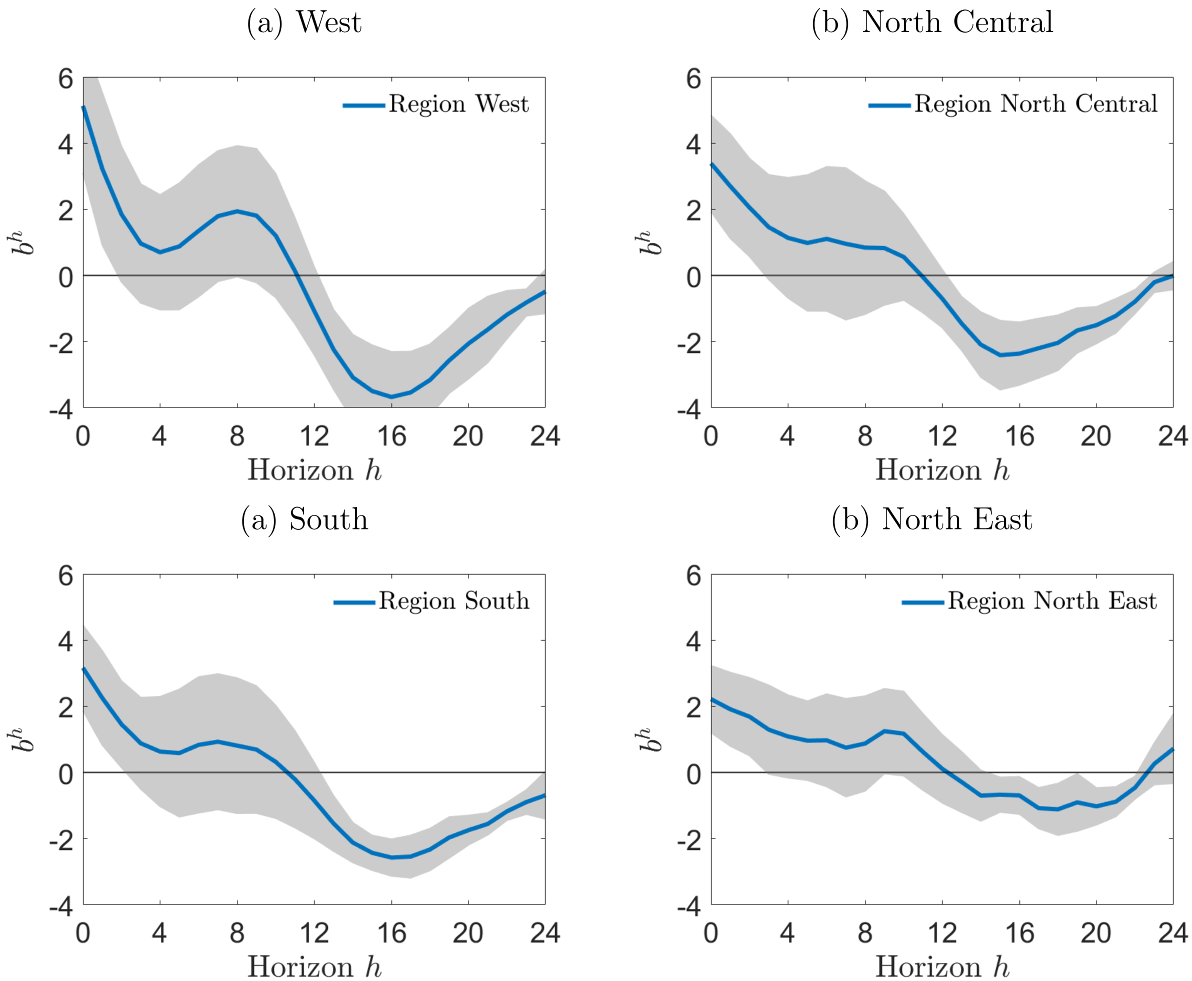

Notes: The figure shows the dynamic response of real capital gain forecast errors across the four different regions (in which cities' housing indices are weighted by their population share) to a one standard deviation innovation in the housing capital gain. The shaded area shows the 90%-confidence intervals, standard errors are robust with respect to serial correlation and heteroskedasticity (Newey-West with $h+1$ lags).

Figure B.3: Regional dynamic forecast error response to realized capital gains (equal weights)

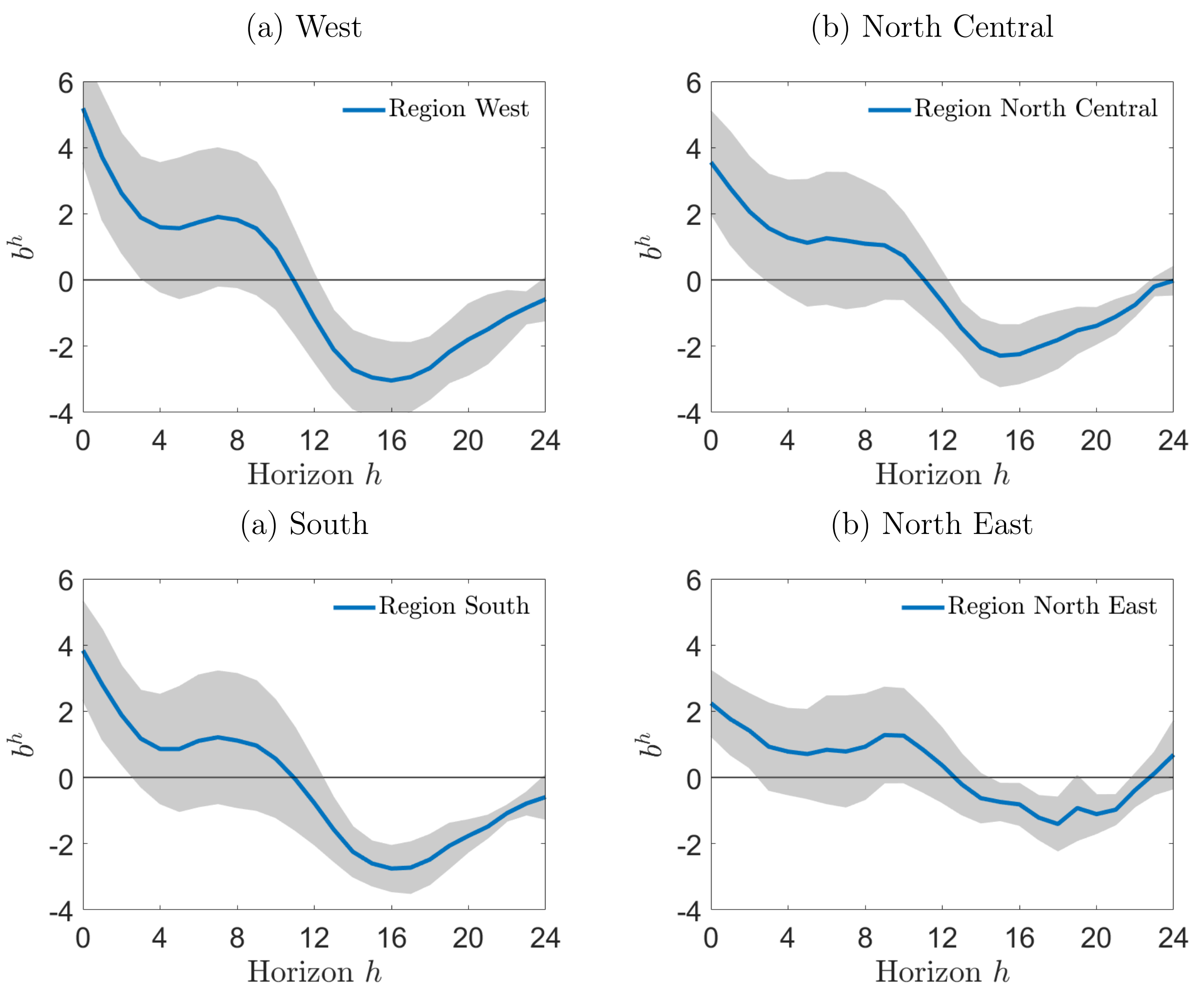

Notes: The figure shows the dynamic response of real capital gain forecast errors across the four different regions (in which cities are equally weighted) to a one standard deviation innovation in the housing capital gain. The shaded area shows the 90%-confidence intervals, standard errors are robust with respect to serial correlation and heteroskedasticity (Newey-West with $h+1$ lags).

B.2 Additional Results for Section 2.3

B.2.1 Dynamic Forecast Error Responses: Housing Price Level

Figure B.4 shows that the simple housing model also not only matches the empirical dynamic forecast error response about capital gains well, but also does a good job in matching the forecast errors about the level of future housing prices. The results are obtained by defining the forecast error X_{t+h} in equation (2.4) as

$$X_{t+h} \equiv q_{t+4+h} - E^{\mathcal{P}}_{t+h}\left[q_{t+4+h}\right] \tag{B.5}$$

and estimating the resulting local projections in the data and the population local projection for the model. Figure B.4 shows that households' expectations about the future level of housing prices initially undershoot and subsequently overshoot, as is the case with expected capital gains.

Figure B.4: Dynamic forecast error responses: housing price levels

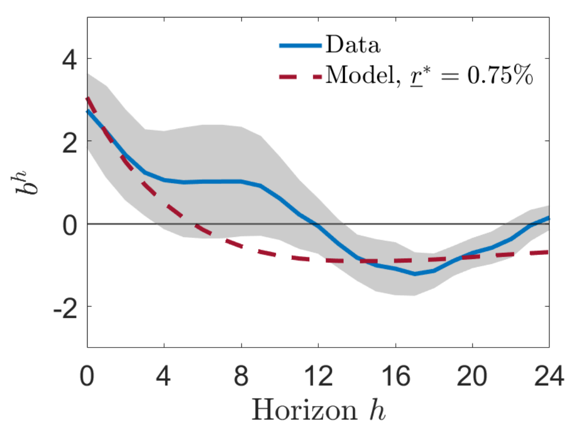

Notes: The figure shows impulse-response functions of housing-price level forecast errors of one-year ahead expectations to a one standard deviation innovation in the housing capital gain from the data and in the data. The shaded area shows the 90%-confidence intervals of the empirical estimates, standard errors are robust with respect to serial correlation and heteroskedasticity (Newey-West with $h + 1$ lags).

B.2.2 Model Response of the PR-Ratio to Housing Demand Shocks

Section 2.3 shows that real housing prices are more sensitive to housing demand shocks at lower levels of the natural rate. Figure B.5 illustrates that the same holds true for the model-implied price-to-rent ratio. The figure depicts the structural impulse response of the price-to-rent ratio (in percent deviations from steady state) to a one standard deviation housing-preference shock. It shows the response for a natural rate of 0.75% (blue line) and 0.25% (red line). The IRFs for the price-to-rent ration look very similar to the ones for real housing prices, shown in Figure 2.4(a).

Figure B.5: Impulse response functions

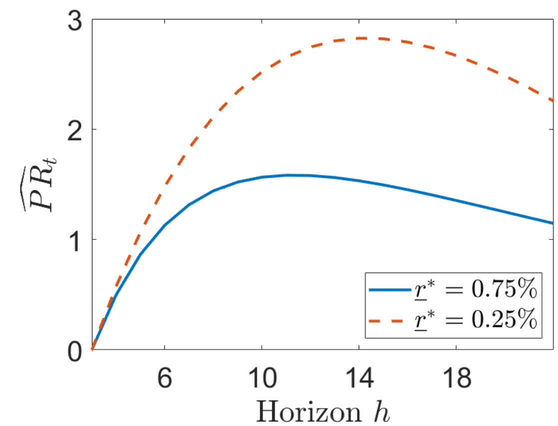

Notes: This figure shows the structural impulse response functions of the price-to-rent ratio (in percent deviations from steady state) to a one standard deviation housing-preference shock for different natural rates.

B.3 The Nonlinear Optimal Policy Problem

We shall consider Ramsey optimal policies in which the policymaker chooses the sequence of policy rates, prices, and allocations to maximize rationally expected household utility, subject to the constraint that prices and allocations constitute an Internally Rational Expectations Equilibrium. Note that the policymaker maximizes utility under a probability measure that is different from the one entertained by households, whenever the latter hold subjective beliefs. Benigno and Paciello (2014) refer to such a policymaker as being 'paternalistic'.

The objective of the policymaker is to maximize household utility. Using equation (2.20) to express the relative quantities demanded of the differentiated goods each period as a function of their relative prices and the linear dependence of utility on the stock of assets, we can write the utility flow to the representative household in the form

$$u(Y_t, q_t^u; \xi_t) - v(Y_t; \xi_t)\Delta_t + \bar{\xi}_t^d \frac{A_t^d}{\tilde{\alpha}} k_t^{\tilde{\alpha}},$$

with

$$\begin{aligned} u(Y_t, q_t^u; \xi_t) &\equiv \tilde{u}(C(Y_t, q_t^u, \xi_t); \xi_t) \\ v(y_t^j; \xi_t) &\equiv \tilde{v}(f^{-1}(y_t^j/A_t); \xi_t), \end{aligned}$$

where Δ_t, defined in equation (2.35), captures the misallocations from price dispersion. The term

$$\bar{\xi}_t^d \equiv \sum_{T=t}^{\infty} E_t[(1-\delta)^{T-t} \beta^{T-t} \xi_T^d]$$

captures the present value contribution from new housing investment. We can use (2.25) and (2.41) to express k_t in terms of Y_t, q_t^u and exogenous shocks. Hence, we can express the policy maker's objective of maximizing (2.14) under rational expectations, as maximizing

$$U = E_0 \sum_{t=0}^{\infty} \beta^t U(Y_t, \Delta_t, q_t^u; \xi_t),$$

where the flow utility is given by

$$\begin{aligned} U(Y_t, \Delta_t, q_t^u; \xi_t) \equiv & \frac{\bar{C}_t^{\tilde{\sigma}^{-1}} C(Y_t, q_t^u, \xi_t)^{1-\tilde{\sigma}^{-1}}}{1-\tilde{\sigma}^{-1}} \\ & - \frac{\lambda}{1+\nu} \bar{H}_t^{-\nu} \left(\frac{Y_t}{A_t}\right)^{1+\omega} \Delta_t \\ & + \frac{A_t^d \bar{\xi}_t^d}{\tilde{\alpha}} \Omega(q_t^u, \xi_t)^{\tilde{\alpha}} \, C(Y_t, q_t^u, \xi_t)^{\frac{\tilde{\alpha}}{1-\tilde{\alpha}} \tilde{\sigma}^{-1}}, \end{aligned} \tag{B.6}$$

which is a monotonically decreasing function of Δ given Y, q^u and ξ, and where $\Omega(q^u, \xi)$ is the function defined in (2.40). The only endogenous variables that are relevant for evaluating the policymaker's objective function are thus Y_t, Δ_t and q_t^u.

The non-linear optimal monetary policy problem is then given by

$$\max_{\{Y_t, q_t^u, p_t^*, w_t(j), P_t, \Delta_t, i_t \geq 0\}} E_0 \sum_{t=0}^{\infty} \beta^t U(Y_t, \Delta_t, q_t^u; \xi_t) \tag{B.7}$$

subject to

$$\left(\frac{p_t^*}{P_t}\right)^{1+\eta(\phi-1)} = \frac{E_t^{\mathcal{P}} \sum_{T=t}^{\infty} (\alpha)^{T-t} Q_{t,T} \frac{\eta}{\eta-1} \phi w_T(j) \left(\frac{Y_T}{A_T}\right)^{\phi} \left(\frac{P_T}{P_t}\right)^{\eta\phi+1}}{E_t^{\mathcal{P}} \sum_{T=t}^{\infty} (\alpha)^{T-t} Q_{t,T}(1-\tau_T) Y_T \left(\frac{P_T}{P_t}\right)^{\eta}} \tag{B.8}$$

$$w_t(j) = \lambda \frac{\bar{H}_t^{-\nu}}{\bar{C}_t^{\tilde{\sigma}^{-1}}} \left(\frac{Y_t}{A_t}\right)^{\phi\nu} C\left(Y_t, q_t^u, \xi_t\right)^{\tilde{\sigma}^{-1}} \left(\frac{p_t^*}{P_t}\right)^{-\eta\phi\nu} \tag{B.9}$$

$$(P_t/P_{t-1})^{\eta-1} = \frac{1-(1-\alpha)\left(\frac{p_t^*}{P_t}\right)^{1-\eta}}{\alpha} \tag{B.10}$$

$$\Delta_t = h(\Delta_{t-1}, P_t/P_{t-1}) \tag{B.11}$$

$$\tilde{u}_C(C(Y_t, q_t^u, \xi_t); \xi_t) = \lim_{T\to\infty} E_t^{\mathcal{P}} \left[\tilde{u}_C(C_T; \xi_T) \beta^T \prod_{k=0}^{T-t} \frac{1+i_{t+k}}{P_{t+k+1}/P_{t+k}}\right] \tag{B.12}$$

$$q_t^u = \xi_t^d + \beta(1-\delta) E_t^{\mathcal{P}} q_{t+1}^u, \tag{B.13}$$

where the initial price level P_{-1} and initial price dispersion Δ_{-1} are given. Equation (B.9) insures that wages clear current labor markets. Similarly, by setting $C_t = C(Y_t, q_t^u, \xi_t)$ on the left-hand side of the consumption Euler equation (B.12), we impose market clearing for output goods in period t. Similarly, setting q_t^u equal to the value defined in (B.13) insures market clearing in the housing market.[2] Firms' subjective expectations about future wages and households' subjectively optimal consumption plans for the future, however, will generally not be consistent with labor market or goods market clearing in the future in all subjectively perceived contingencies, when beliefs deviate from rational ones.

To be able to analyze the policy problem further, it is necessary to be more specific about the beliefs $\mathcal{P}$ entertained by households and firms.

B.4 Derivation of Equation (2.45)

Recall the definition of q_t^u which implies

$$\log q_t^u = \log q_t + \log \widetilde{u}_c(C_t; \xi_t)$$

[2]This holds as long as $D^{\max}$ is chosen sufficiently large, such that it never binds along the equilibrium path.

Under the considered belief setup in which agents learn about risk-adjusted capital gains, the dynamics of risk-adjusted capital gains and beliefs are independent of monetary policy. The response of $\log q_t^u$ to a unexpected change in the path of nominal rates $\boldsymbol{i}$ is thus $\frac{d \log q_t^u}{d\boldsymbol{i}} = 0$, so that

$$\begin{aligned}
\frac{d \log q_t}{d\boldsymbol{i}} &= -\frac{d \log \widetilde{u}_c(C_t; \xi_t)}{d\boldsymbol{i}} \\
&= -\frac{d \log \widetilde{u}_c(C_t; \xi_t)}{d \log C_t} \frac{d \log C_t}{d\boldsymbol{i}} \\
&= -\frac{\widetilde{u}_{cc}(C_t; \xi_t) C_t}{\widetilde{u}_c(C_t; \xi_t)} \frac{d \log C_t}{d\boldsymbol{i}} \\
&= \frac{1}{\widetilde{\sigma}} \frac{d \log C_t}{d\boldsymbol{i}} \qquad \text{(B.14)}
\end{aligned}$$

The optimal housing supply equation (2.25) can be written as

$$\log k_t = \frac{1}{1 - \tilde{\alpha}} \left(\log A_t^d + \log q_t \right).$$

Taking derivatives with respect to $\boldsymbol{i}$ in the previous equation and using (B.14) delivers (2.45).

B.5 Assumptions about Long-Run Beliefs

To insure that the subjectively optimal consumption plans satisfy the transversality condition (2.29), we impose that equation (2.7) describes subjective housing price beliefs for an arbitrarily long but finite amount of time $t < \bar{T} < \infty$ and that households hold rational expectations in the long-run, i.e. for all periods $t \geq \overline{T}$. Agents thus perceive

$$q_t^u = q_t^{u*} \quad \text{for all } t \geq \bar{T}, \mathcal{P} \text{ almost surely,}$$

where $q_t^{u*} = \bar{\xi}_t^d \equiv \sum_{T=t}^{\infty} E_t[(1-\delta)^{T-t} \beta^{T-t} \xi_T^d]$ is the rational expectations housing price. Appendix B.8.3 shows that this assumption is sufficient to insure that the transversality condition is satisfied. The transversality condition may also hold under weaker conditions, but actually showing this turns out to be difficult. The fact that agents will eventually hold rational housing and rental price expectations could be interpreted as agents learning to make rational predictions in the long-run.

B.6 Quadratic Approximation of the Policy Problem

This appendix derives the linear-quadratic approximation to the nonlinear policy problem in Appendix B.3.

B.6.1 Optimal Dynamics and the Housing Price Gap

It will be convenient to determine the welfare-maximizing level of output and the welfare-maximizing housing price under flexible prices, so as to express output and housing prices in terms of gaps relative to these maximizing values. We thus define (Y_t^*, q_t^{u*}) as the values (Y_t, q_t^u) that maximize $U(Y_t, 1, q_t^u; \xi_t)$, which are implicitly defined by[3]

$$U_Y(Y_t^*, 1, q_t^{u*}; \xi_t) = U_{q^u}(Y_t^*, 1, q_t^{u*}; \xi_t) = 0.$$

In particular, we have

$$q_t^{u*} = \overline{\xi}_t^d, \tag{B.15}$$

as shown in Appendix B.8.4. We have

$$\widehat{q}_t^{u,RE} = \widehat{q}_t^{u*}, \tag{B.16}$$

which shows that housing price fluctuations are indeed efficient under RE.

The output gap is defined as

$$y_t^{gap} \equiv \log(Y_t) - \log(Y_t^*) = \widehat{y}_t - \widehat{y}_t^*, \tag{B.17}$$

i.e. the log-difference of output from its dynamically optimal value.

Under subjective beliefs, it follows from equations (2.56) and the linearization of (B.15) (see Appendix B.8.1 below) that

$$\widehat{q}_t^{u,\mathcal{P}} - \widehat{q}_t^{u*} = \left(\frac{1-\beta(1-\delta)}{1-\beta(1-\delta)\beta_t} - \frac{1-\beta(1-\delta)}{1-\beta(1-\delta)\rho_\xi}\right)\widehat{\xi}_t^d + \frac{\beta(1-\delta)(\beta_t - 1)}{1-\beta(1-\delta)\beta_t}. \tag{B.18}$$

Again, for the case where $\beta_t = 1$ and with persistent housing demand shocks ($\rho_\xi \to 1$), the housing price gap under subjective beliefs is equal to the housing price gap under RE. Belief fluctuations, however, now contribute to fluctuations in the housing price gap.

For the real housing price gap, $\widehat{q}_t - \widehat{q}_t^*$, this implies

$$\widehat{q}_t - \widehat{q}_t^* = \left(1 + \widetilde{\sigma}^{-1} C_q\right)\left(\widehat{q}_t^u - \widehat{q}_t^{u*}\right) + \widetilde{\sigma}^{-1} C_Y y_t^{gap}. \tag{B.19}$$

[3] The optimal path for $\{Y_t^*, q_t^{u*}\}$ can then be used to determine optimal dynamics for the remaining variables. In particular, equation (2.41) determines C_t^*, equation (2.25) determines k_t^* and thus D_t^*, and equation (2.19) determines H_t^*.

B.6.2 Quadratically Approximated Welfare Objective

A second-order approximation to the utility function delivers

$$\frac{1}{2}U_{\widehat{Y}\widehat{Y}}\left(\widehat{y}_t-\widehat{y}_t^*\right)^2+\frac{1}{2}U_{\widehat{q}^u\widehat{q}^u}\left(\widehat{q}_t^u-\widehat{q}_t^{u*}\right)^2+\frac{1}{2}\underline{\gamma}^*h_{22}\pi_t^2+t.i.p.,$$

where $t.i.p.$ denotes terms independent of policy and $\underline{\gamma}^*$ is the Lagrange multiplier associated with equation (B.11) at the optimal steady state. See Appendix B.8.5 for a detailed derivation. The dependence of the objective function on inflation follows from a second-order approximation of the constraint (B.11), which allows expressing the second-order utility losses associated with price distortions Δ_t as a function of squared inflation terms.

Since the fluctuations in the housing price gap, $\widehat{q}_t^u-\widehat{q}_t^{u*}$, are either constant (with RE) or determined independently of policy (under subjective beliefs, see Equation (B.18)), the endogenous part of the loss function can be written as

$$\sum_{t=0}^{\infty}\beta^t\frac{1}{2}\left(\Lambda_\pi\pi_t^2+\Lambda_y\left(y_t^{gap}\right)^2\right).$$

B.6.3 New Keynesian Phillips Curve

Linearizing Equations (B.8)-(B.10) delivers the linearized Phillips curve. The condition for the equilibrium wage (B.9) in period T in industry j in which firms last updated their prices in period t is given by

$$w_T(j)=\tilde{w}_T(j)\left(\frac{p_t^j}{P_t}\right)^{-\eta\phi\nu}\left(\frac{P_T}{P_t}\right)^{\eta\phi\nu},$$

where

$$\tilde{w}_T(j)\equiv\lambda\frac{\bar{H}_T^{-\nu}}{\bar{C}_T^{\tilde{\sigma}^{-1}}}\left(\frac{Y_T}{A_T}\right)^{\phi\nu}C\left(Y_T,q_T^u,\xi_T\right)^{\tilde{\sigma}^{-1}}.$$

Since firms' expectations about $w_T(j)$ and P_T are rational, their expectations about $\tilde{w}_T(j)$ are rational as well. Using the expression for $w_T(j)$, noting that $p_t(i)=p_t^j=p_t^*$, and writing out $Q_{t,T}$, it follows that

$$\left(\frac{p_t^*}{P_t}\right)=\left(\frac{E_t^{\mathcal{P}}\sum_{T=t}^{\infty}\left(\alpha\beta\right)^{T-t}\frac{\eta}{\eta-1}\phi\bar{C}_T^{\tilde{\sigma}^{-1}}C_T^{-\tilde{\sigma}^{-1}}\tilde{w}_T(j)\left(\frac{Y_T}{A_T}\right)^{\phi}\left(\frac{P_T}{P_t}\right)^{\eta(1+\omega)}}{E_t^{\mathcal{P}}\sum_{T=t}^{\infty}\left(\alpha\beta\right)^{T-t}\bar{C}_T^{\tilde{\sigma}^{-1}}C_T^{-\tilde{\sigma}^{-1}}(1-\tau_T)Y_T\left(\frac{P_T}{P_t}\right)^{\eta-1}}\right)^{\frac{1}{1+\omega\eta}}. \tag{B.20}$$

Log-linearizing equation (B.20) delivers[4]

$$\widehat{p}_t^* - \widehat{P}_t = \frac{1-\alpha\beta}{1+\omega\eta}\left\{\widehat{\tilde{w}}_t(j) + \phi\left(\widehat{y}_t - \widehat{A}_t\right) - \widehat{\tau}_t - \widehat{y}_t + \alpha\beta E_t^{\mathcal{P}}\left[\frac{1+\omega\eta}{1-\alpha\beta}\left(\widehat{p}_{t+1}^* - \widehat{P}_{t+1} + \pi_{t+1}\right)\right]\right\}. \quad \text{(B.21)}$$

As the expectation in (B.21) is only about variables about which the private agents hold rational expectations, we can replace $E_t^{\mathcal{P}}[\cdot]$ with $E_t[\cdot]$.[5] Therefore, (B.10) can be used in period t and $t+1$, which in its linearized form is given by

$$\widehat{p}_t^* - \widehat{P}_t = \frac{\alpha}{1-\alpha}\pi_t.$$

Substituting $\widehat{\tilde{w}}_t(j)$ by the linearized version of the equilibrium condition (B.9) delivers the linearized New Keynesian Phillips Curve:

$$\pi_t = \kappa_y y_t^{gap} + \kappa_q\left(\widehat{q}_t^u - \widehat{q}_t^{u*}\right) + \beta E_t \pi_{t+1} + u_t, \quad \text{(B.22)}$$

where the coefficients κ are given by

$$\begin{aligned}
\kappa_y &= \frac{1-\alpha}{\alpha}\frac{1-\alpha\beta}{1+\omega\eta}\left(k_y - f_y\right) > 0 \\
\kappa_q &= -\frac{1-\alpha}{\alpha}\frac{1-\alpha\beta}{1+\omega\eta} f_q < 0,
\end{aligned}$$

with $k_y = \partial \log k / \partial \log y$, $f_y = \partial \log f / \partial \log y$, $f_q = \partial \log f / \partial \log q^u$, such that

$$\begin{aligned}
k_y - f_y &= \omega + \tilde{\sigma}^{-1}\frac{(1-g)\underline{Y}}{\underline{C} + \frac{\tilde{\sigma}^{-1}}{1-\tilde{\alpha}}\underline{k}} = \omega + \tilde{\sigma}^{-1} C_Y > 0 \\
f_q &= \tilde{\sigma}^{-1}\frac{\frac{\underline{k}}{1-\tilde{\alpha}}}{\underline{C} + \frac{\tilde{\sigma}^{-1}}{1-\tilde{\alpha}}\underline{k}} = -\tilde{\sigma}^{-1} C_q > 0,
\end{aligned}$$

where $C_q \equiv \frac{q^u}{C}\frac{\partial C}{\partial q^u}$ and $C_Y \equiv \frac{Y}{C}\frac{\partial C}{\partial Y}$, and where the functions $f\left(Y, q^u; \xi\right) \equiv (1-\tau)\,\bar{C}^{\tilde{\sigma}^{-1}} Y C\left(Y, q^u; \xi\right)^{-\tilde{\sigma}^{-1}}$ and $k\left(y; \xi\right) \equiv \frac{\eta}{\eta-1}\lambda\phi\frac{\bar{H}^{-\nu}}{A^{1+\omega}}Y^{1+\omega}$ are the same as in Adam and Woodford (2021), for the current period in which markets clear and the internally rational agents observe this.

---

[4]This follows from the the fact that in steady state, we have $p^* = P$, so that

$$\frac{\eta}{\eta-1}\phi\bar{C}^{\tilde{\sigma}^{-1}}C^{-\tilde{\sigma}^{-1}}\tilde{w}(j)\left(\frac{Y}{A}\right)^{\phi} = \bar{C}^{\tilde{\sigma}^{-1}}C^{-\tilde{\sigma}^{-1}}(1-\tau)Y.$$

The steady state value of the numerator in (B.20) is thus given by $\frac{1}{1-\alpha\beta}\frac{\eta}{\eta-1}\phi\bar{C}^{\tilde{\sigma}^{-1}}C^{-\tilde{\sigma}^{-1}}\tilde{w}(j)\left(\frac{Y}{A}\right)^{\phi}$ and the steady state value of the denominator by $\frac{1}{1-\alpha\beta}\bar{C}^{\tilde{\sigma}^{-1}}C^{-\tilde{\sigma}^{-1}}(1-\tau)Y$.

[5]The subjective consumption plans showing up in the stochastic discount factor drop out at this order of approximation.

The cost-push shock u_t is given by

$$u_t = \frac{(1-\alpha)(1-\alpha\beta)}{\alpha(1+\omega\eta)}\left(\Theta + \widehat{\tau}_t - \widehat{g}_t\right),$$

where

$$\widehat{\tau}_t = -\log\left(\frac{1-\tau_t}{1-\bar{\tau}_t}\right)$$
$$\widehat{g}_t = -\log\left(\frac{1-g_t}{1-\bar{g}_t}\right)$$

define deviations of τ_t and g_t from their second-best steady state values.

As in the standard New Keynesian model, a linearization of (B.11) implies that the state variable Δ_t is zero to first order under the maintained assumption that initial price dispersion satisfies $\Delta_{-1} \sim O(2)$. This constraint, together with the assumption that the Lagrange multipliers are of order $O(1)$, thus drops out of the quadratic formulation of the optimal policy problem. The second-order approximation of (B.11) is, however, important to express the quadratic approximation of utility in terms of inflation.

B.6.4 Linearized IS Equation with Potentially Non-Rational Housing Price Beliefs

We here linearize the constraint (B.12). One difficulty with this constraint is that it features the limiting expectations of the subjectively optimal consumption plan on the right hand side. Generally, this would require solving for the subjectively optimal consumption paths, which is generally difficult.

Under our beliefs specifications, housing prices beliefs are rational in the limit. This insures that we do not have to solve for the subjectively optimal consumption plan, instead can derive the IS equation directly in terms of the output gap.

We can now define the natural rate of interest:

The natural rate $r_t^{*,RE}$ is the one implied by the consumption Euler equation (2.24) or (B.12), rational expectations, and the welfare-maximizing consumption levels under flexible prices $\{C_t^*\}$. It satisfies

$$\tilde{u}_C(C_t^*;\xi_t) = \beta E_t\left[u_C(C_{t+1}^*;\xi_t)(1+r_{t+k}^{*,RE})\right]. \tag{B.23}$$

Using the previous definition, we obtain the linearized Euler equation under potentially subjective housing prices beliefs:

Lemma 3. *For the considered belief specifications, the log-linearized household optimality condition (B.12) implies for all t*

$$y_t^{gap} = \lim_T E_t y_T^{gap} - E_t \left(\sum_{k=0}^{\infty} \varphi \left(i_{t+k} - \pi_{t+1+k} - r_{t+k}^{*,RE} \right) \right) - \frac{C_q}{C_Y} \left(\widehat{q}_t^u - \widehat{q}_t^{u*} \right), \tag{B.24}$$

where $\lim_T E_t y_T^{gap}$ *is the (rational) long-run expectation of the output gap, and* $\varphi \equiv -\frac{\tilde{u}_c}{\tilde{u}_{cc} C} \frac{1}{C_Y} > 0$*. The coefficients* $C_q < 0$ *and* $C_Y > 0$ *are the ones defined in the derivation of the linearized Phillips Curve.*

The proof can be found in Appendix B.8.2

B.6.5 Lagrangian Formulation of the Approximated Ramsey Problem

Collecting results from the previous sections, we obtain the following Lagrangian formulation of the Ramsey problem

$$\max_{\{\pi_t, y_t^{gap}, i_t \geq \underline{i}\}} \min_{\{\varphi_t, \lambda_t\}} \tag{B.25}$$

$$E_0 \sum_{t=0}^{\infty} \beta^t \Bigg\{ -\frac{1}{2} \left(\Lambda_\pi \pi_t^2 + \Lambda_y \left(y_t^{gap} \right)^2 \right)$$

$$+ \varphi_t \left[\pi_t - \kappa_y y_t^{gap} - \kappa_q \left(\widehat{q}_t^u - \widehat{q}_t^{u*} \right) - u_t - \beta E_t \pi_{t+1} \right] \tag{B.26}$$

$$+ \lambda_t \left[y_t^{gap} - \lim_T E_t y_T^{gap} + \varphi E_t \sum_{k=0}^{\infty} \left(i_{t+k} - \pi_{t+1+k} - r_{t+k}^{*,RE} \right) + \frac{C_q}{C_Y} \left(\widehat{q}_t^u - \widehat{q}_t^{u*} \right) \right] \tag{B.27}$$

$$- \varphi_{-1} \pi_0 - \lambda_{-1} \left(\varphi \pi_0 - y_0^{gap} - \frac{C_q}{C_Y} \left(\widehat{q}_0^u - \widehat{q}_0^{u*} \right) \right) \Bigg\},$$

where the process for $(\widehat{q}_t^u - \widehat{q}_t^{u*})$ can be treated as exogenous for the purpose of monetary policy and where the initial Lagrange multipliers $(\varphi_{-1}, \lambda_{-1})$ capture initial pre-commitments. In order to numerically solve the optimal policy problem in (B.25), we recursify the problem as proposed in Marcet and Marimon (2019) and solve for the associated value functions and optimal policies. Details of the recursive formulation can be found in Appendix B.8.6.

B.7 The Volatility of PR-Ratio and of the Natural Rate

Figure B.6 shows the evolution of natural rates of interest and price-to-rent ratios the U.S., Canada, France, Germany, and the United Kingdom, which we use in Section 2.2. The natural rates are estimated by Holston et al. (2017) and Fujiwara et al. (2016). The price-to-rent ratios are taken from the OECD. We convert the quarterly series of natural rates to annual series by taking arithmetic averages and the quarterly series or PR-ratios to annual series by taking harmonic averages.

Figure B.6: Natural rates and price-to-rent ratios

(a) Natural Rates of Interest

(b) Price-to-rent ratios

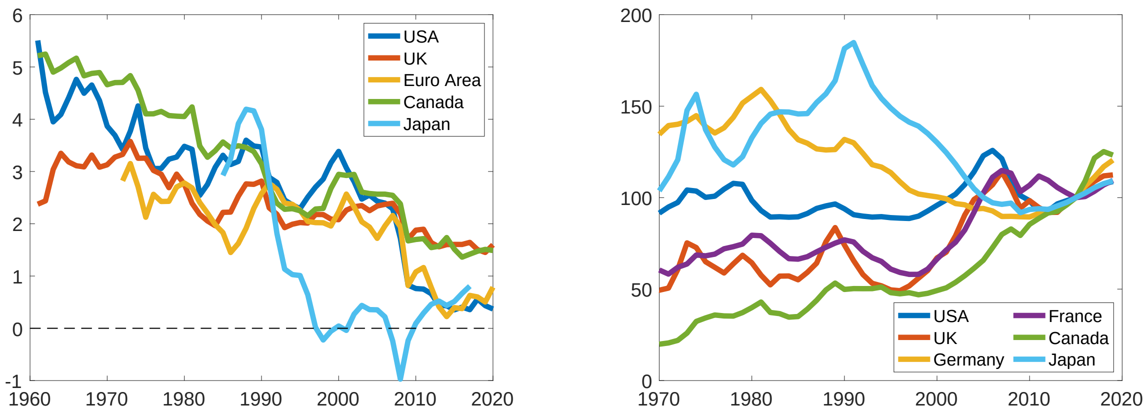

Notes: This figure shows the evolution of the natural rate of interest (left panel) and price-to-rent ratios (right panel) for different advanced economies over the period 1961-2020 and 1970-2019, respectively.

Figures 2.4(b) and 2.6 in the Section 2.2 document that the fall in the level of the natural rates of interest across several advanced economies was accompanied by an increase in the volatility of the price-to-rent ratio and in the volatility of natural rates. These trends are consistent with the subjective belief model, outlined in Sections 2.3 and 2.4.

Figure B.7 plots the volatility of the price-to-rent ratio (left panel) and the standard deviation of the natural rate (right panel), respectively before 1990 (blue bars) and after 1990 (red bars), along with 90% confidence bands. The reported volatilities of the price-to-rent ratios are the standard deviations relative to the period-specific mean values, in line with the model. The reported volatilities of the natural rates of interest are the standard deviations of the fluctuations around a linear time trend, in order to isolate high-frequency volatility that can be related to natural rate fluctuations in the model around a fixed steady state value of the natural rate. Figure B.9 shows the volatility price-to-rent ratio using the same linear detrending approach. The p-values below the respective bars are for the null hypothesis of no change in the volatility. The increase in the volatility of the PR ratio and the natural rate were statistically significant in most of the advanced economies. The evidence is not always statistically significant due to the high serial correlation of the price-to-rent ratio and the natural rate, which makes it difficult to estimate standard deviations precisely. Figure B.8 shows that the reported volatility increases are not driven by the exact point where we split the data, instead looks often similar for other split points.

Figure B.7: Volatility of the PR ratio and natural rates pre and post 1990.

(a) Standard Deviation of the Price-to-Rent Ratio Pre and Post 1990

(b) Standard Deviation of Natural Rate Pre and Post 1990

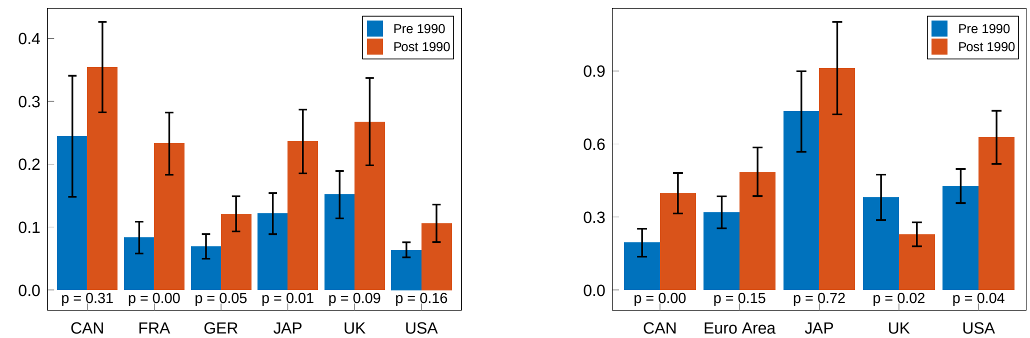

Notes: The black lines denote the 90%-confidence bands. The p-value corresponds to the test whether or not the values changed from pre to post 1990. The reported volatilities of the price-to-rent ratios are the standard deviations relative to the period-specific mean values. The reported volatilities of the natural rates of interest are the standard deviations of the fluctuations around a linear time trend.

Figure B.8: Robustness of housing and natural rate volatility increases with different sample splits

(a) Volatility of the Price-to-Rent ratios

(b) Volatility of natural rates

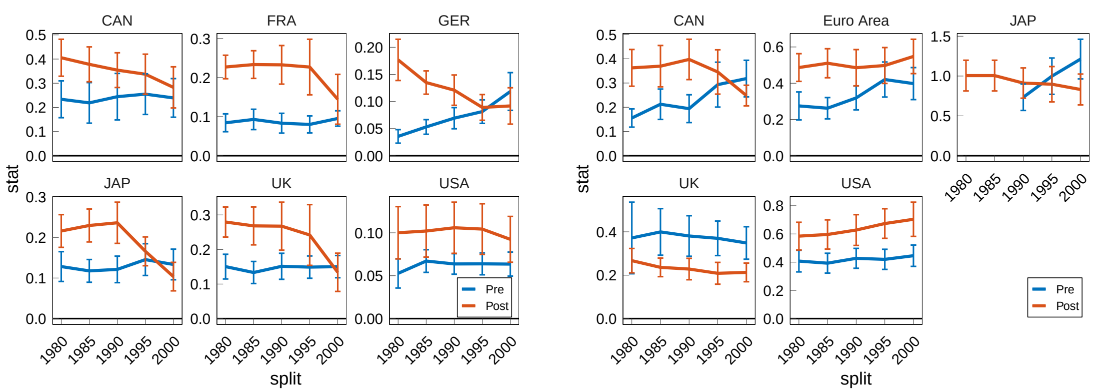

Notes: Panel (a) shows the standard deviation of the price-to-rent ratio, and panel (b) shows the standard deviation of the natural rate for different advanced economies, computed for varied subsamples. The blue lines show the estimates for the pre-period, and the red lines for the post-period, when the sample is split at the year marked on the horizontal axis. The whiskers denote 90%-confidence bands.

Figure B.9: Standard deviation of the detrended PR ratio pre and post 1990

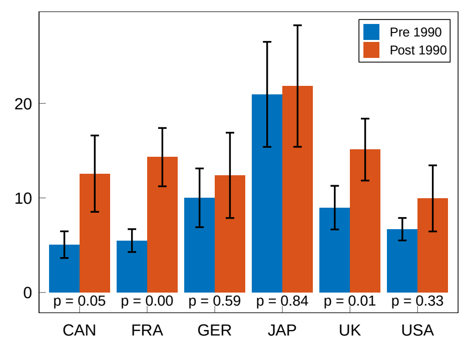

Note: The black lines denote the 90%-confidence bands. The p-value corresponds to the test whether or not the values changed from pre to post 1990.

B.8 Proofs

B.8.1 Results in Section 2.5

Proof of Results in Section 2.5.1. Result (2.50) follows from iterating forward on (2.27). Log linearizing (2.50), we have

$$\widehat{q}_t^u = \widehat{\overline{\xi}}_t^d,$$

and log-linearizing (2.12) delivers

$$\widehat{\xi}_t^d = \rho_\xi \widehat{\xi}_{t-1}^d + \varepsilon_t^d.$$

Since the steady-state value of $\overline{\underline{\xi}}^d$ is $\overline{\underline{\xi}}^d = \frac{\underline{\xi}^d}{1-\beta(1-\delta)}$, the log-linearization of $\overline{\xi}_t^d$ delivers

$$\begin{aligned}
\widehat{\overline{\xi}}_t^d &= (1-\beta(1-\delta))\left[\widehat{\xi}_t^d + \beta(1-\delta)E_t\widehat{\xi}_{t+1}^d + ...\right] \\
&= (1-\beta(1-\delta))\left[\widehat{\xi}_t^d + \beta(1-\delta)\rho_\xi\widehat{\xi}_t^d + ...\right] \\
&= (1-\beta(1-\delta))\sum_{T=t}^{\infty}(\beta(1-\delta)\rho_\xi)^{T-t}\,\widehat{\xi}_t^d \\
&= \widehat{\xi}_t^d\frac{1-\beta(1-\delta)}{1-\beta(1-\delta)\rho_\xi}.
\end{aligned}$$

The results for the price-to rent ration follow by noticing that equation (2.26) implies

$$PR_t \equiv \frac{q_t}{R_t} = \frac{q_t^u}{\xi_t^d}. \tag{B.28}$$

■

Proof of Results in Section 2.5.2. From equation (2.27), which has to hold with equality in equilibrium, and equation (2.8) we get

$$q_t^{u,\mathcal{P}} = \frac{1}{1-\beta(1-\delta)\beta_t}\xi_t^d$$

The percent deviation of housing prices from the steady state, in which $\beta_t = 1$ and $\xi_t^d = \underline{\xi}^d$, is then given by

$$\begin{aligned}
\widehat{q}_t^{u,\mathcal{P}} &= \frac{\frac{1}{1-\beta(1-\delta)\beta_t}\xi_t^d - \frac{1}{1-\beta(1-\delta)}\underline{\xi}^d}{\frac{1}{1-\beta(1-\delta)}\underline{\xi}^d} \\
&= \frac{1-\beta(1-\delta)}{1-\beta(1-\delta)\beta_t}\frac{\xi_t^d}{\underline{\xi}^d} - 1 \\
&= \frac{1-\beta(1-\delta)}{1-\beta(1-\delta)\beta_t}\left(1+\widehat{\xi}_t^d\right) - 1
\end{aligned}$$

$$= \quad \frac{1-\beta(1-\delta)}{1-\beta(1-\delta)\beta_t}\widehat{\xi}_t^d + \frac{\beta(1-\delta)(\beta_t-1)}{1-\beta(1-\delta)\beta_t} \tag{B.29}$$

Note, that we can decompose the housing price under subjective beliefs into the housing price under RE and terms that are driven by beliefs:

$$\widehat{q}_t^{u,\mathcal{P}} = \widehat{q}_t^{u,RE} + \frac{\beta(1-\delta)\left(\beta_t-1\right)}{1-\beta(1-\delta)\beta_t} + \frac{\left(1-\beta(1-\delta)\right)\left(\beta(1-\delta)\left(\beta_t-\rho_\xi\right)\right)}{\left(1-\beta(1-\delta)\beta_t\right)\left(1-\beta(1-\delta)\rho_\xi\right)}\widehat{\xi}_t^d. \tag{B.30}$$

Note, that

$$E_t^{\mathcal{P}}\left[q_{t+1}^{u,\mathcal{P}}\right] = \beta_t q_t^{u,\mathcal{P}}.$$

Therefore, a log-linear approximation around the optimal steady state, in which $\underline{\beta} = 1$, yields

$$E_t^{\mathcal{P}}\left[\widehat{q}_{t+1}^{u,\mathcal{P}}\right] = \widehat{q}_t^{u,\mathcal{P}} + \left(\beta_t - 1\right).$$

From this, we can add and subtract on the right-hand side

$$E_t\left[\widehat{q}_{t+1}^{u,RE}\right] = \rho_\xi \widehat{\xi}_t^d \frac{1-\beta(1-\delta)}{1-\beta(1-\delta)\rho_\xi},$$

which, after plugging in the expression from (B.29), delivers

$$\begin{aligned} E_t^{\mathcal{P}}\left[\widehat{q}_{t+1}^{u,\mathcal{P}}\right] &= E_t\left[\widehat{q}_{t+1}^{u,RE}\right] + \left(\beta_t-1\right)\left[1+\frac{\beta(1-\delta)}{1-\beta(1-\delta)\beta_t}\right] \\ &\quad + \left(1-\beta(1-\delta)\rho_\xi - \left(1-\beta(1-\delta)\beta_t\right)\rho_\xi\right)\frac{\left(1-\beta(1-\delta)\right)}{\left(1-\beta(1-\delta)\beta_t\right)\left(1-\beta(1-\delta)\rho_\xi\right)}\widehat{\xi}_t^d. \end{aligned}$$

In the limit $\rho_\xi \to 1$, this boils down to

$$E_t^{\mathcal{P}}\left[\widehat{q}_{t+1}^{u,\mathcal{P}}\right] = E_t\left[\widehat{q}_{t+1}^{u,RE}\right] + \left(\beta_t-1\right)\left[1+\frac{\beta(1-\delta)}{1-\beta(1-\delta)\beta_t}\left(1+\widehat{\xi}_t^d\right)\right]$$

Log-linearizing equation (B.28), which holds true independent of the belief specification, yields

$$\widehat{PR}_t^{\mathcal{P}} = \widehat{q}_t^{u,\mathcal{P}} - \widehat{\xi}_t^d.$$

■

Proof of Lemma 2. Under the proposed policy that sets $i_t - E_t\pi_{t+1}$ equal to the natural rate defined in equation (2.59), we have

$$
\begin{aligned}
y_t^{gap} &= \lim_T E_t y_T^{gap} - E_t\left(\sum_{k=0}^{\infty}\varphi\left(i_{t+k} - \pi_{t+1+k} - r_{t+k}^{*,RE}\right)\right) - \frac{C_q}{C_Y}\left(\widehat{q}_t^u - \widehat{q}_t^{u*}\right) \\
&= \lim_T E_t y_T^{gap} - E_t\left(\sum_{k=0}^{\infty}\varphi\left(r_{t+k}^{*,RE} - \frac{1}{\varphi}\frac{C_q}{C_Y}\left(\left(\widehat{q}_{t+k}^u - \widehat{q}_{t+k}^{u*}\right) - E_{t+k}\left(\widehat{q}_{t+k+1}^u - \widehat{q}_{t+k+1}^{u*}\right)\right) - r_{t+k}^{*,RE}\right)\right) \\
&\quad -\frac{C_q}{C_Y}\left(\widehat{q}_t^u - \widehat{q}_t^{u*}\right) \\
&= \lim_T E_t y_T^{gap} + E_t\left(\sum_{k=0}^{\infty}\left(\frac{C_q}{C_Y}\left(\left(\widehat{q}_{t+k}^u - \widehat{q}_{t+k}^{u*}\right) - \left(\widehat{q}_{t+k+1}^u - \widehat{q}_{t+k+1}^{u*}\right)\right)\right)\right) - \frac{C_q}{C_Y}\left(\widehat{q}_t^u - \widehat{q}_t^{u*}\right) \\
&= \lim_T E_t y_T^{gap} + E_t\left(\frac{C_q}{C_Y}\left(\left(\widehat{q}_t^u - \widehat{q}_t^{u*}\right) - \lim_k E_t\left(\widehat{q}_{t+k+1}^u - \widehat{q}_{t+k+1}^{u*}\right)\right)\right) - \frac{C_q}{C_Y}\left(\widehat{q}_t^u - \widehat{q}_t^{u*}\right) \\
&= \lim_T E_t y_T^{gap} + \left(\frac{C_q}{C_Y}\left(\widehat{q}_t^u - \widehat{q}_t^{u*}\right)\right) - \frac{C_q}{C_Y}\left(\widehat{q}_t^u - \widehat{q}_t^{u*}\right) \\
&= \lim_T E_t y_T^{gap},
\end{aligned}
$$

which proves that with this policy, the output gap is indeed constant, and $r^{*,\mathcal{P}}$ is the real rate that implies a constant output gap. ∎

B.8.2 Log-linearized Euler equation

Proof of Lemma 3. Log-linearizing equation (B.12) around the optimal steady state delivers

$$
\tilde{u}_{CC}C\widehat{c}_t + \tilde{u}_{C\xi}\underline{\xi}\widehat{\xi}_t = E_t^{\mathcal{P}}\sum_{k=0}^{\infty}\tilde{u}_C\left(i_{t+k} - \pi_{t+1+k}\right) + \lim_{T\to\infty}E_t^{\mathcal{P}}\left(\tilde{u}_{CC}C\widehat{c}_T + \tilde{u}_{C\xi}\underline{\xi}\widehat{\xi}_T\right),
$$

and log-linearizing (B.23) gives

$$
\tilde{u}_{CC}C\widehat{c}_t^* + \tilde{u}_{C\xi}\underline{\xi}\widehat{\xi}_t = E_t\sum_{k=0}^{\infty}\tilde{u}_C r_{t+k}^{*,RE} + \lim_{T\to\infty}E_t\left(\tilde{u}_{CC}C\widehat{c}_T^* + \tilde{u}_{C\xi}\underline{\xi}\widehat{\xi}_T\right).
$$

Subtracting the previous equation from (B.31) delivers

$$
\widehat{c}_t - \widehat{c}_t^* = E_t^{\mathcal{P}}\sum_{k=0}^{\infty}\frac{\tilde{u}_C}{\tilde{u}_{CC}C}\left(i_{t+k} - \pi_{t+1+k} - r_{t+k}^{*,RE}\right) + \lim_{T\to\infty}E_t^{\mathcal{P}}\left(\widehat{c}_{T+1} - \widehat{c}_{T+1}^*\right), \tag{B.31}
$$

where we used $E_t^{\mathcal{P}}\xi_T = E_t\xi_T$ and $E_t^{\mathcal{P}}\widehat{c}_{T+1}^* = E_t\widehat{c}_{T+1}^*$, which hold because agents hold rational expectations about fundamentals.

In all periods in which the subjectively optimal plan is consistent with market clearing in the goods sector, the plan satisfies equation (2.41). Log-linearizing equation (2.41) delivers

$$\widehat{c}_t = C_Y \widehat{y}_t + C_q \widehat{q}_t^u + C_\xi \widehat{\xi}_t, \tag{B.32}$$

where $\widehat{\xi}_t$ is a vector of exogenous disturbances (involving $A_t^d, \bar{C}_t, g_t$). Evaluating this equation at the optimal dynamics defines the optimal consumption gap $\widehat{c}_t^*$:

$$\widehat{c}_t^* \equiv C_Y \widehat{y}_t^* + C_q \widehat{q}_t^{u*} + C_\xi \widehat{\xi}_t.$$

Subtracting the previous equation from (B.32) delivers

$$\begin{aligned} \widehat{c}_t - \widehat{c}_t^* &= C_Y \left(\widehat{y}_t - \widehat{y}_t^*\right) + C_q \left(\widehat{q}_t^u - \widehat{q}_t^{u*}\right) \\ &= C_Y y_t^{gap} + C_q \left(\widehat{q}_t^u - \widehat{q}_t^{u*}\right) \end{aligned} \tag{B.33}$$

Since the current consumption market in period t clears, equation (B.33) holds in period t and can be used to substitute the consumption gap on the left-hand side of equation (B.31). Similarly, since housing price expectations are rational in the limit, the consumption market also clears in the limit under the subjectively optimal plans, i.e., equation (2.41) holds for $t \geq T'$. We can thus use equation (B.33) also to substitute the consumption gap on the r.h.s. of equation (B.31). Using the fact that housing price expectations are rational in the limit ($\lim_T E_t^P \left(\widehat{q}_t^u - \widehat{q}_t^{u*}\right) = 0$), we obtain

$$y_t^{gap} = \lim_T E_t^{\mathcal{P}} y_T^{gap} - E_t \left(\sum_{k=0}^{\infty} \varphi \left(i_{t+k} - \pi_{t+1+k} - r_{t+k}^{*,RE} \right) \right) - \frac{C_q}{C_Y} \left(\widehat{q}_t^u - \widehat{q}_t^{u*}\right).$$

Since we assumed that agents' beliefs about profits and taxes are given by equations (2.37) and (2.38), respectively, evaluated using rational income expectations, the household holds rational expectations about total income. This can be seen by substituting (2.37) and (2.38) into the budget constraint (2.15). We thus have $\lim_T E_t^{\mathcal{P}} y_T^{gap} = \lim_T E_t y_T^{gap}$ in the previous equation, which delivers (B.24). ■

B.8.3 Transversality Condition Satisfied with Subjective Housing Price Beliefs

This appendix shows that under the considered subjective belief specifications, the optimal plans satisfy the transversality constraint (2.29). Since $D_t \in [0, D^{\max}]$ and $E_t^{\mathcal{P}} q_T^u = E_t \overline{\xi}_T^d$ for $T \geq T'$, we have $\lim_{T\to\infty} \beta^T E_t^{\mathcal{P}} \left(D_T q_T^u\right) = 0$. We thus only need to show that $\lim_{T\to\infty} \beta^T E_t^{\mathcal{P}} \frac{\bar{C}_T^{\tilde{\sigma}^{-1}}}{C_T^{\tilde{\sigma}^{-1}}} B_T = 0$.

Combining the budget constraint (2.15) with (2.37) and (2.38) we obtain

$$C_t + B_t + \left(D_t - (1-\delta)D_{t-1} - \tilde{d}(k_t;\xi_t)\right) q_t^u \frac{C_t^{\tilde{\sigma}^{-1}}}{\bar{C}_t^{\tilde{\sigma}^{-1}}} + k_t = (1-g_t)\, Y_t + B_{t-1}.$$

For $t \geq T'$ the subjectively optimal plans satisfy market clearing in the housing market, i.e.,

$$D_t - (1-\delta)D_{t-1} - \tilde{d}(k_t;\xi_t) = 0$$

so that the budget constraint implies

$$C_t + B_t + k_t = (1-g_t)\, Y_t + B_{t-1}. \tag{B.34}$$

Furthermore, for $t \geq T'$ subjectively optimal plans also satisfy market clearing for consumption goods, i.e.,

$$C_t + k_t = (1-g_t)\, Y_t.$$

It thus follows that the subjectively optimal debt level B_t in the budget constraint (B.34) is constant under the subjectively optimal plan, after period $t \geq T'$. Furthermore, the expectations about Y_t in the budget constraint (B.34) is rational under the assumed lump sum transfer expectations, so that the household's subjective consumption expectations are the same as in a rational expectations equilibrium. (The subjectively optimal investment decisions k_t are driven by rational housing price expectations). Since the limit expectations $\bar{C}_T^{\tilde{\sigma}^{-1}}/C_T^{\tilde{\sigma}^{-1}}$ are bounded in the rational expectations equilibrium, it follows that $\lim_{T\to\infty} \beta^T E_t^{\mathcal{P}} \frac{\bar{C}_T^{\tilde{\sigma}^{-1}}}{C_T^{\tilde{\sigma}^{-1}}} B_T = 0$.

B.8.4 Optimal House Price Absent Price Rigidities

The following derivation closely follows Adam and Woodford (2021). We obtain $U_{q^u}\left(Y_t, \Delta_t, q_t^u, \xi_t\right)$ from differentiating equation (B.6) with respect to q_t^u and set it equal to 0:

$$\begin{aligned} U_{q^u}\left(Y_t, \Delta_t, q_t^u, \xi_t\right) = {} & \bar{C}_t^{\tilde{\sigma}^{-1}} C_{q^u}\left(Y_t, q_t^u, \xi_t\right) C\left(Y_t, q_t^u, \xi_t\right)^{-\tilde{\sigma}^{-1}} \\ & + A_t^d \bar{\xi}_t^d \frac{\partial \Omega\left(q_t^u, \xi_t\right)}{\partial q_t^u} \Omega\left(q_t^u, \xi_t\right)^{\tilde{\alpha}-1} C\left(Y_t, q_t^u, \xi_t\right)^{\frac{\tilde{\alpha}}{1-\tilde{\alpha}}\tilde{\sigma}^{-1}} \\ & + \frac{\tilde{\sigma}}{1-\tilde{\alpha}} A_t^d \bar{\xi}_t^d \Omega\left(q_t^u, \xi_t\right)^{\tilde{\alpha}} C\left(Y_t, q_t^u, \xi_t\right)^{\frac{\tilde{\alpha}}{1-\tilde{\alpha}}\tilde{\sigma}^{-1}-1} C_{q^u}\left(Y_t, q_t^u, \xi_t\right) = 0, \end{aligned}$$

where

$$\frac{\partial \Omega\left(q_t^u, \xi_t\right)}{\partial q_t^u} = \frac{1}{q_t^u} \frac{1}{1-\tilde{\alpha}} \Omega\left(q_t^u, \xi_t\right),$$

and when defining $\chi \equiv \frac{\tilde{\sigma}^{-1}}{1-\tilde{\alpha}} - 1$, we get

$$C_{q^u}(Y_t, q_t^u; \xi_t) \equiv \frac{\partial C(Y_t, q_t^u; \xi_t)}{\partial q^u} = \frac{-\frac{1}{q_t^u}\frac{1}{1-\tilde{\alpha}}\Omega(q_t^u, \xi_t) C(Y_t, q_t^u, \xi_t)^{\chi+1}}{1 + (1+\chi)\,\Omega(q_t^u, \xi_t) C(Y_t, q_t^u, \xi_t)^{\chi}}.$$

Taking everything together, we get

$$U_{q^u}\left(Y_t, \Delta_t, q_t^u, \xi_t\right) = \frac{\frac{1}{q_t^u}\frac{1}{1-\tilde{\alpha}}\Omega(q_t^u, \xi_t) C(Y_t, q_t^u, \xi_t)^{\chi+1}}{1 + (1+\chi)\,\Omega(q_t^u, \xi_t) C(Y_t, q_t^u, \xi_t)^{\chi}} \bar{C}_t^{\tilde{\sigma}-1} \left(\frac{\overline{\xi}_t^d}{q_t^u} - 1 \right).$$

In order for U_{q^u} to be zero, we need to have that

$$q_t^{u*} = \overline{\xi}_t^d,$$

as stated in equation (B.15).

B.8.5 Quadratically Approximated Welfare Objective

This derivation follows Adam and Woodford (2021). In the optimal steady state, we have $U_Y = U_{q^u} = U_{Yq^u} = 0$, as well as $U_\Delta + \underline{\gamma}\,(\beta h_1 - 1) = 0$. Given the assumption $\Delta_{-1} \sim O(2)$, it follows $\Delta_t \sim O(2)$ for all $t \geq 0$. Additionally, we have $h_2 \equiv \frac{\partial h(\Delta,\Pi)}{\partial \Pi} = 0$ at the optimal steady state. Therefore, a second-order approximation of the contribution of the variables $(Y_t, \Delta_t, q_t^u, \Pi_t, \xi_t)$ to the utility of the household yields

$$\frac{1}{2} U_{\widehat{Y}\widehat{Y}}\left(\widehat{y}_t - \widehat{y}_t^*\right) + \frac{1}{2} U_{\widehat{q}^u\widehat{q}^u}\left(\widehat{q}_t^u - \widehat{q}_t^{u*}\right) + \frac{1}{2}\underline{\gamma}^* h_{22}\pi_t^2 + t.i.p.,$$

where $t.i.p.$ contains all terms independent of policy. Under rational expectations, we have that $(\widehat{q}_t^u - \widehat{q}_t^{u*}) = 0$ and is thus constant and independent of (monetary) policy. Under subjective beliefs, $(\widehat{q}_t^u - \widehat{q}_t^{u*})$ is purely driven by beliefs β_t and housing demand shocks ξ_t^d, see equation (B.18), both independent of policy. Therefore, we include $\frac{1}{2} U_{\widehat{q}^u\widehat{q}^u}\left(\widehat{q}_t^u - \widehat{q}_t^{u*}\right)$ in $t.i.p.$.

The term $U_{\widehat{Y}\widehat{Y}}$ is given by $U_{\widehat{Y}\widehat{Y}} \equiv Y \frac{\partial}{\partial Y}\left(U_{\widehat{Y}}\right) \equiv Y\frac{\partial}{\partial Y}\left(Y U_Y\right) = \underline{Y}^* U_Y + \left(\underline{Y}^*\right)^2 U_{YY}$. At the optimal steady state, we have

$$\begin{aligned} \Lambda_\pi &= -\frac{1}{2}\underline{\gamma}^* h_{22} > 0 \\ \Lambda_y &= -\frac{1}{2}\left(\underline{Y}^*\right)^2 U_{YY} > 0, \end{aligned}$$

where

$$U_{YY} = -\tilde{\sigma}^{-1}\left(1 - \underline{g}\right) \underline{\bar{C}}^{\tilde{\sigma}^{-1}} C\left(\underline{Y}, \underline{q}^u, \underline{\xi}\right)^{-\tilde{\sigma}^{-1}-1} C_Y \frac{\underline{Y}^*}{C\left(\underline{Y}, \underline{q}^u, \underline{\xi}\right)}$$

$$
\begin{aligned}
&- \frac{\lambda}{1+\nu}\left(1+\omega\right)\omega \frac{\bar{\underline{H}}^{-\nu}}{\underline{A}^{1+\omega}}\underline{Y}^{\omega-1} < 0 \\
h_{22} &= \frac{\alpha\eta\left(1+\omega\right)\left(1+\omega\eta\right)}{1-\alpha} > 0 \\
\underline{\gamma}^{*} &= \frac{U_{\Delta}}{1-\alpha\beta} < 0,
\end{aligned}
$$

with

$$
U_{\Delta} = -\frac{\underline{Y}^{*}\left(1-\underline{g}\right)}{1+\omega}\left(\frac{\bar{C}^{\tilde{\sigma}^{-1}}}{C\left(\underline{Y}^{*}, \underline{q}^{u*}, \underline{\xi}\right)}\right)^{\tilde{\sigma}^{-1}} < 0.
$$

B.8.6 Recursified Optimal Policy Problem with Lower Bound

We numerically solve the quadratically approximated optimal policy problem with forward-looking constraints (B.25). While it would be preferable to solve the fully nonlinear Ramsey problem, as spelled out in Appendix B.3, this is computationally not feasible with sufficient degree of numerical accuracy because the problem features 9 state variables and an occasionally binding constraint. The quadratically approximated problem features 2 state variables less because price dispersion Δ_t is to first order independent of policy and because the Phillips curve reduces from a system involving two forward-looking infinite sums, see equation (B.20), to a system involving only a single infinite sum, see (B.21).

Eggertsson and Singh (2019) compare the exact solution of the standard New Keynesian model with lower bound to the solution of the linear-quadratic approximation with lower bound and show that the quantitative deviations are modest, even for extreme shocks of the size capturing the 2008 recession in the U.S..

To obtain a recursive problem, we apply the approach of Marcet and Marimon (2019) to the problem with forward-looking constraints (B.25). We thereby assume that the Lagrangian defined by problem (B.25) satisfies the usual duality properties that allow interchanging the order of maximization and minimization, which we verify ex-post using the computed value function. We set the terminal value function for $t = T'$ to its RE value function $W^{RE}(\cdot)$. For $t \leq T'$ we have a value function $W_t(\cdot)$ satisfying the following recursion:

$$
\begin{aligned}
& W_t(\varphi_{t-1}, \mu_{t-1}, u_t, r_t^{*,RE}, \beta_t, \xi_t^d, q_{t-1}^u) \\
= \; & \max_{\left(\pi_t, y_t^{gap}, i_t \geq \underline{i}\right)} \min_{(\varphi_t, \lambda_t)} -\frac{1}{2}\left(\Lambda_\pi \pi_t^2 + \Lambda_y \left(y_t^{gap}\right)^2\right) \\
& + \left(\varphi_t - \varphi_{t-1}\right)\pi_t - \varphi_t\left(\kappa_y y_t^{gap} + \kappa_q\left(\widehat{q}_t^u - \widehat{q}_t^{u*}\right) + u_t\right) \\
& + \lambda_t\left[y_t^{gap} - \lim_T E_t y_T^{gap} + \varphi\left(i_t - E_t \sum_{k=0}^{\infty} r_{t+k}^{*,RE}\right) + \frac{C_q}{C_Y}\left(\widehat{q}_t^u - \widehat{q}_t^{u*}\right)\right] \\
& + \mu_{t-1}\varphi\left(i_t - \pi_t\right) + \gamma_t\left(i_t - \underline{i}\right)
\end{aligned}
$$

$$+\beta E_t \left[W_{t+1}(\varphi_t, \underbrace{\beta^{-1}(\lambda_t + \mu_{t-1})}_{=\mu_t}, u_{t+1}, r_{t+1}^{*,RE}, \beta_{t+1}, \xi_{t+1}^d, q_t^u \right] \tag{B.35}$$

where the next period state variables (β_{t+1},q_t^u) are determined by equations (2.9) and (2.54) and ($\widehat{q}_t^u - \widehat{q}_t^{u*}$) is determined by equation (B.18). Here we assume that $r_t^{*,RE}$ follows a Markov process, such that the term $E_t \sum_{k=0}^{\infty} r_{t+k}^{*,RE}$ showing up in the current-period return can be expressed as a function of the current state $r_t^{*,RE}$. The future state variables ($\varphi_t, \mu_t, \beta_{t+1}, q_t^u$) are predetermined in period t. The expectation about the continuation value is thus only over the exogenous states ($u_{t+1}, r_{t+1}^{*,RE}, \xi_{t+1}^d$). The endogenous state variable φ_{t-1} is simply the lagged Lagrange multiplier on the New Keynesian Phillips curve with housing. The endogenous state variable μ_{t-1} is given for all $t \geq 0$ by

$$\mu_t = \beta^{-(t+1)}(\lambda_0 + \mu_{-1}) + \beta^{-t}\lambda_1 + ... + \beta^{-1}\lambda_t.$$

The initial values (φ_{-1}, μ_{-1}) are given at time zero and equal to zero in the case of time-zero-optimal monetary policy.

For periods $t < T'$, where T' is the period from which housing price expectations are rational and the lower bound constraint ceases to bind, the value functions depend on time, thereafter they are time-invariant. Likewise for sufficiently large T', the value functions $W_t(\cdot)$ and $W_{t+1}(\cdot)$ will become very similar.

We can numerically solve for the value function $W_t(\cdot)$ by value function iteration, starting with $W_{T'}$ which is the value function associated with the linear-quadratic problem with RE.

B.8.7 Optimal Targeting Rule

Differentiating (B.35) with respect to $\{\pi_t, y_t^{gap}, i_t\}$ yields:

$$\begin{aligned}
\frac{\partial W_t}{\partial \pi_t} &= -\Lambda_\pi \pi_t + (\varphi_t - \varphi_{t-1}) - \mu_{t-1}\varphi = 0 \\
\frac{\partial W_t}{\partial y_t^{gap}} &= -\Lambda_y y_t^{gap} - \varphi_t \kappa_y + \lambda_t = 0 \\
\frac{\partial W_t}{\partial i_t} &= \gamma_t + \lambda_t \varphi + \mu_{t-1}\varphi = 0 \text{ and } \gamma_t (i_t - \underline{i}) = 0.
\end{aligned}$$

Combining these first-order conditions, we can derive the following targeting rule which characterizes optimal monetary policy

$$\Lambda_\pi \pi_t + \frac{\Lambda_y}{\kappa_y}\left(y_t^{gap} - y_{t-1}^{gap}\right) + \frac{\lambda_{t-1}}{\kappa_y} + \mu_{t-1}\left(\varphi + \frac{1}{\kappa_y}\right) + \frac{\gamma_t}{\varphi \kappa_y} = 0,$$

where γ_t is the Lagrange multiplier associated with the lower bound on interest rates. If the lower bound on the nominal interest rate does not bind in the current period, we have $\gamma_t = 0$.

Furthermore, if the lower bound has not been binding up to period t, the IS equation has not posed a constraint for the monetary policymaker. Thus, $\lambda_{t-1} = \lambda_{t-k} = 0$ for all $k = 0, 1, ..., t$. For an initial value of $\mu_{-1} = 0$, it follows that $\mu_{t-1} = 0$. The targeting rule then collapses to

$$\Lambda_\pi \pi_t + \frac{\Lambda_y}{\kappa_y}\left(y_t^{gap} - y_{t-1}^{gap}\right) = 0,$$

which is the same as in Clarida et al. (1999).

The Lagrange multiplier $\gamma_t \leq 0$ captures the cost of a currently binding lower bound. If $\gamma_t < 0$, the optimal policy requires a compensation in the form of a positive output gap or inflation. The multipliers λ_{t-1} and μ_{t-1} capture promises from past commitments when the lower bound was binding.

Another way to express equation (B.36) is to write it as

$$\Lambda_\pi \pi_t + \frac{\Lambda_y}{\kappa_y}\left(y_t^{gap} - y_{t-1}^{gap}\right) + \frac{1}{\varphi\kappa_y}\left[\gamma_t - \frac{1+\beta+\varphi\kappa_y}{\beta}\gamma_{t-1} + \frac{\gamma_{t-2}}{\beta}\right] = 0. \tag{B.36}$$

House prices do not enter the optimal target criterion directly but larger fluctuations in house prices make the lower bound bind more often and for a longer period of time. The optimal policy, thus, requires larger compensations in terms of positive output gaps and inflation. To implement this, the nominal interest rate needs to be kept longer at the lower bound.

B.8.8 Calibration of C_q/C_Y

To calibrate C_q/C_Y, the ratio of the consumption elasticities to housing prices and income, respectively, note that from appendix "Second-Order Conditions for Optimal Allocation" in Adam and Woodford (2021), we have

$$C_{q^u}(Y_t, q_t^u; \xi_t) \equiv \frac{\partial C(Y_t, q_t^u; \xi_t)}{\partial q^u} = \frac{-\frac{1}{q_t^u}\frac{1}{1-\tilde{\alpha}}\Omega(q_t^u, \xi_t)C(Y_t, q_t^u, \xi_t)^{\chi+1}}{1 + (1+\chi)\,\Omega(q_t^u, \xi_t)C(Y_t, q_t^u, \xi_t)^{\chi}}$$

where $\chi \equiv \frac{\tilde{\sigma}^{-1}}{1-\tilde{\alpha}} - 1$. In our formulation, we have defined

$$C_q \equiv \frac{\partial C(Y_t, q_t^u; \xi_t)}{\partial \ln q_t^u} = \frac{\partial C(Y_t, q_t^u; \xi_t)}{\partial q_t^u}\frac{\partial q_t^u}{\partial \ln q_t^u} = C_{q^u}(Y_t, q_t^u; \xi_t)\frac{q_t^u}{C_t}$$

so that we have

$$C_q = -\frac{\frac{1}{1-\tilde{\alpha}}\Omega(q_t^u, \xi_t)C(Y_t, q_t^u, \xi_t)^{\chi+1}}{C(Y_t, q_t^u, \xi_t) + (1+\chi)\,\Omega(q_t^u, \xi_t)C(Y_t, q_t^u, \xi_t)^{\chi+1}}.$$

From the appendix in Adam and Woodford (2021) we also have

$$C_Y(Y_t, q_t^u, \xi_t) \equiv \frac{\partial C_Y(Y_t, q_t^u, \xi_t)}{\partial Y_t} = \frac{1 - g_t}{1 + \Omega(q_t^u, \xi_t)\,(1+\chi)\,C(Y_t, q_t^u, \xi_t)^{\chi}}$$

so that in our notation

$$C_Y \equiv \frac{\partial C_Y(Y_t, q_t^u, \xi_t)}{\partial \ln Y_t} = \frac{(1-g_t)\,Y_t}{C(Y_t, q_t^u, \xi_t) + \Omega(q_t^u, \xi_t)\,(1+\chi)\,C(Y_t, q_t^u, \xi_t)^{\chi+1}}.$$

We then have

$$\frac{C_q}{C_Y} = \frac{\frac{-\frac{1}{1-\widetilde{\alpha}}\Omega(q_t^u,\xi_t)C(Y_t,q_t^u,\xi_t)^{\chi+1}}{C(Y_t,q_t^u,\xi_t)+(1+\chi)\Omega(q_t^u,\xi_t)C(Y_t,q_t^u,\xi_t)^{\chi+1}}}{\frac{(1-g_t)Y_t}{C(Y_t,q_t^u,\xi_t)+\Omega(q_t^u,\xi_t)(1+\chi)C(Y_t,q_t^u,\xi_t)^{\chi+1}}} = -\frac{1}{1-\widetilde{\alpha}}\frac{\Omega(q_t^u,\xi_t)C(Y_t,q_t^u,\xi_t)^{\chi+1}}{(1-g_t)\,Y_t}.$$

In the steady state, we have $\overline{Y}(1-\overline{g}) = \overline{C} + \overline{\Omega C}^{\chi+1}$, which says that privately consumed output $\overline{Y}(1-\overline{g})$ is divided up into consumption $\overline{C}$ and resources invested in the housing sector, $\overline{\Omega C}^{1+\chi}$. We thus have that

$$\frac{\overline{\Omega C}^{\chi+1}}{\overline{Y}(1-\overline{g})} = 1 - \frac{\overline{C}}{\overline{Y}(1-\overline{g})} = 1 - \frac{\overline{C}}{\overline{C} + \overline{\Omega C}^{\chi+1}} = 1 - \frac{1}{1+\overline{\Omega C}^{\chi}}.$$

Following Adam and Woodford (2021), we set this to the share of housing investment to total consumption, $\overline{\Omega C}^{\chi}$, equal to 6.3%, so that in steady state we have

$$\frac{C_q}{C_Y} = -\frac{1}{1-\widetilde{\alpha}}\left(1 - \frac{1}{1.063}.\right)$$

Finally, following Adam and Woodford (2021), we set the long-run elasticity of housing supply equal to five, which implies $\widetilde{\alpha} = 0.8$, so that

$$\frac{C_q}{C_Y} = -5\left(1 - \frac{1}{1.063}\right) \approx -0.29633.$$

From this, it follows that

$$C_Y = \frac{(1-g)Y}{C + (1+\chi)\,\Omega C^{\chi+1}} = \frac{C+k}{C + \frac{\widetilde{\sigma}^{-1}}{1-\tilde{\alpha}}k} = \frac{1+\frac{k}{C}}{1+\frac{\widetilde{\sigma}^{-1}}{1-\tilde{\alpha}}\frac{k}{C}} = \frac{1+0.063}{1+5\cdot 0.063} = 0.80836$$

and

$$C_q = -0.29633 \cdot 0.80836 = -0.23954.$$

Chapter 3

Inflation – who cares? Monetary Policy in Times of Low Attention

3.1 Introduction

Managing inflation expectations is an important instrument for monetary policy. It offers a powerful substitute to conventional tools if the nominal interest rate is constrained by the lower bound. By making promises about the future conduct of its policy, the monetary authority can shape inflation expectations, and thus, steer the real interest rate even when the nominal rate is constrained. At least this is how it works in theory.[1] But while traditional analyses assume that agents are perfectly informed and have rational expectations, recent empirical evidence suggests that the general public is usually poorly informed about and inattentive to monetary policy and inflation.[2] What do these low levels of attention imply for the conduct of monetary policy? And how has the public's attention to inflation changed over the last fifty years?

To tackle these questions, I propose an approach to quantify *attention to inflation* in the data. This approach is based on a model of optimal attention choice subject to information acquisition costs. The result is a law of motion for inflation expectations in which attention governs how strongly agents update their expectations following an inflation surprise. The optimal degree of attention depends positively on how volatile and persistent inflation is.

Using micro survey data of professional forecasters and consumers in the U.S., I then use this approach to estimate attention to inflation in the data and show that attention was very low in the years just before the Covid-19 crisis. In the 1970s and 1980s, on the other hand, attention to

[1] For optimal monetary policy under full-information rational expectations, see e.g., Clarida et al. (1999), Woodford (2003) or Galí (2015). See Eggertsson and Woodford (2003), Adam and Billi (2006b) or Coibion et al. (2012) for early analyses including the lower bound on nominal interest rates.

[2] Candia et al. (2021) and Coibion et al. (2022a) show that U.S. firms as well as households are usually poorly informed about and quite inattentive to monetary policy. Coibion and Gorodnichenko (2012, 2015a), for example, show that models of limited attention more closely align with empirical patterns of inflation expectations, compared to models with full-information rational expectations. In line with limited attention, D'Acunto et al. (2022) show that forward guidance is quite ineffective in stimulating inflation expectations.

inflation was substantially higher. Consistent with the underlying model, times of higher inflation volatility and persistence are characterized by higher attention to inflation.

How does this decline in attention matter for the conduct of optimal monetary policy? To answer this question, I solve for the Ramsey optimal monetary policy in a standard New Keynesian model augmented with an effective lower bound (ELB) on the nominal interest rate and with inflation expectations that are characterized by limited attention. Both of these ingredients matter greatly for optimal monetary policy and the normative implications of declining attention. Lower attention has a stabilizing effect on inflation expectations and actual inflation, resembling more anchored (short-run) expectations. If we do not account for the ELB, lower attention is thus welfare improving (similar to the findings in Paciello and Wiederholt (2014)). Accounting for the occasionally-binding ELB, however, completely overturns the normative implications of lower attention, as lower levels of attention now lead to a decline in welfare. The reason is that lower attention renders managing expectations more difficult, which is particularly relevant if the nominal interest rate is constrained by the lower bound.[3]

To better understand these results, I first show that under sub-optimal policy, i.e., if monetary policy follows an *ad-hoc* Taylor rule, limited attention can lead to substantially longer periods at the ELB. Even though lower levels of attention attenuate the initial response of inflation expectations to a given shock, the decline in expectations becomes more persistent which hinders actual inflation from recovering quickly. Due to the persistently-low inflation, the monetary authority keeps the interest rate at the ELB for longer. I refer to these periods of long spells at the ELB and persistent declines in inflation and inflation expectations as *inflation-attention traps*. The response of the output gap, on the other hand, is very similar to the one under rational expectations. Thus, low attention offers a potential explanation for why inflation was relatively stable during the Great Recession but was persistently low during the subsequent recovery, seemingly disconnected from output (as documented in Del Negro et al. (2020)).

To deal with these attention traps and the overall inability to manage inflation expectations at low levels of attention, I find that it is *optimal* to induce a higher average inflation rate. The higher inflation rate provides additional space when cutting the interest rate following adverse shocks, thus, mitigates the drawbacks of lower attention, and therefore helps to prevent long spells at the ELB.[4] Given the estimates of attention just before the Covid-19 crisis, the average inflation rate under Ramsey optimal policy is about 2-3 percentage points higher than under rational expectations. This increase in the level of inflation, however, is costly from a welfare

[3]Coibion et al. (2022b) and Coibion et al. (2022a) show that managing expectations by the central bank is indeed a difficult task and the effects of monetary policy are much smaller than in most theoretical models. D'Acunto et al. (2022) find small effects of forward guidance on inflation expectations and durable consumption.

[4]Thus, the reason for the higher inflation rate is different from earlier papers that also consider an occasionally-binding ELB (e.g., Adam and Billi (2006b), Adam et al. (2022)). In these papers, the higher average inflation rate arises due to promises the policymaker makes *at the lower bound*. In the present paper, on the other hand, the higher inflation rate arises due to considerations *before the lower bound* binds, *foreseeing what will happen* at the lower bound.

perspective and it turns out that this level effect dominates the stabilization benefits. Thus, lower attention is welfare deteriorating when accounting for the ELB.

Another instrument to mitigate the drawbacks of limited attention are negative interest rate policies. Allowing for negative interest rates up to -0.5% (annualized) lowers the necessary increase in the optimal inflation target. As attention declines, however, the effectiveness of negative interest rate policies decreases and the optimal increase in the inflation target is close to the one without negative rates.

Related literature. The empirical part of this paper is related to recent findings in Jørgensen and Lansing (2021) who show that inflation expectations have become more anchored over the last decades.[5] My measure of attention is inversely related to their definition of anchoring, but attention is concerned with short-run expectations whereas anchoring usually refers to the stabilization of long-run expectations. I complement their empirical analysis along several dimensions which I detail more closely in the empirical Section 3.2. Additionally, I offer new insights in how stabilized expectations matter when nominal interest rates are constrained by a lower-bound constraint and what this implies for optimal monetary policy.

My limited-attention model of inflation expectations is closely related to the general information choice problem in Mackowiak et al. (2021). In contrast to their model, and the rational inattention literature more generally, agents in my model have a perceived law of motion of inflation that can potentially differ from the actual law of motion.[6] The reduced-form of the model that I bring to the data is close to the one in Vellekoop and Wiederholt (2019). In contrast to their paper, I focus on how attention changed over the last fifty years and show that attention tends to be higher in times of volatile and persistent inflation.

Ball et al. (2005), Adam (2007), Paciello and Wiederholt (2014) and Gáti (2022) characterize optimal monetary policy in models with different forms of limited attention. I contribute to this literature by allowing for an occasionally-binding lower-bound constraint. I show that accounting for the lower bound can lead to *qualitatively* different welfare implications due to changes in the optimal *level* of inflation. Wiederholt (2015) and Gabaix (2020) examine how information rigidities and inattention matter at the zero lower bound. Angeletos and Lian (2018) study the implications of relaxing the common knowledge assumption for forward guidance and show that the effects of forward guidance are attenuated in such a setting. My paper complements these three papers by studying the Ramsey optimal policy in a fully stochastic setup and focuses on the implications for the optimal inflation target. To the best of my knowledge, the present paper is the first to study the trade off of lower attention in a fully stochastic model with an occasionally binding

[5] Similarly, Gáti (2022) documents that anchoring of long-run inflation expectations is time varying and has substantially increased recently.

[6] Mackowiak et al. (2021) provide a recent overview of this literature, which was inspired by the seminal paper Sims (2003). For further developments in this literature, see, among others, Mackowiak and Wiederholt (2009), Paciello and Wiederholt (2014), Maćkowiak et al. (2018), Afrouzi and Yang (2021), and see Gabaix (2019) for an overview of *behavioral* inattention.

lower-bound constraint and to characterize the Ramsey optimal monetary policy in such a setting.

Outline. The rest of the paper is structured as follows. The empirical strategy to quantify attention, the description of the data and the empirical results are presented in Section 3.2. In Section 3.3, I show how limited attention can lead to inflation-attention traps, before I then study optimal policy in Section 3.4. Section 3.5 concludes.

3.2 Quantifying Attention

In this section, I derive an expectations-formation process under limited attention that provides a straightforward approach to measure attention to inflation empirically. The model is an application of Mackowiak et al. (2021), who study a general problem of optimal information acquisition. I relegate all the details and derivations to appendix C.1.

The main difference to Mackowiak et al. (2021) is that agents in my model do not exactly know the underlying process of inflation but have a simplified view of how inflation evolves.[7] In particular, the agent believes that (demeaned) inflation tomorrow, π', depends on (demeaned) inflation today, π, as follows

$$\pi' = \rho_\pi \pi + \nu,$$

where $\rho_\pi \in [0, 1]$ denotes the perceived persistence of inflation and $\nu \sim i.i.N.(0, \sigma_\nu^2)$. This assumption is supported by empirical evidence (see e.g., Faust and Wright (2013) or Canova (2007)).[8] Note, that the perceived volatility and persistence need not be the same as their actual counterparts, consistent with the empirical evidence on inflation expectations (see, e.g., Table C.1 in the Appendix). The agent wants to minimize her expected forecast error but inflation in the current period is unobservable and acquiring and processing information is costly. The agent thus faces a trade off how attentive she wants to be. I follow the literature on rational inattention and assume that the loss arising from making mistakes in her forecasts is quadratic with a scaling factor r, and the cost of information acquisition and processing is linear in mutual information with a scaling parameter λ.

In this setup and with a normal prior, the optimal signal takes the form

$$s = \pi + \varepsilon,$$

with $\varepsilon \sim i.i.N.(0, \sigma_\varepsilon^2)$ (see Matějka and McKay (2015)).

[7] In the monetary model, later on, I will disentangle the effects from this potentially misperceived law of motion from the assumption of costly attention and show that most of the results are driven by limited attention and not the misperceived law of motion (see section 3.4.4).

[8] Fulton and Hubrich (2021) show that simple models such as AR(1) models are hard to beat when forecasting inflation in real time.

The optimal forecast is given by $\pi^e = \rho_\pi E[\pi|s]$, and Bayesian updating implies

$$\pi^e = \rho_\pi (1-\gamma)\hat{\pi} + \rho_\pi \gamma s, \tag{3.1}$$

where $\gamma = 1 - \frac{\sigma^2_{\pi|s}}{\sigma^2_\pi} \in [0,1]$ measures how much attention the agent pays to inflation, and $\hat{\pi}$ denotes the prior mean of π.

Solving for the optimal γ and writing the cost of information relative to the stakes, $\tilde{\lambda} \equiv \frac{\lambda}{r}$, yields the *optimal* level of attention, summarized in the following Lemma.

Lemma 4. *The optimal level of attention is given by*

$$\gamma = max\left(0, 1 - \frac{\tilde{\lambda}}{2\rho^2_\pi \sigma^2_\pi}\right), \tag{3.2}$$

which shows that the optimal level of attention γ is

(i) decreasing in the relative cost of information acquisition, $\tilde{\lambda} \equiv \frac{\lambda}{r}$,

(ii) increasing in inflation volatility, σ_π, and

(iii) increasing in inflation persistence, ρ_π.

From Lemma 4, we see that aside from the relative information cost, $\tilde{\lambda}$, the persistence, ρ_π, and the volatility of inflation, σ_π, are crucial drivers of attention. The model predicts a positive relationship between attention and σ_π, as well as between attention and ρ_π. In the following, I will first estimate attention γ, asses how it has changed over time and then test whether there is indeed evidence for these positive relations.

3.2.1 Bringing the Model to the Data

To estimate attention in the data, I extend the law of motion of inflation expectations, equation (3.1), to a dynamic setup. The agent believes that inflation π follows

$$\pi_t = (1-\rho_\pi)\bar{\pi} + \rho_\pi \pi_{t-1} + \nu_t,$$

where $\bar{\pi}$ is the agent's long-run belief about inflation and ρ_π is the perceived persistence of inflation. I assume that the error term ν_t is normally distributed with mean zero and variance σ^2_ν.

The agent receives a signal about inflation of the form

$$s_{it} = \pi_t + \varepsilon_{it},$$

where the noise ε_{it} is assumed to be normally distributed with variance σ^2_ε.

Given these assumptions, it follows from the (steady state) Kalman filter that optimal updating is given by

$$\pi^e_{t+1|t,i} = (1-\rho_\pi)\bar{\pi} + \rho_\pi \pi^e_{t|t-1,i} + \rho_\pi \gamma \left(\pi_t - \pi^e_{t|t-1,i}\right) + u_{i,t}, \tag{3.3}$$

where the updating gain γ captures the agent's level of attention. From equation (3.3), we observe that lower attention, i.e., a lower γ, implies that the agent updates her expectations to a given forecast error, $(\pi_t - \pi^e_{t|t-1,i})$, less strongly. Lower attention is reflected in more noisy signals, and more noise means the agent trusts her received signals less and thus, puts less weight on these signals. Hence, her expectations remain more strongly anchored at her prior beliefs.

In the estimation of equation (3.3), I allow for individual-specific intercepts. This can either reflect a mean bias in the perceived inflation rate, $\bar{\pi}_i \neq \bar{\pi}$, or that the agent believes her signals are biased on average, as in Vellekoop and Wiederholt (2019).

3.2.2 Data

I focus on the Survey of Professional Forecasters (SPF) from the Federal Reserve Bank of Philadelphia, as well as the Survey of Consumers from the University of Michigan (SoC). In the Appendix, I show that the findings extend to other data sets as well. For the SPF, I consider individual and aggregate forecasts. The main focus is on expectations about the quarter-on-quarter percentage change in the GDP deflator, which is available since 1969. I drop forecasters for which I have less than eight observations. As a robustness check, I will show that the results are robust to using expectations about the consumer price index, CPI. This data series, however, is only available since 1979.

While the SPF provides data on expectations about the next quarter, the SoC only provides one-year-ahead expectations. Therefore, I will compare them to the actual year-on-year changes in the CPI.[9] As the SoC does not have a panel dimension, I consider average (and median) expectations. Additionally, I estimate attention using the Survey of Consumer Expectations from the Federal Reserve Bank of New York (SCE). The SCE, launched in 2013, has a panel structure and thus allows me to estimate attention using individual-consumer data, at least for the period after 2013. Data on actual inflation comes from the FRED database from the Federal Reserve Bank of St. Louis. Throughout, I focus on the pre-Covid period and end the sample in 2019Q4. Appendix C.2 provides summary statistics and plots the discussed time series.

Jørgensen and Lansing (2021) estimate how strongly anchored inflation expectations in the U.S. are and how this changed over the last fifty years. I extend their empirical strategy in several dimensions. First, I allow the persistence of perceived inflation to change over time and do not restrict it to follow a random walk. Second, I show that not only aggregate professional forecasters' expectations have become more anchored, but also consider individual-specific expectations, as well

[9] The question in the SoC is not explicitly about the CPI but about "prices". Using GDP deflator inflation instead of CPI inflation barely affects the results.

as consumers' inflation expectations. Third, I do not impose the structure of the New Keynesian Phillips Curve on the data but directly estimate attention simply based on the proposed law of motion for inflation expectations. That said, we will see that my results are consistent with the results in Jørgensen and Lansing (2021).

3.2.3 Estimation Results

Before estimating attention, I rewrite the updating equation (3.3) as

$$\pi^e_{t+1|t,i} = \beta_i + \beta_1 \pi^e_{t|t-1,i} + \beta_2 \left(\pi_t - \pi^e_{t|t-1,i}\right) + u_{i,t}, \tag{3.4}$$

where $\beta_i = (1-\rho_\pi)\bar{\pi}_i$, $\beta_1 = \rho_\pi$ and $\frac{\beta_2}{\beta_1} = \gamma$. For the SPF, where I can use individual forecasts, I estimate (3.4) using a forecaster-fixed-effects regression. Since the dependent variable shows up with a lag on the right-hand side, however, the exogeneity assumption is violated. Therefore, I apply the estimator proposed by Blundell and Bond (1998) (BB for short) and I use all available lags of the dependent variable as instruments.[10] All reported standard errors are robust with respect to heteroskedasticity and serial correlation. As an alternative to the BB estimator, I also estimate (3.4) using pooled OLS. For the Survey of Consumers I cannot include individual fixed effects as I use average and median expectations and thus, I apply the Newey-West estimator using four lags (Newey and West (1987)).

To examine how attention changed over time, I run regression (3.4) for the period before and after 1990, separately. The results are robust to different split points (see Appendix C.3). Later on, I will estimate (3.4) using rolling-windows of ten years each and show that the general patterns I document are robust.

Table 3.1: Regression Results of Equation (3.4)

	Professional Forecasters		Consumers	
	Blundell Bond	Pooled OLS	Averages	Median
$\widehat{\gamma}_{pre}$	0.70	0.44	0.75	0.43
s.e.	(0.1005)	(0.0397)	(0.1574)	(0.0970)
$\widehat{\gamma}_{post}$	0.41	0.22	0.31	0.24
s.e.	(0.0522)	(0.0290)	(0.0881)	(0.0601)
N	3566	3566	120	120

Note: This table shows the results from regression (3.4) for professional forecasters (SPF) as well as for consumers. For the SPF, I use the Blundell and Bond (1998) (BB) estimator (first two columns), as well as pooled OLS (columns 3-4). For the Survey of Consumer, I consider average expectations (columns 5-6) and median expectations (columns 7-8). $\widehat{\gamma}_{pre}$ and $\widehat{\gamma}_{post}$ denote the estimated attention parameters for the period pre 1990 and post 1990, respectively. The standard errors, reported in parentheses, are robust with respect to heteroskedasticity and serial correlation.

Table 3.1 shows the results. Here, $\widehat{\gamma}_{pre}$ and $\widehat{\gamma}_{post}$ denote the estimated attention parameters

[10] Appendix C.3 shows that the results are robust to using fewer lags.

for the period pre 1990 and post 1990, respectively. We see that attention is substantially lower after 1990 compared to the period before 1990.[11] This is true for professional forecasters and for consumers, and as I show in Appendix C.3, also for other datasets. The point estimates after 1990 are basically half of what they were before the 1990s and these differences are statistically significant at the 1% level.

The decline in attention is even more pronounced when focusing on the most recent decade. To show this, I run regression (3.4) for the period between 2010 and 2020. Additionally, I also use data from the New York Fed Survey of Consumer Expectations, starting in 2013. The advantage of this survey compared to the Michigan Survey is that it surveys the same consumers up to twelve times in a row, providing a much larger sample size and a panel dimension. Table 3.2 shows the results. We see that overall, attention declined substantially compared to earlier periods and is between 0.04 and 0.17 during this period of low and stable inflation. Furthermore, the results from the SCE lie in the same ballpark as the ones from the Michigan Survey, which indicates that using average (or median) consumer expectations does not fundamentally affect the results. In fact, the estimated attention parameter for the average expectations from the Michigan Survey is 0.04 when restricting the sample to 2013-2020, which is exactly the same as the estimate obtained from the New York Fed Survey.

Table 3.2: Attention since 2010

	Professional Forecasters		Consumers		NY Fed Survey
	Blundell Bond	Pooled OLS	Averages	Median	Pooled OLS
$\widehat{\gamma}$	0.17	0.07	0.12	0.09	0.04
s.e.	(0.0729)	(0.0333)	(0.0658)	(0.0616)	(0.0316)
N	1322	1322	40	40	74229

Note: This table shows the results from regression (3.4) for the period between 2010 and 2020 for professional forecasters (SPF) as well as for consumers. For the SPF, I use the Blundell and Bond (1998) (BB) estimator (first column), as well as pooled OLS (column 2). For the Survey of Consumer, I consider average expectations (column 3) and median expectations (columns 4). The standard errors, reported in parentheses, are robust with respect to heteroskedasticity and serial correlation. Additonally, column 5 shows the results for consumer inflation expectations from the New York Fed Survey of Consumer Expectations.

When is Attention High?

Lemma 4 states that two key drivers of attention are inflation volatility and inflation persistence. To examine these relationships empirically, I estimate regression (3.4), using a rolling-window approach in which every window is 10 years long, and estimate one attention parameter for every window, $\widehat{\gamma}_t$, as well as the period-specific inflation volatility, $\widehat{\sigma}_{\pi,t}$ and the persistence parameter,

[11] I test the validity of the instruments in the Blundell-Bond estimation by testing for autocorrelation of order one and two in the first-differenced error terms. The respective p-values are 0.000 (order 1) and 0.973 (order 2) for the period before 1990 and 0.000 (order 1) and 0.737 (order 2) for the period after 1990. This indicates that the instruments used in the estimation are valid.

$\widehat{\rho}_{\pi,t}$. In particular, I use the window-specific standard deviation of inflation as my measure of $\widehat{\sigma}_{\pi,t}$ and the first-order autocorrelation of inflation for $\widehat{\rho}_{\pi,t}$. Appendix C.3 shows that the following results also hold for different window lengths or when using the standard deviation and persistence of *expected* inflation for $\widehat{\sigma}_{\pi,t}$ and $\widehat{\rho}_{\pi,t}$.

Figure 3.1 summarizes the results graphically. The left scatterplot shows the inflation volatility on the horizontal axis and the estimated attention parameter, $\widehat{\gamma}$, on the vertical axis. The attention parameters shown in the figure are the ones for individual professional forecasters, obtained via pooled OLS. The right panel shows the relationship between attention and inflation persistence (on the horizontal axis). In both cases, we see that there is a clear positive relationship, just as the limited-attention model predicts.

Figure 3.1: Attention, Inflation Volatility and Inflation Persistence

(a) Inflation Volatility (b) Inflation Persistence

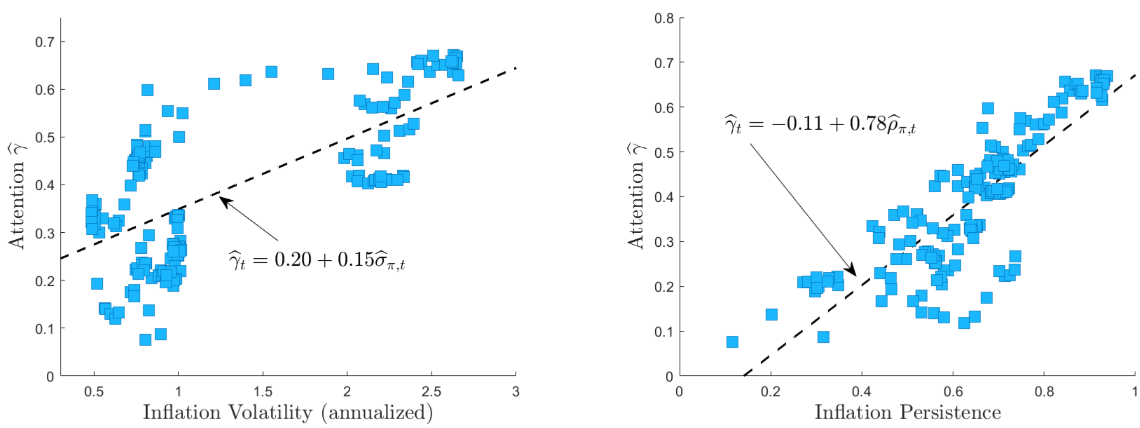

Notes: The right panel shows the relationship between inflation volatility and the estimated attention parameter, $\widehat{\gamma}$. The left panel shows the relationship between the persistence of inflation and the estimated attention parameter. Both panels report the results for individual professional forecasters where the attention parameter was estimated via pooled OLS.

To check if these findings are statistically significant, I regress attention on inflation volatility or on inflation persistence as follows

$$\widehat{\gamma}_t = \alpha_1 + \beta\widehat{\sigma}_{\pi,t} + u_t \tag{3.5}$$

$$\widehat{\gamma}_t = \alpha_2 + \zeta\widehat{\rho}_{\pi,t} + v_t. \tag{3.6}$$

Table 3.3 reports the results. Standard errors are robust with respect to heteroskedasticity. We see that the observed patterns in Figure 3.1 are indeed statistically significant. Attention to inflation is positively correlated with inflation volatility and inflation persistence, as Lemma 4 predicts. This is true for professional forecasters as well as for consumers. Overall, the magnitudes of the estimates indicate that the results are somewhat stronger for professional forecasters. Appendix C.3 shows that these results hold when regressing attention on inflation volatility and

Table 3.3: Attention, Inflation Volatility and Inflation Persistence

	Professional Forecasters		Consumers
Estimator	Blundell-Bond	Pooled OLS	OLS
$\widehat{\beta}$	0.13***	0.15***	0.09***
s.e.	(0.0155)	(0.0105)	(0.0272)
$\widehat{\zeta}$	0.71***	0.78***	0.56***
s.e.	(0.0568)	(0.0349)	(0.0714)
N	165	165	163

Note: This table shows the results of regression (3.5) and (3.6). Standard errors are robust with respect to heteroskedasticity. *** : p-value < 0.01, ** : p-value < 0.05, * : p-value < 0.1.

persistence jointly. I also show that the results are robust when controlling for the *average level* of inflation and the average level of inflation has no significantly-positive effect on attention when controlling for volatility and persistence.[12] Given the decline in inflation volatility and inflation persistence over the last fifty years (see Table C.1 in Appendix C.2), the positive correlation with attention supports the findings in Table 3.1: attention declined as the volatility and persistence of inflation declined.[13]

Additional evidence and robustness. In Appendix C.3.1, I document a decline in news coverage of inflation in popular news papers as well as in books, thus, providing additional, complementary evidence on the decline in attention after the 1980s.[14] In Appendix C.3, I show that the presented results are robust to using different data sources, different sample splits, different specifications of the BB estimator, allowing for time-fixed effects to account, for example, for varying trend inflation, as well as constructing a quasi panel of consumers, based on their income.

3.3 Monetary Policy Implications of Limited Attention

How does the decline in attention to inflation affect the conduct of monetary policy? To answer this question, I augment the standard New Keynesian model with inflation expectations that are characterized by limited attention, and a lower-bound constraint on the nominal interest rate. To do this, I replace the inflation expectations under full-information rational expectations (FIRE) with their limited-attention counterpart derived in the previous section.[15] This way of

[12] From a conceptual perspective, it is also not clear why the *level* of attention should directly matter. Consider an economy with a very high but constant inflation rate. The agents could then simply set their expectations to this average value and never update their expectations, i.e., be completely inattentive. Since inflation is constant, however, they would predict inflation correctly and paying zero attention would nevertheless be optimal.

[13] Benati (2008) documents a decline in inflation persistence in advanced economies, especially for countries that introduced inflation targeting regimes.

[14] Carroll (2003) proposes a micro-foundation of sticky information models (as in Mankiw and Reis (2002)) that relies on news coverage of inflation. Lamla and Lein (2014) and Pfajfar and Santoro (2013) test these predictions empirically.

[15] One way to *micro-found* this would be similar to Section V in Angeletos and La'o (2010), where, for example, firms have two managers that do not communicate with each other. In my setting, this would require that one

incorporating limited attention in the model allows me to study the Ramsey optimal policy in the presence of an occasionally-binding lower bound on the policy rate. Apart from the formation of inflation expectations, I build on the standard New Keynesian model without capital, with rigid prices in the spirit of Calvo (1983) and Yun (1996) and with a lower bound on the nominal interest rate. The government pays a subsidy to intermediate-goods producers to eliminate steady state distortions arising from market power.

The linearized model yields an aggregate supply equation, the *New Keynesian Phillips Curve*, and an aggregate Euler (or IS) equation:

$$\pi_t = \beta \pi^e_{t+1|} + \kappa y^{gap}_t + u_t, \tag{3.7}$$

$$y^{gap}_t = E_t y^{gap}_{t+1} - \varphi \left(i_t - \pi^e_{t+1|t} - r^n_t \right), \tag{3.8}$$

where κ measures the sensitivity of aggregate inflation to changes in the output gap, y^{gap}_t, $\beta \in (0,1)$ denotes the time discount factor of the representative household, and u_t are cost-push shocks, following an AR(1) process with persistence $\rho_u \in [0,1]$ and innovations $\varepsilon^u \sim i.i.N.(0, \sigma^2_u)$. The output gap is the log deviation of output from its efficient counterpart that would prevail under flexible prices. Altogether, equation (3.7) summarizes the aggregate supply side of the economy. Equation (3.8), together with monetary policy, determines aggregate demand in this model. Here, $\varphi > 0$ measures the real rate elasticity of output, i_t is the nominal interest rate which is set by the monetary authority, and r^n_t is the natural interest rate. The natural interest rate is the real rate that prevails in the economy with fully flexible prices and is exogenous. It follows an AR(1) process with persistence $\rho_r \in [0,1]$ and innovations $\varepsilon^r \sim i.i.N.(0, \sigma^2_r)$, independent of ε^u. The nominal interest rate and the natural rate are both expressed in absolute deviations of their respective steady state values, $\bar{i}$ and $\bar{r}^n$, with $\bar{i} = \bar{r}^n$, as I linearize the model around the zero-inflation steady state. E_t denotes the full-information rational expectations operator.

Inflation expectations are characterized by limited attention and are given by

$$\pi^e_{t+1|t} = (1 - \rho_\pi)\bar{\pi} + \rho_\pi \pi^e_{t|t-1} + \rho_\pi \gamma \left(\pi_t - \pi^e_{t|t-1} \right), \tag{3.9}$$

where the notation is the same as in Section 3.2. Given the representative agent assumption, I abstract from noise shocks in (3.9). The law of motion of inflation expectations, equation (3.9), is driven by two main assumptions. First, the perceived law of motion follows an AR(1) process and second, paying attention is costly. In Section 3.4.4, I will disentangle the two and show that it is mainly the second assumption that drives the implications for optimal monetary policy. For the

manager sets the firm's price and makes its employment decision for some *given inflation expectation* and the other manager provides this inflation forecast, according to the steps outlined in Section 3.2, without taking the other manager's decisions into account (and similary for the representative household). An approach similar to mine, namely replacing the expectations after the derivation of the equilibrium conditions, is taken in Jørgensen and Lansing (2021) or parts of the *learning* literature (see, for example, Evans and Honkapohja (2003) or Milani (2007) for early contributions).

most part, I will focus on $\rho_\pi = 1$ in which case average inflation expectations align with actual average inflation and long-run beliefs $\bar{\pi}$ are irrelevant. I discuss the case with $\rho_\pi < 1$ in Appendix C.6.4. For empirically-realistic values of ρ_π the results are very similar to the case with $\rho_\pi = 1$. This belief formation process is empirically plausible, in the sense that it is consistent with recent empirical findings, documented in Angeletos et al. (2021): after a shock, expectations initially underreact, followed by a delayed overreaction (see Appendix C.6.3).

As is standard in the rational inattention literature, I assume that the attention parameter γ is constant.[16] The usual assumption to obtain this is that in period $t = 0$ the agent chooses her level of attention and then obtains all future signals at this point. This leaves conditional second moments time-invariant and thus, the optimal level of attention constant. I will, however, compare economies with different levels of attention.

Table 3.4: Model Parameterization

Parameter	Value	Source/Target
	Preferences and technology	
β	0.9975	Average natural rate of 1%
φ	1	Adam and Billi (2006b)
κ	0.057	Adam and Billi (2006b)
	Exogenous shock processes	
ρ_r	0.8	Adam and Billi (2006b)
σ_r	0.2940%	Adam and Billi (2006b)
ρ_u	0	Adam and Billi (2006b)
σ_u	0.154%	Adam and Billi (2006b)

Calibration. I calibrate the model to quarterly frequency. I assume an annualized steady state natural rate of 1%. The rest of the calibration is taken from Adam and Billi (2006b). Table 3.4 summarizes the calibration. The attention parameter γ will be varied to understand its role for monetary policy.

Attention and the Phillips Curve. Inflation expectations are a crucial driver of actual inflation through the supply side of the economy, and thus, changing levels of attention affect the Phillips Curve. The following Proposition summarizes these effects.

Proposition 7. *The New Keynesian Phillips Curve under limited attention is given by*

$$\pi_t = \frac{\beta(1-\gamma)}{1-\beta\gamma}\pi^e_{t|t-1} + \frac{\kappa}{1-\beta\gamma}y^{gap}_t + \frac{1}{1-\beta\gamma}u_t.$$

Proof. See Appendix C.5. ■

[16]See, e.g., Mackowiak and Wiederholt (2009), Maćkowiak et al. (2018); Mackowiak et al. (2021).

From Proposition 7, we see that for a given prior inflation expectation, $\pi^e_{t|t-1}$, and a given realization of the exogenous shock u_t, inflation becomes less sensitive to changes in the output gap at lower levels of attention. Put differently, a decrease in attention γ resembles a flatter Phillips Curve. We also see that for a given output gap, inflation becomes less sensitive to cost-push shocks u_t at lower levels of attention.

Thus, Proposition 7 captures the stabilizing effects of lower attention. As attention declines, firms' inflation expectations react less to changes in actual inflation. Through the Phillips Curve, this muted reaction of expectations in turn stabilizes inflation itself. What cannot be seen from Proposition 7, but what will be crucial in the subsequent analysis, is that lower attention not only affects the initial response of inflation and inflation expectations but the dynamics as well. Changes in inflation and inflation expectations become more persistent at low levels of attention, even though the initial response is muted. I now show in a numerical example that at low levels of attention the economy can get stuck in an *inflation-attention trap*.[17]

3.3.1 Inflation-Attention Traps

To close the model from Section 3.3, I assume for now that, away from the lower bound, the monetary authority sets the nominal interest rate according to a Taylor rule

$$\tilde{i}_t = \rho_i \tilde{i}_{t-1} + (1 - \rho_i) \left(\phi_\pi \pi_t + \phi_y y_t^{gap}\right), \tag{3.10}$$

where $\rho_i \in [0, 1)$ captures persistence in the interest rate, $\phi_\pi > 1$ and $\phi_y \geq 0$ denote the reaction coefficients to inflation and the output gap, respectively. The actual interest rate, i_t, however is constrained by the lower bound

$$i_t = max\{\tilde{i}_t, -\bar{i}\}, \tag{3.11}$$

where I set the lower bound (in levels) to zero.

I set the persistence parameter of the nominal interest rate to 0.7, and the reaction coefficients $\phi_\pi = 2$ and $\phi_y = 0.5$, as in Andrade et al. (2019). In Appendix C.6.2, I show that the exact specification of the Taylor rule is inconsequential for the following results.

Figure 3.2 plots the impulse response functions of the model's main variables to a negative natural rate shock of three standard deviations that pushes the nominal interest rate to the lower bound. The black-dashed-dotted lines are the IRFs in the model under FIRE and the blue-dashed lines are the ones under limited attention for the case $\gamma = 0.3$.

In both cases, the shock is large enough to push the economy to the lower bound. While the reaction of the output gap is very similar in both economies, the responses of inflation and inflation expectations are strikingly different. Initially, the muted response of inflation expectations to the adverse shock is reflected in a smaller downturn of inflation itself under limited attention. This

[17] In Appendix C.4, I show analytically in a stylized example how lower levels of attention mute the effects of forward guidance.

Figure 3.2: Impulse Response Functions to a Negative Natural Rate Shock

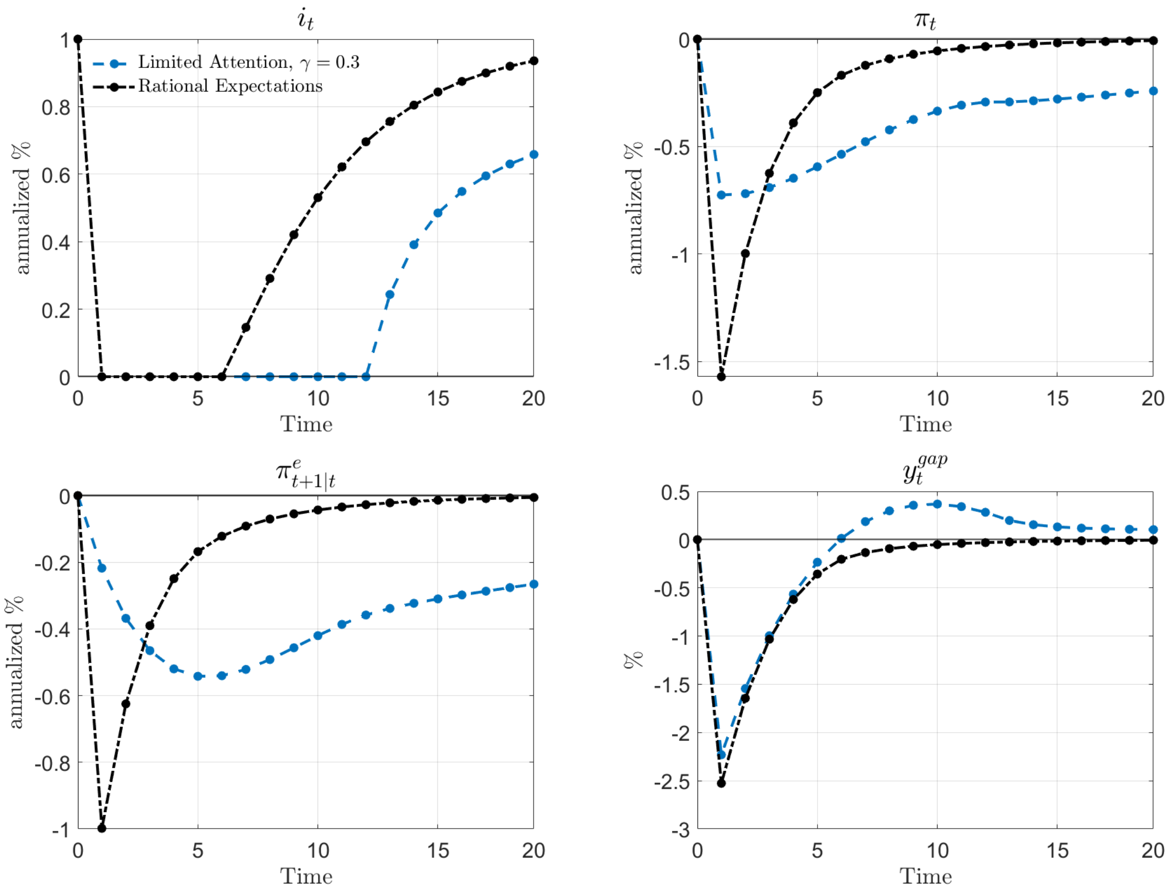

Note: This figure shows the impulse-response functions of the nominal interest rate (upper-left panel), inflation (upper-right panel), inflation expectations (lower-left) and the output gap (lower-right) to a negative natural rate shock of three standard deviations. The blue-dashed lines show the case for the limited-attention model and the black-dashed-dotted lines for the rational expectations model. Everything is in terms of percentage deviations from the respective steady state levels, except the nominal rate is in levels.

captures the stabilizing effects that come with lower attention. The sluggish adjustment of inflation expectations in the following, however, leads to a very persistent undershooting of inflation. Even five years after the shock, inflation and inflation expectations are still substantially below their steady state levels of zero. The result is a prolonged period of a binding lower bound. While the economy under rational expectations escapes the ELB six periods after the shock, the economy under limited attention is stuck for twice as long. This is what I label *inflation-attention trap*. A side-effect of these traps is that the long ELB period leads to an output boom. As discussed earlier, this (expected) output boom in the future, however, has rather small effects on the economy today if people are inattentive.

Overall, limited attention to inflation offers a possible explanation for why several advanced economies were stuck at the ELB after the financial crisis, as well as inflation that undershot the central banks' inflation targets, even though the initial decrease was muted and output recovered quite strongly after declining severely initially (Del Negro et al. (2020)). In other words, the limited-attention model can explain the *missing deflation puzzle* as well as the *missing inflation puzzle* (Coibion and Gorodnichenko (2015b), Constancio (2015)).

Non-rational output gap expectations. In appendix C.6.1, I relax the assumption that agents have rational expectations with respect to the output gap and instead assume that they are also inattentive to the output gap.

In that case, inflation-attention traps get exacerbated. The economy remains stuck at the ELB even longer and, additionally, inflation, inflation expectations, and now also the output gap stay below their initial values very persistently. Thus, the results reported in the following are likely to be even stronger when allowing for non-rational output gap expectations.

3.4 Optimal Monetary Policy

To understand how monetary policy should optimally deal with declining attention, I now derive the Ramsey optimal monetary policy in this economy. I focus on the case of $\rho_\pi = 1$, in which average inflation expectations coincide with the actual inflation average. Appendix C.6.4 reports the results when relaxing this assumption.

The policymaker's objective is to maximize the representative household's utility, taking the household's and firms' optimal behavior, including their attention choice, as given. Thus, the policymaker cannot exploit the private agent's lack of information. Nevertheless, the policymaker can affect inflation expectations by influencing inflation itself and he can set the average inflation expectations by setting the average inflation rate.

The policymaker is paternalistic in the sense of Benigno and Paciello (2014) and evaluates the household's utility under rational expectations. A second-order approximation to the household's utility function yields the policymaker's objective

$$-\frac{1}{2}E_0\sum_{t=0}^{\infty}\beta^t\left[\pi_t^2+\chi\left(y_t^{gap}\right)^2\right], \tag{3.12}$$

where χ is the relative weight of the output gap. Following Adam and Billi (2006b), I set $\chi = 0.007$. In the following, I refer to (3.12) as *welfare*.

In sum, the optimal policy problem is given by

$$\max_{\pi_t, y_t^{gap}, i_t} \quad -\frac{1}{2}E_0\sum_{t=0}^{\infty}\beta^t\left[\pi_t^2+\chi\left(y_t^{gap}\right)^2\right] \tag{3.13}$$

subject to

$$\pi_t = \beta\pi_{t+1|t}^e + \kappa y_t^{gap} + u_t \tag{3.14}$$

$$y_t^{gap} = E_t y_{t+1}^{gap} - \varphi \left(i_t - \pi_{t+1|t}^e - r_t^n \right) \tag{3.15}$$
$$\pi_{t+1|t}^e = \pi_{t|t-1}^e + \gamma \left(\pi_t - \pi_{t|t-1}^e \right) \tag{3.16}$$
$$u_t = \rho_u u_{t-1} + \varepsilon_t^u \tag{3.17}$$
$$r_t^n = \rho_r r_{t-1}^n + \varepsilon_t^r \tag{3.18}$$
$$i_t \geq -\bar{i}, \tag{3.19}$$

with $\varepsilon_t^u \sim i.i.N.\left(0, \sigma_u^2\right)$ and $\varepsilon_t^r \sim i.i.N.\left(0, \sigma_{r^n}^2\right)$ and (3.19) is the lower-bound constraint.[18] All variables are in percent deviations from their respective steady state, except the nominal interest rate and the natural rate which are in absolute deviations.

3.4.1 The Optimal Inflation Target

What do low levels of attention imply for inflation volatility and the optimal inflation target? For this, I solve the Ramsey problem for different levels of attention, namely $\gamma \in \{0.05, 0.1, 0.2, 0.3\}$. An attention parameter of 0.3 is close to the estimates for consumers' attention after 1990, and the lower levels of 0.05 and 0.1 are close to the ones observed since 2010.

Figure 3.3: Optimal Inflation and Inflation Volatility

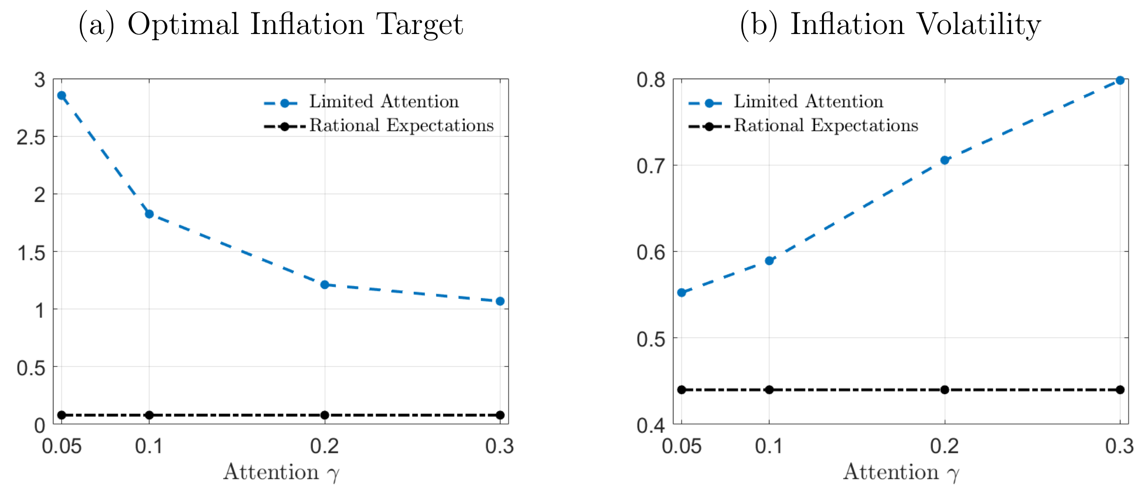

Notes: This figure shows the average inflation rate under Ramsey optimal policy (left panel) and inflation volatility for different attention levels. The blue-dashed lines show the results for the model under limited attention, and the black-dashed-dotted lines for the rational-expectations model.

Figure 3.3 shows the results. The average inflation rate under Ramsey optimal policy—what I refer to as the optimal inflation target—is plotted in the left panel and the inflation volatility in the right panel. The blue-dashed lines show the results for the model under limited attention, and the black-dashed-dotted lines for the rational-expectations model. We see that the optimal inflation target increases substantially as attention declines. At the levels of attention estimated

[18] I solve this numerically by recursifying the constrained optimization problem, as in Marcet and Marimon (2019).

just before the Covid crisis, $\gamma \in \{0.05, 0.1\}$, the inflation target is about 2-3 percentage points higher than under rational expectations due to the discussed ineffectiveness of forward-guidance policies. By increasing the average level of inflation the average nominal interest rate increases and thus makes it less likely that the ELB becomes binding. Indeed, the frequency of a binding ELB decreases substantially (not shown). For $\gamma = 0.3$, the ELB is binding 22% of the time under optimal policy, whereas this value shrinks to 1.4% for an attention level of 0.05.

While lower attention renders forward guidance, make-up policies and other policies that work (partly) through inflation expectations less effective, lower attention also stabilizes inflation, as can be seen from the right panel in Figure 3.3. First, because lower attention mutes the inflation response to shocks and output (see Proposition 7). Second, the lower ELB frequency further stabilizes the economy. Thus, lower attention to inflation can help stabilizing actual inflation and reduces the number of binding-ELB periods. The lower inflation volatility at lower levels of attention in fact justifies these low attention levels, as optimal attention depends positively on inflation volatility (see Section C.1). This low volatility, however, requires an increase in the inflation target, which is costly. Thus, it is not clear *a priori* whether lower attention leads to welfare gains or not.

3.4.2 Welfare

What are the effects of declining attention on overall welfare? Welfare is given by equation (3.12) and from the previous discussion, we know that lower attention poses a trade off. On the one hand, inflation volatility decreases and the ELB binds less frequently, when attention is low. This raises welfare. On the other hand, lower attention complicates managing inflation expectations and thus, the optimal average *level* of inflation increases, which is costly. Which effect dominates?

Panel (a) in Figure 3.4 shows that the cost of the level effect outweighs the stabilization benefits. As attention falls, welfare decreases. This is especially pronounced at low levels of attention, where the optimal inflation target increases substantially (see Figure 3.3).

Absent the lower-bound constraint, the complications in managing inflation expectations due to limited attention are much less pronounced since managing expectations is particularly important at the lower bound. In fact, lower attention is welfare improving in the case without an ELB. Panel (b) in Figure 3.4 shows this graphically. The stabilization benefits that arise from lower attention—which is reflected in more anchored expectations—lead to an increase in welfare.

These findings show that accounting for the ELB is crucial for making a normative statement about costs and benefits of stabilizing inflation expectations. The ELB highlights the drawbacks that arise from the stabilization of expectations due to the fall in attention, as the management of expectations becomes particularly relevant when the ELB binds.

Figure 3.4: Welfare and Attention

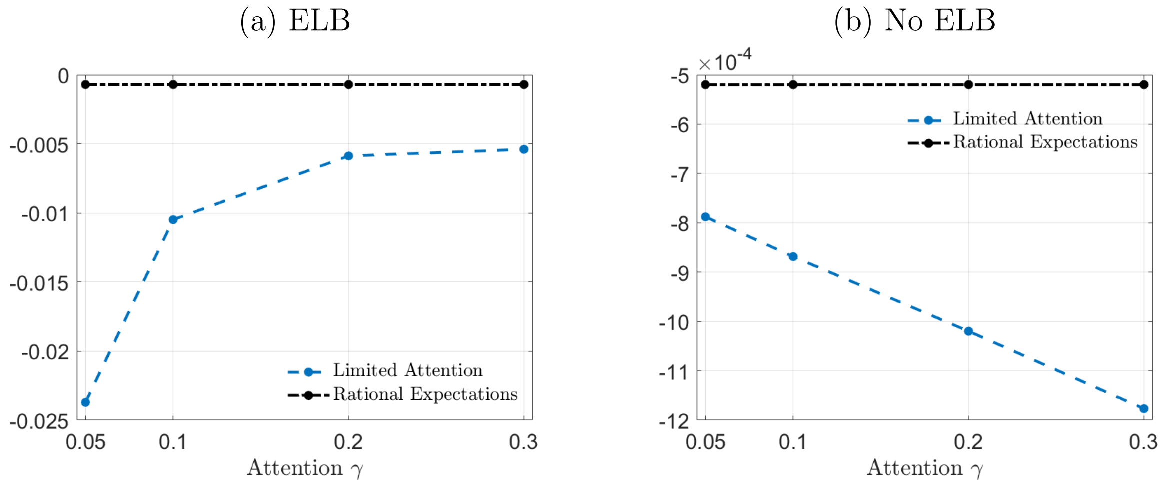

Notes: This figure shows welfare (3.12) under Ramsey optimal policy for different levels of attention. The left panel shows the results for the case with an occasionally-binding ELB, and the right panel without an ELB. The blue-dashed lines show the results for the model under limited attention, and the black-dashed-dotted lines for the rational-expectations model.

3.4.3 Negative Interest Rate Policies

In recent years, several central banks in advanced economies have implemented negative interest rate policies (NIRP).[19] Could negative rates limit the negative consequences of declining attention? In order to answer this question, I solve the same Ramsey optimal policy problem as above, but set the effective lower bound to -0.5% (annualized).

Figure 3.5 reports the outcomes. Panel (a) shows the optimal inflation target (red-dashed-dotted line) and compares it to the case with an ELB at 0 (blue-dashed line). We see that the additional policy space due to the negative lower bound indeed calls for a lower inflation target. However, the decline in attention also weakens the effectiveness of NIRP. We see this by observing that the optimal inflation target under NIRP gets closer to the one without negative rates as attention declines. To see these gaps clearly, panel (b) shows the difference in the optimal inflation targets, defined as $\pi^{*,ZLB} - \pi^{*,NIRP}$, where π^* denotes the optimal inflation target and the superscripts ZLB and $NIRP$ denote the two cases where the ELB is at 0% or -0.5%, respectively. Overall, allowing for negative policy rates can help limiting the drawbacks of low attention but these policies itself become less effective as attention declines.

3.4.4 Full Attention

There are two main assumptions underlying the law of motion of inflation expectations (3.9). First, the assumption that agents perceive inflation to follow an AR(1) process. Second, that paying attention to inflation is costly and thus, their attention is limited. To disentangle these two effects,

[19]See Brandão-Marques et al. (2021) for a recent survey on negative interest rate policies and its effectiveness.

Figure 3.5: Negative Interest Rate Policies and Attention

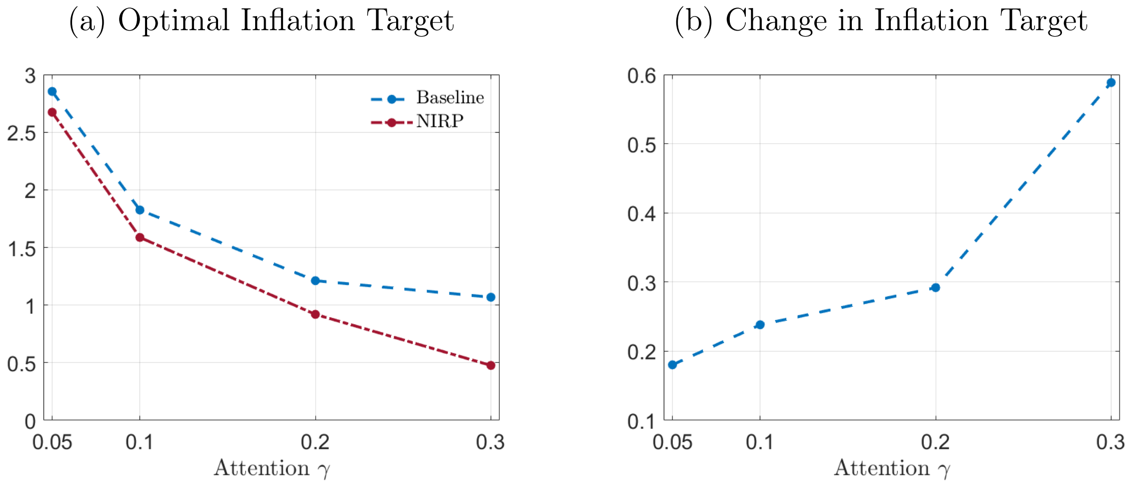

Notes: The left panel shows the average inflation rate under optimal policy for different degrees of attention γ. The blue-dashed lines show the results for the benchmark model where the lower bound is at 0, and the red-dashed-dotted lines show the results when allowing for negative interest rates up to -0.5% (annualized). The right panel shows the difference in the optimal inflation targets, defined as $\pi^{*,ZLB} - \pi^{*,NIRP}$, where π^* denotes the optimal inflation target and the superscripts ZLB and $NIRP$ denote the two cases where the ELB is at 0% or -0.5%, respectively.

Figure C.10 in Appendix C.6.5 shows the implications of shutting down the second channel, i.e., if we set $\gamma = 1$. Put differently, how much of the results in previous sections is solely due to the misperception of the law of motion of inflation? Panel (a) shows that the optimal average inflation rate is practically identical to the one under rational expectations. Thus, as argued above, it is really the lack of attention that pushes up the optimal inflation rate, whereas the assumption of a misperceived law of motion is rather innocuous from this perspective.

Panel (b) in Figure C.10 reports the inflation volatility under Ramsey optimal policy for different levels of attention, including the full attention case, $\gamma = 1$. Inflation is substantially more volatile than under rational expectations. Thus, while the misperception of the law of motion of inflation is rather inconsequential for the optimal inflation target, it predicts a higher inflation volatility. Indeed, there is a trade-off. On the one hand, higher attention increases the response of inflation to shocks. On the other hand, if agents are fully attentive, inflation also reacts more strongly to changes in expected future output. Thus, to achieve a certain effect on today's inflation rate, the policymaker is required to make smaller promises about its policies in the future which in turn stabilizes inflation and output already today. While the first channel dominates at lower levels of γ, the second effect pushes inflation volatility down at higher levels of γ.

Finally, panel (c) in Figure C.10 shows the welfare implications of the misperceived law of motion. Welfare is slightly more negative when agents have a misperceived law of motion of inflation compared to fully rational agents. Given the results in panels (a) and (b), we see that these additional welfare losses are mainly due to increased inflation volatility rather than its level.

3.5 Conclusion

With the stabilization of inflation in advanced economies since the Great Inflation period, inflation has become less important in people's lives. In this paper, I quantify this using a limited-attention model of inflation expectations. In line with this model, I show that attention to inflation decreased together with inflation volatility and inflation persistence since the 1970s. Especially in the period between 2010 and 2020, the general public's attention to inflation was close to zero.

For monetary policy the decline in attention was desirable at first, since lower attention stabilizes inflation expectations and hence, stabilizes actual inflation. With the outbreak of the Great Recession and nominal rates at their lower bound, however, managing inflation expectations became a central tool for monetary policy. But managing inflation expectations is difficult when people are inattentive.

The optimal policy response is a substantial increase in the inflation target. This increases the average nominal rate and thus, binding ELB periods become less likely. The cost of this increase in inflation, however, outweighs the stabilization benefits of lower attention. Lower attention, therefore, decreases welfare if we account for the lower bound. This stands in stark contrast to the case without an ELB in which case lower attention leads to welfare gains through the stabilization of inflation expectations and inflation. My paper thus shows that accounting for the ELB is crucial when assessing the role of the public's attention to inflation.

Appendix C

Appendix to Chapter 3

C.1 A Limited-Attention Model of Inflation Expectations

In this section, I derive the expectations-formation process under limited attention sketched in Section 3.2. The agent believes that (demeaned) inflation tomorrow, π', depends on (demeaned) inflation today, π, as follows

$$\pi' = \rho_\pi \pi + \nu,$$

where $\rho_\pi \in [0, 1]$ denotes the perceived persistence of inflation and $\nu \sim i.i.N.(0, \sigma_\nu^2)$. Inflation in the current period is unobservable, so before forming an expectation about future inflation, the agent needs to form an expectation about today's inflation. I denote this nowcast $\widetilde{\pi}$, and the resulting forecast about next period's inflation $\pi^e = \rho_\pi \widetilde{\pi}$. Given her beliefs, the full-information forecast π^{e*} is

$$\pi^{e*} \equiv \rho_\pi \pi.$$

But since π is not perfectly observable, the actual forecast will deviate from the full-information forecast. Deviating, however, is costly, as this causes the agent to make mistakes in her decisions.

The agent's choice is not only about how to form her expectations given certain information, but about how to choose this information optimally, while taking into account how this will later affect her forecast. That is, she chooses the form of the signal s she receives about current inflation. Since acquiring information is costly, it cannot be optimal to acquire different signals that lead to an identical forecast. Due to this one-to-one relation of signal and forecast, we can directly work with the joint distribution of π^e and π, $f(\pi^e, \pi)$, instead of working with the signal.

Let $U(\pi^e, \pi)$ denote the negative of the loss that is incurred when the agent's forecast deviates from the forecast under full information, and $C(f)$ the cost of information. Then, the agent's problem is given by

$$\max_f \int U(\pi^e, \pi) f(\pi^e, \pi) d\pi d\pi^e - C(f) \tag{C.1}$$

$$\text{subject to} \int f(\pi^e, \pi) d\pi^e = g(\pi), \text{ for all } \pi,$$

where $g(\pi)$ is the agent's prior, which is assumed to be Gaussian; $\pi \sim N(\hat{\pi}, \sigma_\pi^2)$. $C(.)$ is the cost function that captures how costly information acquisition is. It is linear in *mutual information* $I(\pi; \pi^e)$, i.e., the expected reduction in entropy of π due to knowledge of π^e:

$$C(f) = \lambda I(\pi; \pi^e) = \lambda \left(H(\pi) - E\left[H(\pi|\pi^e) \right] \right),$$

where $H(x) = -\int f(x) log(f(x)) dx$ is the entropy of x and λ is a parameter that measures the cost of information.

The objective function $U(.)$ is assumed to be quadratic:

$$U(\pi^e, \pi) = -r \left(\rho_\pi \pi - \pi^e \right)^2,$$

where r measures the stakes of making a mistake.[1,2]

In this setup, Gaussian signals are optimal (and in fact the unique solution, see Matějka and McKay (2015)). The optimal signal thus has the form

$$s = \pi + \varepsilon,$$

with $\varepsilon \sim i.i.N.(0, \sigma_\varepsilon^2)$.[3] The problem (C.1) now reads

$$\max_{\sigma_{\pi|s}^2 \leq \sigma_\pi^2} E_\pi \left[E_s \left[-r\rho_\pi^2 \left(\pi - E[\pi|s] \right)^2 \right] \right] - \lambda I(\pi; \pi^e) = \max_{\sigma_{\pi|s}^2 \leq \sigma_\pi^2} \left(-r\rho_\pi^2 \sigma_{\pi|s}^2 - \frac{\lambda}{2} log \frac{\sigma_\pi^2}{\sigma_{\pi|s}^2} \right). \tag{C.2}$$

The optimal forecast is given by $\pi^e = \rho_\pi E[\pi|s]$, and Bayesian updating implies

$$\pi^e = \rho_\pi (1 - \gamma) \hat{\pi} + \rho_\pi \gamma s, \tag{C.3}$$

where $\gamma = 1 - \frac{\sigma_{\pi|s}^2}{\sigma_\pi^2} \in [0, 1]$ measures how much attention the agent pays to inflation, and $\hat{\pi}$ denotes the prior mean of π.

[1] A quadratic loss function is usually derived from a second-order approximation of the household's utility function or the firm's profit function (see, e.g., Mackowiak and Wiederholt (2009)).

[2] These stakes (or also the information cost parameter λ) can be interpreted as a way to incorporate other variables to which the agent might pay attention. For example, a household might not only want to forecast inflation but also her own income stream going forward. In this case, a smaller r could capture an increase in her idiosyncratic income volatility. Thus, paying attention to inflation is relatively less beneficial, as the relative importance of her idiosyncratic income increases. Such an interpretation also explains why professional forecasters might not be fully informed about inflation, given that they usually forecast a whole array of variables.

[3] In this case, the entropy becomes $H(x) = \frac{1}{2} log(2e\sigma_x^2)$, where σ_x^2 is the variance of x. Note, that here denotes the number "pi" and not inflation.

An equivalent way of writing γ is

$$\gamma = \frac{\sigma_\pi^2}{\sigma_\pi^2 + \sigma_\varepsilon^2}. \tag{C.4}$$

Now, since the agent *chooses* the level of attention, we can re-formulate (C.2) as

$$\max_{\gamma \in [0,1]} \left(-r\rho_\pi^2 (1-\gamma)\sigma_\pi^2 - \frac{\lambda}{2} log \frac{1}{1-\gamma} \right). \tag{C.5}$$

Writing the cost of information relative to the stakes, $\tilde{\lambda} \equiv \frac{\lambda}{r}$, and solving the optimization problem (C.5) yields the *optimal* level of attention, presented in Lemma 4.

C.2 Data and Summary Statistics

Figure C.1 shows the main time series that are used in the empirical analyses of Section 3.2. Apart from the apparent decrease in the level and volatility of inflation as well as inflation expectations, we see that expectations became more and more detached from actual inflation. First, consumer expectations seem to be biased on average in the most recent decades, as can be seen in the lower panel. While these expectations closely tracked inflation in the 70s and 80s, this is not the case anymore.[4] Second, professional forecasters' expectations seem to perform quite well on average. In the last twenty years, however, they barely react to actual changes in inflation anymore. Overall, these observations suggest that attention decreased in the last decades.

Figure C.1: Inflation and Inflation Expectations

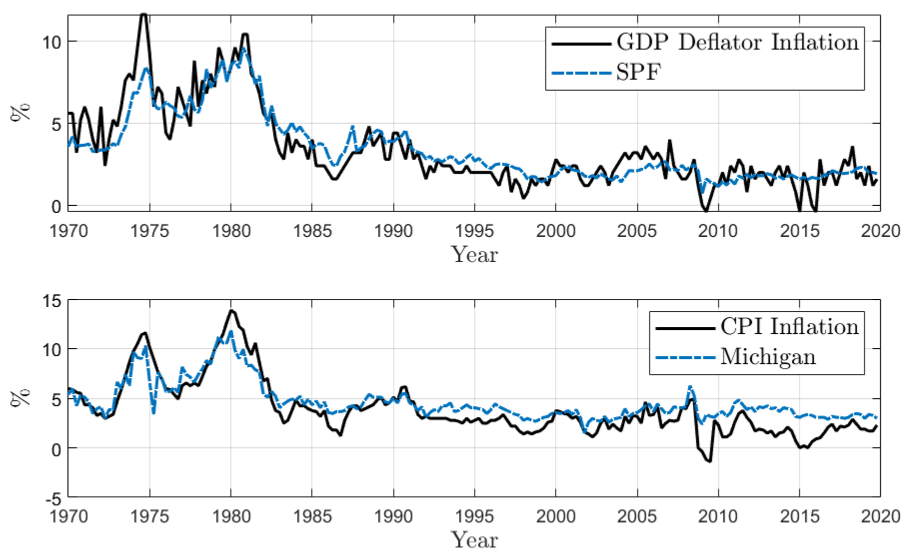

Note: This figure shows the raw time series of inflation, as well as survey expectations about future inflation. Everything is in annualized percentages.

Table C.1 shows the summary statistics, for the period before and after the 1990s, separately. For professional forecasters, the perceived persistence is higher than the actual one. This is especially the case when the actual persistence is relatively low, as was the case after 1990. Afrouzi et al. (2023) document a similar finding in an experimental setting. This might point towards lower attention since the 1990s. Note, that in the empirical analysis I account for changes in the perceived persistence.

[4] In the empirical analysis I account for this mean bias by including an intercept in the regressions.

Table C.1: Summary Statistics

	GDP Deflator Inflation		SPF Expectations	
	1968-1990	1990-2020	1968-1990	1990-2020
Mean (%)	5.44	2.00	5.18	2.15
Std. Dev. (%)	2.43	0.90	1.87	0.63
Persistence	0.84	0.55	0.93	0.92

	CPI Inflation		Consumer Expectations	
	1968-1990	1990-2020	1968-1990	1990-2020
Mean (%)	6.09	2.42	6.00	3.62
Std. Dev. (%)	3.00	1.26	2.17	0.68
Persistence	0.96	0.77	0.85	0.70

Note: This table shows the summary statistics of the data. The upper panel shows the statistics for the quarter-on-quarter GDP deflator inflation (left) and the corresponding inflation expectations from the Survey of Professional Forecasters (right). The lower panel shows the year-on-year CPI inflation (left) and the corresponding inflation expectations from the Survey of Consumers from the University of Michigan. All data are annualized.

C.3 Robustness and Additional Evidence

In this section, I show that the empirical results are robust along several dimensions.

Additional Data Sources

In Table C.2, I show how attention changed over time for different data sources. The first two columns show the results for the Greenbook forecasts, columns 3-4 for the Livingston Survey, and columns 5-6 and 7-8 are for CPI forecasts from the SPF instead of forecasts about the GDP deflator. As in the main text, I use two different estimators. First, the Blundell-Bond estimator (columns 5-6) and pooled OLS (columns 7-8). All standard errors are robust with respect to heteroskedasticity and serial correlation. We see that the main finding of lower attention in inflation expectations in the period after 1990 compared to the period before is robust to these changes in the data source and/or exact variable.

Table C.2: Regression Results of Equation (3.4)

	Greenbook		Livingston		SPF CPI BB		SPF CPI OLS	
	< 1990	≥ 1990	< 1990	≥ 1990	< 1990	≥ 1990	< 1990	≥ 1990
$\widehat{\gamma}$	0.39	0.24	0.28	0.17	0.36	0.23	0.17	0.13
s.e.	(0.0851)	(0.0715)	(0.0554)	(0.0624)	(0.1444)	(0.0328)	(0.0409)	(0.0142)
N	84	100	83	61	550	3,577	550	3,577

Note: This table shows the results from regression (3.4) for different data sources. The standard errors, reported in parentheses, are robust with respect to heteroskedasticity and serial correlation.

Different Sample Splits

Table 3.1 in the main body of the paper shows that attention to inflation declined by focusing on a sample split in 1990. To show that this is robust to the exact split point, Tables C.3 and C.4 show that the result holds when splitting the sample in 1985 or 1995, respectively. In fact, the decline in attention is even somewhat more pronounced when splitting the sample in 1985. This is in line with the theoretical prediction of the limited-attention model. Namely, the period between 1985 and 1990 was a period of relatively low and stable inflation compared to the period pre 1985 (see Figure C.1), and thus, a period in which the model would predict a relatively low level of attention.

Table C.3: Regression Results of Equation (3.4), pre 1985 vs. post 1985

	Professional Forecasters				Consumers			
	Blundell Bond		Pooled OLS		Averages		Median	
	< 1985	≥ 1985	< 1985	≥ 1985	< 1985	≥ 1985	< 1985	≥ 1985
$\widehat{\gamma}$	0.75	0.37	0.45	0.25	0.77	0.31	0.50	0.26
s.e.	(0.1247)	(0.0399)	(0.0403)	(0.0338)	(0.1688)	(0.0811)	(0.0955)	(0.0561)
N	1914	3887	1914	3887	64	140	27	140

Note: This table shows the results from regression (3.4) for professional forecasters (SPF) as well as for consumers. For the SPF, I use the Blundell and Bond (1998) (BB) estimator (first two columns), as well as pooled OLS (columns 3-4). For the Survey of Consumer, I consider average expectations (columns 5-6) and median expectations (columns 7-8). The standard errors, reported in parentheses, are robust with respect to heteroskedasticity and serial correlation.

Table C.4: Regression Results of Equation (3.4), pre 1995 vs. post 1995

	Survey of Professional Forecasters				Survey of Consumers			
	Blundell Bond		Pooled OLS		Averages		Median	
	< 1995	≥ 1995	< 1995	≥ 1995	< 1995	≥ 1995	< 1995	≥ 1995
$\widehat{\gamma}$	0.70	0.41	0.44	0.21	0.72	0.27	0.43	0.22
s.e.	(0.0907)	(0.0654)	(0.0379)	(0.0344)	(0.1473)	(0.0962)	(0.0819)	(0.0654)
N	2708	3093	2708	3093	104	100	67	100

Note: This table shows the results from regression (3.4) for professional forecasters (SPF) as well as for consumers. For the SPF, I use the Blundell and Bond (1998) (BB) estimator (first two columns), as well as pooled OLS (columns 3-4). For the Survey of Consumer, I consider average expectations (columns 5-6) and median expectations (columns 7-8). The standard errors, reported in parentheses, are robust with respect to heteroskedasticity and serial correlation.

Different Specifications of the BB Estimator

In the baseline estimation, reported in Table 3.1, I included all potential lags for the Blundell-Bond estimation. To show that the results are robust to this specification, I show in Table C.5 that for maximum lag lengths of 20 and 10 periods, the estimated attention parameter $\widehat{\gamma}$ is in all cases higher before 1990 compared to the period after 1990.

Table C.5: Different Maximum Lag Lengths

	All Lags		20 Lags		10 Lags	
	< 1990	≥ 1990	< 1990	≥ 1990	< 1990	≥ 1990
$\widehat{\gamma}$	0.70	0.41	0.74	0.51	0.84	0.69
s.e.	(0.1005)	(0.0522)	(0.1086)	(0.0632)	(0.1247)	(0.1127)
N	2235	3566	2235	3566	2235	3566

Note: This table shows the results from regression (3.4) for different numbers of lags included in the BB estimation. The standard errors, reported in parentheses, are robust with respect to heteroskedasticity and serial correlation.

Time Fixed Effects

To account for potential changes in trend inflation, I include time-fixed effects in regression (3.4). To do so, recall that (3.4) is given by

$$\pi^e_{t+1|t,i} = \beta_i + \beta_1 \pi^e_{t|t-1,i} + \beta_2 \left(\pi_t - \pi^e_{t|t-1,i}\right) + u_{i,t}. \tag{C.6}$$

To include time fixed effects, I first compute a period-specific persistence parameter, ρ_π. Note, that in (C.6), β_1 measures this persistence. Therefore, I subtract $\widehat{\rho}_\pi \pi^e_{t|t-1,i}$ from both sides and then to directly estimate γ, I further divide both sides by $\widehat{\rho}_\pi$:

$$\frac{\pi^e_{t+1|t,i} - \widehat{\rho}_\pi \pi^e_{t|t-1,i}}{\widehat{\rho}_\pi} = \delta_i + d_t + \gamma \left(\pi_t - \pi^e_{t|t-1,i}\right) + v_{i,t}, \tag{C.7}$$

where d_t captures time-fixed effects, $\delta_i = \frac{\beta_i}{\widehat{\rho}_\pi}$ and $v_{i,t} = \frac{u_{i,t}}{\widehat{\rho}_\pi}$. I do this transformation for the period before and after 1990 separately. Note, that this transformation also deals with the endogeneity problem explained in Section 3.2.

The estimated attention levels are 0.75 (s.e. 0.0327) for the period before 1990 and 0.61 (s.e. 0.0295) after 1990 if I use the first-order autocorrelation of expected inflation as my measure of ρ_π. If I use the estimate of β_1 from equation (C.7) as my measure of ρ_π, the estimated attention before the 1990s is 0.68 (s.e. 0.0252) and the one after the 1990s is 0.46 (s.e. 0.0242). Thus, we see that the decrease in attention is robust to controlling for time-fixed effects, even though the decline is somewhat muted.

When using the first-order autocorrelation of expected inflation as my measure of ρ_π, estimating equation (3.5) in this way, delivers a point estimate of 0.06 (s.e. 0.0111) that is statistically significant on all conventional significance levels. The estimate for ζ in regression (3.6) is 0.23 (s.e. 0.0220), statistically significant on all conventional significance levels. When using $\hat{\beta}_1$ from (C.7) as the measure of ρ_π, the point estimate of β in equation (3.5) is 0.06 (s.e. 0.0074) and the estimate of ζ in (3.6) is 0.29 (s.e. 0.0306), both statistically significant on all conventional levels of significance. Thus, the positive relationships between attention and volatility, as well as between attention and inflation persistence, are robust to controlling for time fixed effects.

Professional Forecasters in the Aggregate

When estimating attention of professional forecastors' average expectations instead of individual ones, we obtain a value of 0.24 (s.e. 0.0481) for the period before 1990 and of 0.09 (s.e. 0.0353) after 1990. Consistent with the main results, attention substantially decreased in recent decades and is about half after 1990 compared to before.

Estimating regression (3.5) on aggregate SPF data delivers a coefficient of 0.15 (p-value of 0.000) and the estimate of ζ in regression (3.6) is 0.69 (p-value of 0.000). Thus, the results reported in the main text are robust.

Joint Regressions

Instead of running regressions (3.5) and (3.6) separately, I estimate

$$\widehat{\gamma}_t = \alpha + \beta \widehat{\sigma}_{\pi,t} + \zeta \widehat{\rho}_{\pi,t} + u_t. \tag{C.8}$$

Table C.6 shows that the results are robust to this change in specification.

Table C.6: Attention, Inflation Volatility and Inflation Persistence

	Survey of Professional Forecasters		Michigan Survey
Estimator	Blundell-Bond	Pooled OLS	OLS
$\widehat{\beta}$	0.04***	0.05***	0.06***
s.e.	(0.0153)	(0.0128)	(0.0150)
$\widehat{\zeta}$	0.59***	0.65***	0.31***
s.e.	(0.0597)	(0.0499)	(0.0772)
N	165	165	163

Note: This table shows the results of regression (C.8). Standard errors are robust with respect to heteroskedasticity. *** : p-value < 0.01, ** : p-value < 0.05, * : p-value < 0.1.

Controlling for Average Inflation

A potential confounder in regression (C.8) above is the average level of inflation. Thus, I now control for the average level of inflation, computed as the average inflation rate in the respective 10-year window. In particular, I run the following regression

$$\widehat{\gamma}_t = \alpha + \beta\widehat{\sigma}_{\pi,t} + \zeta\widehat{\rho}_{\pi,t} + \omega\widehat{\overline{\pi}}_t + u_t, \tag{C.9}$$

where $\widehat{\overline{\pi}}_t$ is the estimated average inflation rate. Table C.7 reports the results for the professional forecasters. We see that the volatility and the persistence of inflation are positively related with attention and that these relationships are statistically significant even when controlling for the average level of inflation. The average level of inflation, on the other hand, does not have a positive, statistically-significant, effect on the estimated attention when we control for the volatility and persistence of inflation. These results are consistent with the underlying theoretical model.

Table C.7: Controlling for Average Inflation

	Survey of Professional Forecasters	
Estimator	Blundell-Bond	Pooled OLS
$\widehat{\beta}$	0.13***	0.06***
s.e.	(0.0264)	(0.0200)
$\widehat{\zeta}$	0.69***	0.82***
s.e.	(0.1153)	(0.0824)
$\widehat{\omega}$	−0.02**	0.01
s.e.	(0.0095)	(0.0082)
N	165	165

Note: This table shows the results of regression (C.9). Standard errors are robust with respect to heteroskedasticity. *** : p-value < 0.01, ** : p-value < 0.05, * : p-value < 0.1.

Quasi-Panel of Consumers

The Survey of Consumers does not follow consumers over time. Therefore, I could not allow for individual-specific fixed effects but rather consider average and/or median inflation expectations. I now group the survey respondents into four groups, based on their income. The SoC provides data on this starting in the last quarter of 1979.

Table C.8 shows the results. The first two columns report the results for the split point in 1990, and the third and fourth column for the split point in 1995. We see that the estimated attention levels using this quasi panel are similar to the ones obtained using average expectations (Table 3.1).

Table C.8: Regression Results of Equation (3.4), Quasi-Panel

	Survey of Consumers			
	< 1990	≥ 1990	< 1995	≥ 1995
$\widehat{\gamma}$	0.77	0.33	0.70	0.29
s.e.	(0.0933)	(0.0263)	(0.1078)	(0.0289)
N	160	480	240	400

Note: This table shows the results from regression (3.4), estimated using the Blundell and Bond (1998) estimator, for consumers grouped into four groups, based on their income. The standard errors, reported in parentheses, are robust with respect to heteroskedasticity and serial correlation.

Table C.9 shows the results of regressions (3.5) and (3.6) (first column), as well as of the joint regression (C.8), using this quasi panel of consumers. We see that the results are robust and that there is indeed a significantly positive relation between attention and inflation volatility, as well as between attention and inflation persistence.

Table C.9: Attention, Inflation Volatility and Inflation Persistence

	Survey of Consumers	
Estimator	Separate	Joint
$\widehat{\beta}$	0.13^{***}	0.13^{***}
s.e.	(0.0106)	(0.0126)
$\widehat{\zeta}$	0.20^{***}	0.12^{***}
s.e.	(0.0787)	(0.0620)
N	121	121

Note: This table shows the results of regressions (3.5), (3.6) (first column) and (C.8) (second column) using a quasi panel of consumers. The attention parameters have been estimated using the BB-estimator. Standard errors are robust with respect to heteroskedasticity. *** : p-value < 0.01, ** : p-value < 0.05, * : p-value < 0.1.

Volatility and Persistence of Inflation Expectations

Table C.10 shows the results of regressions (3.5) and (3.6) using the volatility and persistence of inflation expectations instead of actual inflation as independent variables. Standard errors are robust with respect to heteroskedasticity.

Table C.10: Attention, Expected Inflation Volatility and Persistence

	Survey of Professional Forecasters		Michigan Survey
Estimator	Blundell-Bond	Pooled OLS	OLS
$\widehat{\beta}$	0.14^{***}	0.16^{***}	0.13^{***}
s.e.	(0.0153)	(0.0098)	(0.0172)
$\widehat{\zeta}$	1.06^{***}	1.21^{***}	0.23^{***}
s.e.	(0.1272)	(0.0838)	(0.0685)
N	165	165	163

Note: This table shows the results of regressions (3.5) and (3.6) using the volatility and persistence of inflation expectations instead of actual inflation as dependent variables. Standard errors are robust with respect to heteroskedasticity. *** : p-value < 0.01, ** : p-value < 0.05, * : p-value < 0.1.

Window Length

As predicted by the underlying model of optimal information acquisition, I showed that there is indeed a positive relationship between attention to inflation and inflation volatility, as well as between attention and inflation persistence. In the baseline specification, I relied on a rolling-window approach in which every window was 10 years. Tables C.11 and C.12 show that these results are robust to using different window lengths, namely 5 and 15 years.

Table C.11: Attention, Inflation Volatility and Inflation Persistence: Five Year Window

	Survey of Professional Forecasters		Michigan Survey
Estimator	Blundell-Bond	Pooled OLS	OLS
$\widehat{\beta}$	-0.01	0.06***	0.13***
s.e.	(0.1643)	(0.0185)	(0.0411)
$\widehat{\zeta}$	0.73	0.44***	0.40***
s.e.	(0.6731)	(0.0551)	(0.1547)
N	185	185	183

Note: This table shows the results of regression (C.8) using windows of 5 years each. Standard errors are robust with respect to heteroskedasticity. *** : p-value < 0.01, ** : p-value < 0.05, * : p-value < 0.1.

Table C.12: Attention, Inflation Volatility and Inflation Persistence: Fifteen Year Window

	Survey of Professional Forecasters		Michigan Survey
Estimator	Blundell-Bond	Pooled OLS	OLS
$\widehat{\beta}$	0.01	0.01	0.07***
s.e.	(0.0124)	(0.0115)	(0.0136)
$\widehat{\zeta}$	0.90***	1.00***	0.43***
s.e.	(0.0603)	(0.0552)	(0.0706)
N	145	145	143

Note: This table shows the results of regression (C.8) using windows of 15 years each. Standard errors are robust with respect to heteroskedasticity. *** : p-value < 0.01, ** : p-value < 0.05, * : p-value < 0.1.

Attention over Time

Figure C.2 shows the estimated attention levels, γ, (black-solid line) from the SPF consensus forecasts, together with the volatility of GDP deflator inflation (blue-dashed lines). We clearly see the aforementioned decrease in attention over time, as well as the positive correlation of attention and inflation volatility.

Figure C.2: Attention and Inflation Volatility over Time

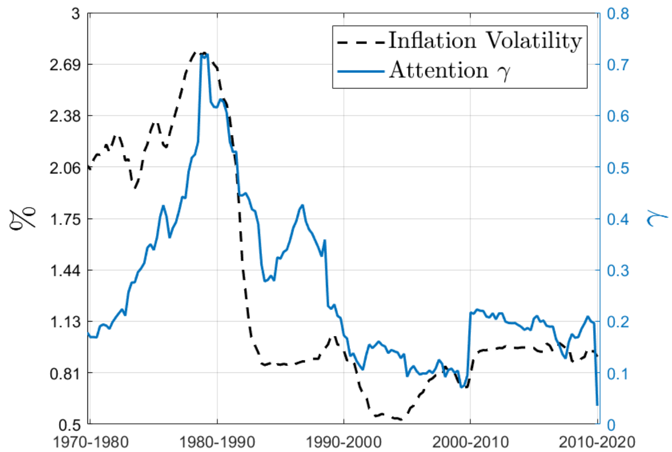

Notes: This figure shows the estimated attention levels, γ, (black-solid line) from the SPF consensus forecasts, together with the volatility of GDP deflator inflation (blue-dashed lines).

C.3.1 News and Book Coverage of Inflation

In this section, I provide complementary evidence to the one presented in Section 3.2 based on news coverage. Figure C.3 shows the relative frequency of the word "inflation" among all words in two major U.S. newspapers (blue-dashed lines), the New York Times (left panel) and the Washington Post (right panel), together with the annual U.S. CPI inflation (black-solid lines). It is evident that news coverage is higher in times of high and volatile inflation as was the case during the 1970s and early 1980s. Moreover, the figure suggests that the public's attention to inflation—proxied here by news coverage—has not always been as low as in recent years, but declined over time.

In Figure C.4, we see that a similar picture emerges when looking at the coverage of "inflation" in books, according to *Google Books Ngram Viewer*. In the left panel, we see that "inflation" is covered more frequently in English books written in times of high inflation. But this is not simply a U.S. phenomenon. To see this, I show the same statistic for books written in Spanish for the word "inflación". To contrast this with inflation, the black solid line shows the average inflation (in logs) of the four largest Spanish-speaking countries, weighted by their 2020 population size. These are Argentina, Colombia, Mexico and Spain. Again, we observe that attention to inflation—measured by book coverage—is higher in times of high and volatile inflation.

Figure C.3: News Coverage of Inflation

(a) New York Times (b) Washington Post

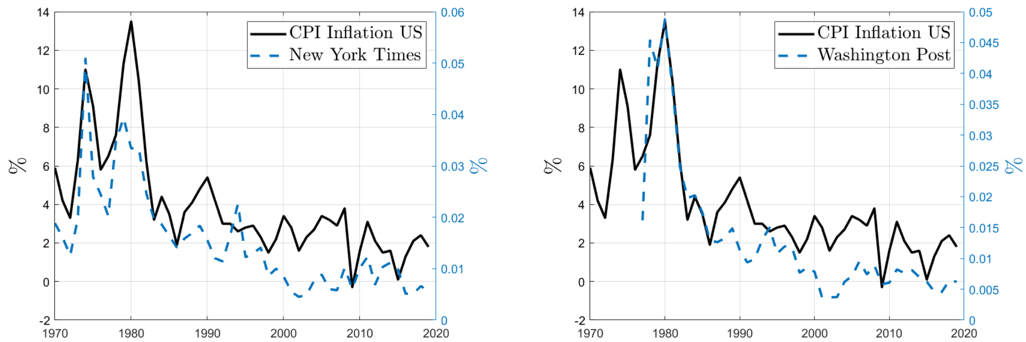

Notes: This figure shows the relative frequency (blue dashed lines, right axis) of the word "inflation" in the New York Times (left) and the Washington Post (right). The black solid line shows annual U.S. CPI inflation (left axis).

Figure C.4: Book Coverage of Inflation

(a) Google Books English (b) Google Books Spanish

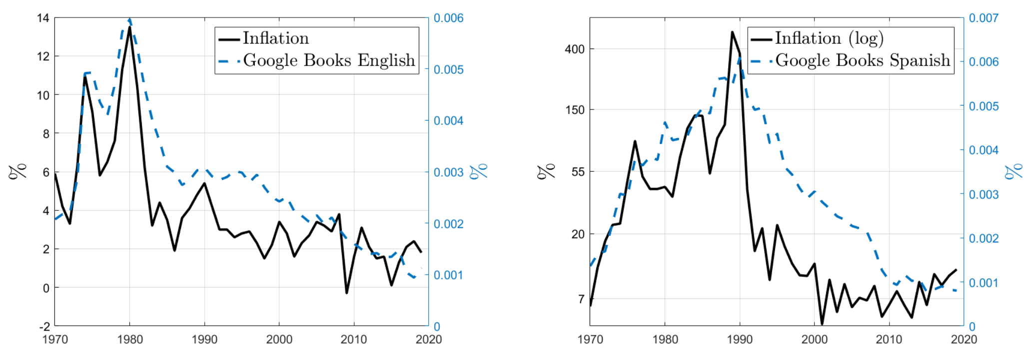

Notes: The blue dashed lines show the frequency of the words "inflation" and "inflación", respectively, in English and Spanish books, according to Google Books Ngram Viewer. The black solid line shows the corresponding inflation rates.

C.4 Forward Guidance and Attention

To see how lower attention weakens the effectiveness of forward guidance, consider the following stylized experiment.[5] The economy is hit by a negative natural rate shock in period $t = 0$ that pushes the nominal interest rate to the effective lower bound, i.e., $r_0^n < 0$ and $i_0 = -\underline{i}$. In $t = 1$, the natural rate returns to its steady state value and stays there indefinitely, $r_t^n = 0$ for all $t \geq 1$. From period $t = 2$ onwards, the output gap, and the real rate are back at their steady states, $y_t^{gap} = 0$ and $i_t - \pi_{t+1|t}^e = 0$ for all $t \geq 2$.

To model forward guidance, the real rate is assumed to be below the natural rate in $t = 1$. To make it comparable across different degrees of attention, I impose that

$$r_1 \equiv i_1 - \pi_{2|1}^e < 0$$

is the same for all γ and known in advance.[6] Hence, forward guidance here means to announce a certain value for the real rate. I discuss the implications of forward guidance via the nominal rate in section C.4.1. In the following, I assume that $(-\underline{i} - r_0^n + r_1)$ is negative, which means that the announced policy, captured by $r_1 < 0$, *makes up* for the binding lower bound in $t = 0$, captured by $-\underline{i} - r_0^n > 0$.

Given the real rate r_1 and the fact that $y_2^{gap} = 0$, the Euler equation in $t = 1$ determines the output gap in period 1 as

$$y_1^{gap} = -\varphi\left(r_1\right) > 0. \tag{C.10}$$

Equation (C.10) captures the *make-up policy:* by keeping the real rate below the natural rate, output is (expected to be) above potential after the lower-bound constraint stops to be binding.

In $t = 0$, the ELB binds and the natural rate is negative. Thus, the Euler equation in $t = 0$ yields

$$y_0^{gap} = \underbrace{-\varphi\left(r_1\right)}_{=E_0 y_1^{gap}} - \varphi\left(-\underline{i} - \pi_{1|0}^e - r_0^n\right).$$

Substituting the law of motion for inflation expectations

$$\pi_{1|0}^e = (1-\gamma)\pi_{0|-1}^e + \gamma\pi_0,$$

into the Phillips Curve

$$\pi_0 = \frac{\beta}{1-\beta\gamma}(1-\gamma)\pi_{0|-1}^e + \frac{\kappa}{1-\beta\gamma} y_0^{gap}$$

[5] For clarity, I focus on the case $\rho_\pi = 1$. I discuss the general case with $\rho_\pi < 1$ in section C.4.1.

[6] This is different to Angeletos and Lian (2018), where private agents are uncertain about future policies.

yields an expression for inflation expectations:

$$\pi^e_{1|0} = \frac{1-\gamma}{1-\beta\gamma}\pi^e_{0|-1} + \frac{\kappa\gamma}{1-\beta\gamma}y^{gap}_0.$$

Putting everything together, we arrive at the following result.

Proposition 8. *The output gap in the period when the shock hits, $t = 0$, is given by*

$$y^{gap}_0 = -\frac{\varphi\left(1-\beta\gamma\right)}{1-\gamma(\beta+\varphi\kappa)}\left[-\underline{i} - r^n_0 + r_1\right] + \frac{\varphi(1-\gamma)}{1-\gamma(\beta+\varphi\kappa)}\pi^e_{0|-1} \tag{C.11}$$

and inflation in $t = 0$ is given by

$$\pi_0 = -\frac{\kappa\varphi}{1-\gamma(\beta+\varphi\kappa)}\left[-\underline{i} - r^n_0 + r_1\right] + (1-\gamma)\left[\frac{\beta}{1-\beta\gamma} + \frac{\varphi}{1-\gamma(\beta+\varphi\kappa)}\right]\pi^e_{0|-1}. \tag{C.12}$$

Proposition 8 captures the effectiveness of forward guidance on the output gap and inflation in the period when the shock hits. Assuming $(1-\gamma\left(\beta+\varphi\kappa\right))$ is positive makes sure that forward guidance has a stimulating effect on output and inflation in $t = 0$. Proposition 8 captures several channels how a change in attention affects the economy's response to forward guidance, which I discuss in the following two corollaries.

Corollary 5. *Lower attention weakens*

(i) the negative effects of the shock,

(ii) the positive effects of forward guidance,

(iii) the positive effects of a decrease in the lower bound $-\underline{i}$

on the output gap and inflation.

Corollary 5 follows from the fact that the terms $\frac{\varphi(1-\beta\gamma)}{1-\gamma(\beta+\varphi\kappa)}$ and $\frac{\kappa\varphi}{1-\gamma(\beta+\varphi\kappa)}$ in front of $[-\underline{i} - r^n_0 + r_1]$ are both increasing in γ. Points (*i*) and (*ii*) capture the main trade off of lower attention. While lower attention has a stabilizing effect via more anchored inflation expectations (point (*i*)), it renders forward guidance less effective (point (*ii*)). The reason why forward guidance becomes less effective as attention declines is because inflation expectations increase less in response to the announced policy, and thus, the real rate remains higher. Point (*iii*) illustrates an additional drawback of lower attention. A reduction of the effective lower bound, $-\underline{i}$, is less stimulating if agents in the economy are less attentive. Thus, going from a zero lower bound to a lower bound in negative territory, as conducted in several advanced economies over the last ten years, becomes less effective in terms of stimulating output and inflation if the public is inattentive. Away from the lower bound, point (*iii*) implies that the effectiveness of conventional monetary policy via the nominal interest rate becomes less effective as attention declines.

How attention matters for the transmission of prior inflation expectations on the output gap and inflation is ambiguous, as the following Corollary shows.

Corollary 6. *Lower attention*

(i) weakens the positive effect of higher prior inflation beliefs, $\pi^e_{0|-1}$, on the output gap if and only if,

$$(\beta + \varphi\kappa) > 1, \tag{C.13}$$

(ii) weakens the positive effect of higher prior inflation beliefs on inflation if and only if

$$\frac{\beta\,(\beta-1)}{(1-\beta\gamma)^2} + \frac{\varphi\left((\beta+\varphi\kappa)-1\right)}{\left(1-\gamma(\beta+\varphi\kappa)\right)^2} > 0. \tag{C.14}$$

Overall, the role of attention for the effects of higher prior beliefs on output and inflation is ambiguous. This is mainly the case because, on the one hand, lower attention implies that agents put more weight on their prior beliefs. On the other hand, as discussed previously, lower attention leads to more stable inflation overall, thus, weakening the effects of prior beliefs. This can also be seen in the discussion of the Phillips Curve, see Proposition 7.

Given the calibration in Table 3.4, conditions (C.13) and (C.14) both hold for all $\gamma < 0.99$. The effects of changes in γ, however, are numerically small. Thus, an increase in the average inflation rate—which increases average prior beliefs—is a promising monetary instrument to combat the loss of control via forward guidance as attention declines. By *ex-ante* increasing the average inflation rate, the policymaker not only supports higher inflation expectations and thus, lower real rates for a given nominal rate, but also gains additional policy space through the increase in the average nominal rate. Higher average inflation, however, is also costly.

C.4.1 Extensions

I now show that all the results go through when relaxing the assumption that $\rho_\pi = 1$ and also discuss how forward guidance via the nominal (instead of the real) interest rate changes the results and I also allow for attention heterogeneity across firms and households. We consider the same stylized experiment but now the law of motion for inflation expectations is given by

$$\pi^e_{1|0} = (1-\rho_\pi)\bar{\pi} + \rho_\pi(1-\gamma)\pi^e_{0|-1} + \rho_\pi\gamma\pi_0,$$

which can be substituted into the Phillips Curve:

$$\pi_0 = \frac{\beta}{1-\beta\rho_\pi\gamma}\left((1-\rho_\pi)\bar{\pi} + \rho_\pi(1-\gamma)\pi^e_{0|-1}\right) + \frac{\kappa}{1-\beta\rho_\pi\gamma} y^{gap}_0.$$

Thus, inflation expectations are given by

$$\pi^e_{1|0} = \frac{1-\rho_\pi}{1-\beta\rho_\pi\gamma}\bar{\pi} + \frac{\rho_\pi(1-\gamma)}{1-\beta\rho_\pi\gamma}\pi^e_{0|-1} + \frac{\kappa\rho_\pi\gamma}{1-\beta\rho_\pi\gamma}y^{gap}_0.$$

Putting everything together, we arrive at the following Proposition.

Proposition 9. *The output gap in the period when the shock hits, $t = 0$, is given by*

$$\begin{aligned} y^{gap}_0 &= -\frac{\varphi\left(1-\beta\rho_\pi\gamma\right)}{1-\rho_\pi\gamma(\beta+\varphi\kappa)}\left[-\underline{i} - r^n_0 + r_1\right] \\ &+ \frac{\varphi}{1-\rho_\pi\gamma(\beta+\varphi\kappa)}\left[(1-\rho_\pi)\bar{\pi} + \rho_\pi(1-\gamma)\pi^e_{0|-1}\right] \end{aligned}$$

and inflation in $t = 0$ is given by

$$\begin{aligned} \pi_0 = &-\frac{\kappa\varphi}{1-\rho_\pi\gamma(\beta+\varphi\kappa)}\left[-\underline{i} - r^n_0 + r_1\right] + (1-\rho_\pi)\left[\frac{\beta}{1-\beta\rho_\pi\gamma} + \frac{\varphi}{1-\rho_\pi\gamma(\beta+\varphi\kappa)}\right]\bar{\pi} \\ &+ \rho_\pi(1-\gamma)\left[\frac{\beta}{1-\beta\rho_\pi\gamma} + \frac{\varphi}{1-\rho_\pi\gamma(\beta+\varphi\kappa)}\right]\pi^e_{0|-1}. \end{aligned} \tag{C.15}$$

Proposition 9 captures the effectiveness of forward guidance on the output gap and inflation in the period when the shock hits. The assumption that $(1-\rho_\pi\gamma\left(\beta+\varphi\kappa\right))$ is positive, makes sure that forward guidance, i.e, a lower r_1 has a stimulating effect on output and inflation in $t = 0$. Proposition 9 captures several channels how a change in attention affects the economy's response to forward guidance, which I now collect in a series of corollaries.

Corollary 7. *Lower attention*

(i) weakens the negative effect of the shock on impact,

(ii) weakens the effects of forward guidance on the output gap and inflation,

(iii) weakens the stimulative effects of a decrease in the lower bound $-\underline{i}$.

Corollary 7 follows from the fact that the terms $\frac{\varphi(1-\beta\rho_\pi\gamma)}{1-\rho_\pi\gamma(\beta+\varphi\kappa)}$ and $\frac{\kappa\varphi}{1-\rho_\pi\gamma(\beta+\varphi\kappa)}$ in front of $[-\underline{i} - r^n_0 + r_1]$ are both increasing in γ. Points (*i*) and (*ii*) capture the main trade off of lower attention. While lower attention has a stabilizing effect via more anchored inflation expectations (point (*i*)), it renders forward guidance less effective (point (*ii*)). The reason why forward guidance becomes less effective as attention declines is because inflation expectations increase less in response to the announced policy, and thus, the real rate remains higher. Point (*iii*) illustrates an additional drawback of lower attention. A reduction of the effective lower bound, $-\underline{i}$, is less stimulating if agents in the economy are less attentive. Thus, going from a zero lower bound to a lower bound in negative territory, as conducted in several advanced economies over the last ten years, becomes less effective in terms of stimulating output and inflation if the public is inattentive. Note, that a decrease in the perceived inflation persistence, ρ_π, has the exact same implications as a decrease in γ.

The next corollary discusses how changes in attention affect the role of long-run inflation beliefs on the output gap and inflation.

Corollary 8. *Lower attention weakens the positive effects of higher long-run inflation beliefs $\bar{\pi}$ on output and inflation,*

Corollary 8 says that higher long-run beliefs have a positive effect on inflation and the output gap, but lower attention weakens these effects. However, as long as $\gamma(\beta + \varphi\kappa) < 1$, a higher ρ_π mutes the effects of $\bar{\pi}$ on the output gap. Since this condition is usually satisfied and because ρ_π is in general close to 1, the role of high long-run inflation beliefs is quite weak. In the limit case $\rho_\pi \to 1$, long-run beliefs become irrelevant.

How attention matters for the transmission of prior inflation expectations on the output gap and inflation is ambiguous, as the following Corollary shows.

Corollary 9. *Lower attention*

(i) weakens the positive effect of higher prior inflation beliefs, $\pi^e_{0|-1}$, on the output gap if and only if,

$$\rho_\pi(\beta + \varphi\kappa) > 1, \tag{C.16}$$

(ii) weakens the positive effect of higher prior inflation beliefs on inflation if and only if

$$\frac{\rho_\pi\beta(\rho_\pi\beta - 1)}{(1 - \beta\rho_\pi\gamma)^2} + \frac{\rho_\pi\varphi(\rho_\pi(\beta + \varphi\kappa) - 1)}{(1 - \rho_\pi\gamma(\beta + \varphi\kappa))^2} > 0. \tag{C.17}$$

Overall, the role of attention for the effects of higher prior beliefs on output and inflation is ambiguous. This is mainly the case because, on the one hand, lower attention implies that agents put more weight on their prior beliefs. On the other hand, as discussed previously, lower attention leads to more stable inflation overall, thus, weakening the effects of prior beliefs. This can also be seen in the discussion of the Phillips Curve, see Proposition 7.

Given the calibration in Table 3.4, conditions (C.16) and (C.17) both hold. The effects of changes in γ, however, are numerically small. Thus, an increase in the average inflation rate—which also increases average prior beliefs—is a promising monetary instrument to combat the loss of control via forward guidance as attention declines. By *ex-ante* increasing the average inflation rate, the policymaker not only supports higher inflation expectations and thus, lower real rates for a given nominal rate, but also gains additional policy space through the increase in the average nominal rate. Higher average inflation, however, is also costly. In the analysis of optimal policy, later on, I will explore this trade off and characterize the optimal inflation target for different levels of attention.

Forward Guidance via Nominal Interest Rates

So far, forward guidance was characterized as a promise to keep the *real* rate low. Now, assume that forward guidance is conducted via promising lower *nominal* rates instead. Thus, i_1 will be fixed across different γ. For simplicity, I focus on the case with $\rho_\pi = 1$ and $\pi^e_{0|-1} = 0$. It follows from the Euler equation in $t = 1$ that

$$y_1^{gap} = -\varphi\left(i_1 - (1-\gamma)\,\gamma\pi_0 - \gamma\pi_1\right).$$

The Phillips Curve in $t = 1$ yields

$$\pi_1 = \frac{(1-\gamma)\gamma}{1-\beta\gamma}\pi_0 + \frac{\kappa}{1-\beta\gamma}y_1^{gap},$$

so that we get an expression for y_1^{gap} in terms of π_0:

$$y_1^{gap} = -\frac{\varphi\,(1-\beta\gamma)}{1-\gamma\,(\beta+\varphi\kappa)}i_1 + \varphi(1-\gamma)\gamma\frac{1+\gamma(1-\beta)}{1-\gamma(\beta+\varphi\kappa)}\pi_0. \tag{C.18}$$

Given $\pi^e_{1|0} = \gamma\pi_0$, the Phillips Curve in $t = 0$ yields

$$\pi_0 = \frac{\kappa}{1-\beta\gamma}y_0^{gap},$$

and hence, $\pi^e_{1|0} = \frac{\kappa\gamma}{1-\beta\gamma}y_0^{gap}$. Plugging this into the Euler equation in $t = 0$ gives

$$y_0^{gap} = \mathbb{E}_0 y_1^{gap} - \varphi\left(-\underline{i} - \frac{\kappa\gamma}{1-\beta\gamma}y_0^{gap} - r_0^n\right).$$

Solving for y_0^{gap} leads to the following Lemma.

Lemma 10. *Forward guidance via the nominal interest rate yields the following output gap*

$$y_0^{gap} = A_1\left[-\frac{\varphi\,(1-\beta\gamma)}{1-\gamma\,(\beta+\varphi\kappa)}i_1 - \varphi\left(-\underline{i} - r_0^n\right)\right], \tag{C.19}$$

and inflation

$$\pi_0 = \frac{\kappa}{1-\beta\gamma}A_1\left[-\frac{\varphi\,(1-\beta\gamma)}{1-\gamma\,(\beta+\varphi\kappa)}i_1 - \varphi\left(-\underline{i} - r_0^n\right)\right], \tag{C.20}$$

where

$$A_1 \equiv \frac{1}{1-\varphi(1-\gamma)\gamma\frac{1+\gamma(1-\beta)}{1-\gamma(\beta+\varphi\kappa)}\frac{\kappa}{1-\beta\gamma} - \frac{\varphi\kappa\gamma}{1-\beta\gamma}}. \tag{C.21}$$

Given the calibration in Table 3.4, A_1 is positive and increasing in γ. Thus, promising lower future nominal interest rates can indeed stimulate the economy. But similar to the case in which the policy maker commits to a certain future *real* rate, forward guidance becomes less effective when agents are less attentive. In fact, all three results from Corollary 7 go through.

Recall equation (C.18):

$$y_1^{gap} = -\frac{\varphi\,(1-\beta\gamma)}{1-\gamma\,(\beta+\varphi\kappa)} i_1 + \varphi(1-\gamma)\gamma\frac{1+\gamma(1-\beta)}{1-\gamma(\beta+\varphi\kappa)}\pi_0.$$

Note, that the first term becomes less negative as γ declines. Given the calibration in Table 3.4, also the second term decreases as attention declines. Thus, for a given π_0, a particular i_1 has weaker effects on the output gap in $t = 1$ at lower levels of attention. Since lower attention also weakens the positive effects of forward guidance on π_0, the output gap (and inflation) stay lower also in $t = 1$.

Since inflation in $t = 0$ and $t = 1$ is lower at smaller values of γ, also $\pi^e_{2|1}$ will be lower and thus, for a given nominal rate i_1, the *real* rate, $r_1 \equiv i_1 - \pi^e_{2|1}$, will be higher. Hence, to achieve a certain forward guidance in terms of the real interest rate, the promise in terms of the nominal rate needs to be larger when firms and households are inattentive. Combining this with the findings on the effectiveness of forward guidance via the *real* rate (Proposition 9) shows how lower attention renders forward guidance less powerful *even though* the promise in terms of the *nominal rate* is stronger.

Heterogeneous Attention

So far, I assumed that firms and households are equally attentive. But what if firms and households differ in their attention to inflation? Let us denote firms' attention by γ_F and households' attention by γ_H with $\gamma_F \neq \gamma_H$. For clarity, I focus on the case with $\rho_\pi = 1$ and $\pi^{e,j}_{0|-1} = 0$ for $j \in \{F, H\}$.

Lemma 11. *With heterogeneous attention to inflation, the output gap in $t = 0$ is given by*

$$y_0^{gap} = \frac{-\varphi\,(1-\beta\gamma_F)}{1-\beta\gamma_F-\kappa\varphi\gamma_H}\left[-\underline{i}+r_1-r_0^n\right], \tag{C.22}$$

and inflation by

$$\pi_0 = \frac{-\varphi\kappa}{1-\beta\gamma_F-\kappa\varphi\gamma_H}\left[-\underline{i}+r_1-r_0^n\right], \tag{C.23}$$

where $r_1 \equiv i_1 - \pi^{e,H}_{2|1}$ is the real rate given the households' expectations.

Lemma 11 shows that a similar result as in Corollary 7 holds under heterogeneous attention levels.

Corollary 12. *Lower attention of either firms or households*

(i) weakens the negative effect of the shock on the output gap and inflation on impact,

(ii) weakens the effects of forward guidance on the output gap and inflation,

(iii) weakens the stimulative effects of a decrease in the lower bound $-\underline{i}$ on the output gap and inflation.

The parts concerning the output gap in Corollary 12 follow because the term in front of the brackets in equation (C.22) becomes more negative as either of $\{\gamma_F, \gamma_H\}$ increases:

$$\frac{\partial\left[\frac{-\varphi(1-\beta\gamma_F)}{1-\beta\gamma_F-\kappa\varphi\gamma_H}\right]}{\partial\gamma_F} = -\frac{\beta\kappa\varphi^2\gamma_H}{(1-\beta\gamma_F-\kappa\varphi\gamma_H)^2} < 0$$

$$\frac{\partial\left[\frac{-\varphi(1-\beta\gamma_F)}{1-\beta\gamma_F-\kappa\varphi\gamma_H}\right]}{\partial\gamma_H} = -\frac{\varphi^2(1-\beta\gamma_F)}{(1-\beta\gamma_F-\kappa\varphi\gamma_H)^2} < 0,$$

and the parts concerning inflation because the term $\frac{-\varphi\kappa}{1-\beta\gamma_F-\kappa\varphi\gamma_H}$ in equation (C.23) becomes more negative as either of $\{\gamma_F, \gamma_H\}$ increases, too.

Thus, if either firms or households (or both) become less attentive, forward guidance becomes less effective. In fact, the two degrees of attention reinforce each other, as the following Corollary shows.

Corollary 13. *Lower levels of households' attention to inflation weaken the effectiveness of forward guidance, especially when firms' attention to inflation is low, and vice-versa.*

To see this, note that

$$\frac{\partial^2\left[-\frac{\varphi\kappa}{1-\beta\gamma_F-\kappa\varphi\gamma_H}\right]}{\partial\gamma_F\partial\gamma_H} = \frac{-2\varphi^2\kappa^2\beta}{(1-\beta\gamma_F-\kappa\varphi\gamma_H)^3} < 0,$$

$$\frac{\partial^2\left[\frac{-\varphi(1-\beta\gamma_F)}{1-\beta\gamma_F-\kappa\varphi\gamma_H}\right]}{\partial\gamma_F\partial\gamma_H} = \frac{-\beta\kappa\varphi^2\left[1-\beta\gamma_F+\kappa\varphi\gamma_H\right]}{(1-\beta\gamma_F-\kappa\varphi\gamma_H)^3} < 0.$$

C.5 Proof of Proposition 7

Proof. The New Keynesian Phillips Curve is given by

$$\pi_t = \beta \pi^e_{t+1|t} + \kappa y^{gap}_t + u_t.$$

Substituting

$$\pi^e_{t+1|t} = \pi^e_{t|t-1} + \gamma \left(\pi_t - \pi^e_{t|t-1}\right)$$

for $\pi^e_{t+1|t}$ yields

$$\begin{aligned}
&\pi_t = \beta \left(\pi^e_{t|t-1} + \gamma \left(\pi_t - \pi^e_{t|t-1}\right)\right) + \kappa y^{gap}_t + u_t \\
\Leftrightarrow &\pi_t(1-\beta\gamma) = \beta \pi^e_{t|t-1} (1-\gamma) + \kappa y^{gap}_t + u_t \\
\Leftrightarrow &\pi_t = \frac{\beta \pi^e_{t|t-1} (1-\gamma) + \kappa y^{gap}_t + u_t}{(1-\beta\gamma)} \\
\Leftrightarrow &\pi_t = \frac{\beta (1-\gamma)}{(1-\beta\gamma)} \pi^e_{t|t-1} + \frac{\kappa}{(1-\beta\gamma)} y^{gap}_t + \frac{u_t}{(1-\beta\gamma)}.
\end{aligned}$$

Now, taking derivatives with respect to y^{gap}_t, u_t, and $\pi^e_{t|t-1}$, respectively, yields the results (i), (ii), and (iii). ■

C.6 Additional Numerical Results

C.6.1 Non-Rational Output Gap Expectations

During most of the analysis, I assumed that agents have rational expectations about the output gap even though they are inattentive to inflation. I now relax this assumption and instead assume that output gap expectations are given by

$$y_{t+1|t}^{gap,e} = y_{t|t-1}^{gap,e} + \gamma^{y}\left(y_{t}^{gap} - y_{t|t-1}^{gap,e}\right). \tag{C.24}$$

With these expectations, the aggregate IS equation is given by

$$y_{t}^{gap} = y_{t+1|t}^{gap,e} - \varphi\left(i_t - \pi_{t+1|t}^{e} - r_{t}^{n}\right), \tag{C.25}$$

whereas the Phillips Curve and the Taylor rule remain unchanged. Following Jørgensen and Lansing (2021), I assume that γ^{y} is about a third of γ and thus, set them to $\gamma^{y} = 0.1$ and $\gamma = 0.3$.

Figure C.5 shows the impulse response functions of the main variables in this economy after a negative three-standard deviation natural rate shock. We see that the inflation-attention traps get exacerbated. The reason is that now make-up policies are even less effective because not only inflation expectations are backward looking but also output gap expectations. Thus, even though there is interest-rate smoothing in the Taylor rule which features some form of make-up policy, this is not effective in stimulating expectations and thus, the economy remains stuck at the ELB even longer. Furthermore, inflation, inflation expectations, and now also the output gap stay below their initial values very persistently. Overall, this exercise illustrates that the results reported in the main text might be even starker if we allow for non-rational output gap expectations.

C.6.2 Different Taylor Rule

To show that the exact specification of the Taylor rule is not essential for the occurence of inflation-attention traps, Figure C.6 shows the impulse-response functions of the nominal interest rate, inflation, inflation expectations and the output gap for the model in which the Taylor rule absent the ELB is given by

$$i_t = 1.5\pi_t. \tag{C.26}$$

Figure C.5: Impulse Response Functions with Non-Rational Output Gap Expectations

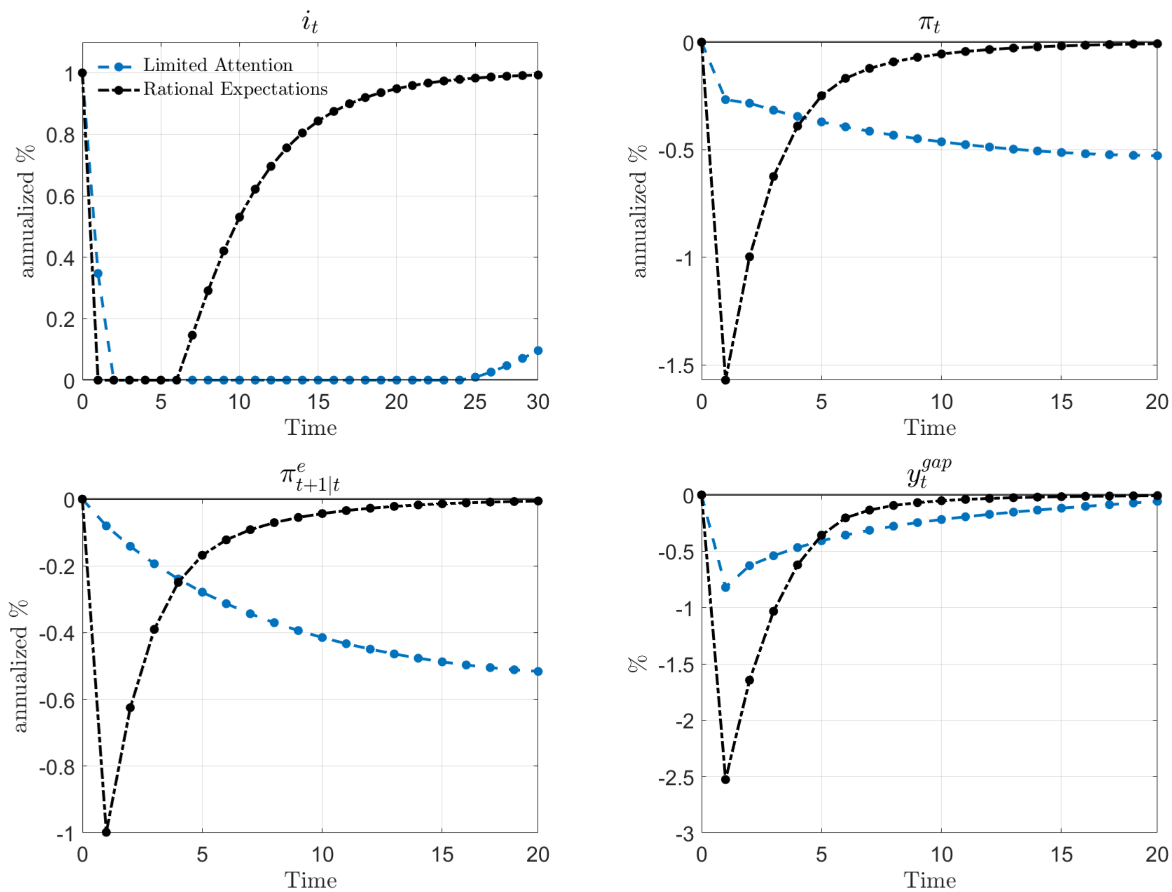

Note: This figure shows the impulse-response functions of the nominal interest rate (upper-left panel), inflation (upper-right panel), inflation expectations (lower-left) and the output gap (lower-right) to a negative natural rate shock of three standard deviations in the case where output gap expectations are given by equation (C.24). The blue-dashed lines show the case for the limited-attention model and the black-dashed-dotted lines for the rational expectations model. Everything is in terms of percentage deviations from the respective steady state levels, except the nominal rate is in levels.

Figure C.6: IRFs to Natural Rate Shock for Taylor rule (C.26)

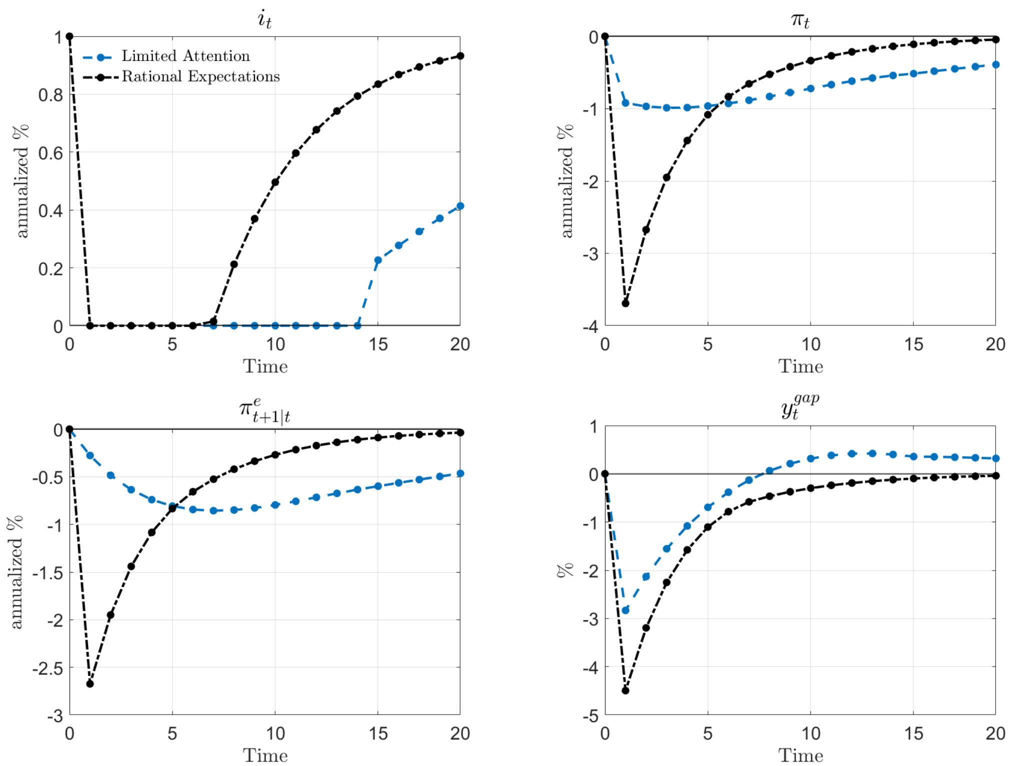

Note: This figure shows the impulse-response functions of the nominal interest rate (upper-left panel), inflation (upper-right panel), inflation expectations (lower-left) and the output gap (lower-right) to a negative natural rate shock of three standard deviations. The blue-dashed lines show the case for the limited-attention model and the black-dashed-dotted lines for the rational expectations model. Everything is in terms of percentage deviations from the respective steady state levels, except the nominal rate is in levels.

C.6.3 Forecast Errors

Angeletos et al. (2021) propose a new test of models that deviate from FIRE. Namely, that expectations should initially underreact but overshoot eventually. A straightforward way to test this is to look at the model-implied impulse response functions of the forecast error, $\pi_{t+1} - \pi^e_{t+1|t}$, to an exogenous shock. Figure C.7 shows these IRFs. The left panel shows the IRF of the forecast error after a positive natural rate shock and the right panel shows the corresponding IRF to a negative natural rate shock. In both cases, we see an underreaction in expectations, which manifests itself in a positive forecast error after a shock that increases the forecasted variable, and vice-versa following a negative shock. After about 5-6 periods, the forecast error response, however, flips sign. This is exactly the eventual overreaction, mentioned above and documented in Angeletos et al. (2021). Thus, my model of inflation expectations matches these empirical findings.

Figure C.7: Impulse Response Functions of Forecast Errors

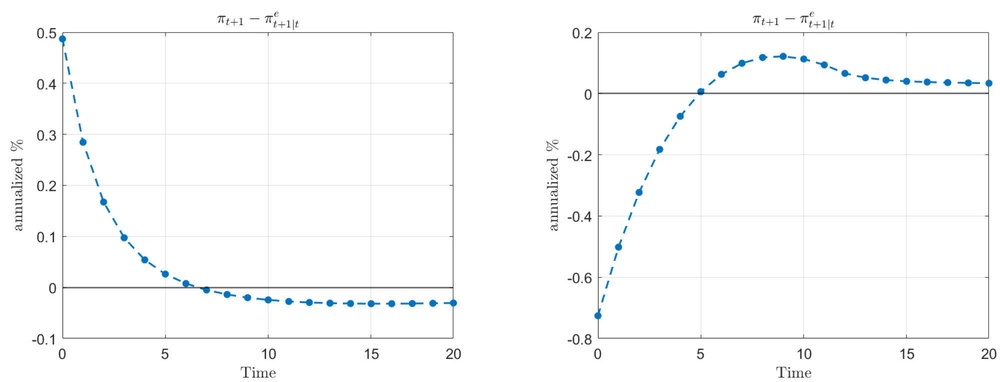

Note: This figure shows the impulse-response functions of inflation forecast errors after a three-standard deviation positive (left) and negative (right) natural rate shock.

C.6.4 No Random Walk

A potential concern with the results stated in Section 3.3, in particular the *inflation-attention trap* in Figure 3.2, is that these findings are driven by the random walk assumption in the belief process of the agents. Relaxing the random-walk assumption requires to take a stand on the perceived average inflation. In this case, where I solve the model around the zero inflation steady state, this is quite innocuous. But later on, when I focus on Ramsey optimal policy, this cannot be done anymore without distorting the results, in the sense that agents might have a mean bias.

Figure C.8 shows the same impulse response functions as reported in Figure 3.2 for the case of $\rho_\pi = 0.95$ and an average inflation of 0. We see a similar pattern, even though somewhat less pronounced. Inflation is persistently lower under limited attention due to slowly-adjusting inflation expectations. Expectations are updated even more sluggishly when $\rho_\pi < 1$. Further, this also dampens the initial response in inflation expectations, and thus, of inflation itself. Therefore, the attention trap is somewhat mitigated and the economy escapes the lower bound faster than with $\rho_\pi = 1$. Nevertheless, the nominal interest rate is low for longer due to the slow recovery of inflation.

Figure C.8: Impulse Response Functions to a Negative Natural Rate Shock

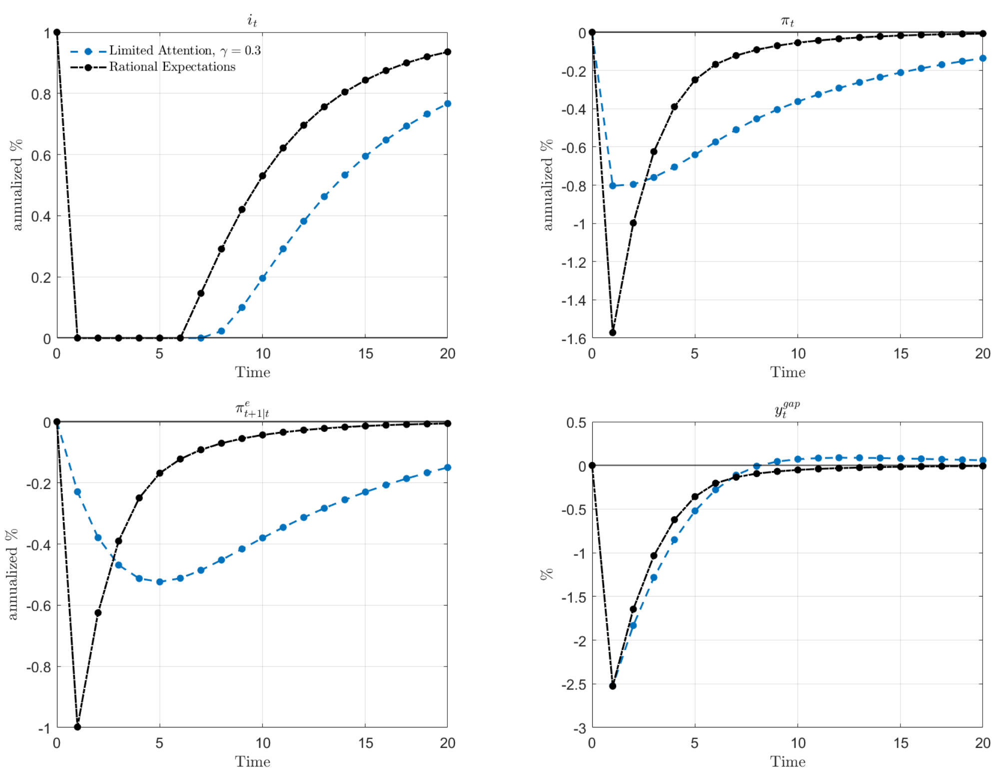

Note: This figure shows the impulse-response functions of the nominal interest rate (upper-right panel), inflation (upper-left panel), inflation expectations (lower-right) and the output gap (lower-left) to a negative natural rate shock of three standard deviations. The blue-dashed lines show the case for the limited-attention model and the black-dashed-dotted lines for the rational expectations model. Everything is in terms of percentage deviations from the respective steady state levels, expect the nominal rate is in levels.

Optimal policy with a bias in inflation expectations. In the main analysis, I have assumed that agents believe that inflation follows a random walk. Under this assumption, inflation expectations and inflation coincide on average. In the following, I relax this assumption and assume that the perceived persistence parameter is less than 1, $\rho_\pi < 1$. As discussed earlier, this yields the following inflation-expectations formation

$$\pi^e_{t+1|t} = (1-\rho_\pi)\bar{\pi} + \rho_\pi \pi^e_{t|t-1} + \rho_\pi \gamma \left(\pi_t - \pi^e_{t|t-1}\right),$$

where $\bar{\pi}$ captures the long-run expectations of the agent. I set $\rho_\pi = 0.95$ and compare economies with different $\bar{\pi}$, namely $\bar{\pi} \in \{0\%, 2\%, 4\%\}$ (annualized).

Figure C.9 shows the optimal inflation target (left panel) and welfare (3.12) (right panel) under Ramsey optimal policy for different levels of attention and different mean beliefs, $\bar{\pi}$. The blue-dashed lines show the results for the case with $\rho_\pi = 1$ (which is the baseline case discussed above), the gray-dashed-dotted lines show the results for $\rho_\pi = 0.95$ and $\bar{\pi} = 0\%$, the black-solid lines for $\rho_\pi = 0.95$ and $\bar{\pi} = 2\%$, and the red-dotted lines for $\rho_\pi = 0.95$ and $\bar{\pi} = 4\%$.

Figure C.9: Mean Bias, Optimal Inflation Target and Welfare

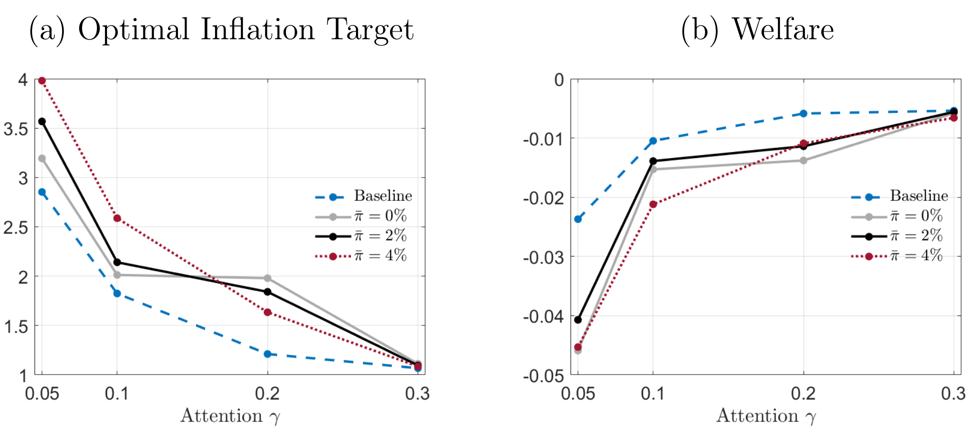

Notes: This figure shows the average inflation rate under Ramsey optimal policy (left panel) and welfare (3.12) (right panel) under Ramsey optimal policy for different levels of attention and different mean beliefs, $\bar{\pi}$. The blue-dashed lines show the results for the case with $\rho_\pi = 1$ (which is the baseline case), the gray-dashed-dotted lines show the results for $\rho_\pi = 0.95$ and $\bar{\pi} = 0\%$, the black-solid lines for $\rho_\pi = 0.95$ and $\bar{\pi} = 2\%$, and the red-dotted lines for $\rho_\pi = 0.95$ and $\bar{\pi} = 4\%$.

We see that introducing a mean bias in general leads to an increase in the optimal inflation target and additional welfare losses, independent of $\bar{\pi}$. This mainly comes from the fact that ρ_π is now below 1, which dampens the degree of updating captured by γ. Thus, once the economy gets stuck at the ELB and the policymaker tries to decrease real rates by increasing inflation expectations, actual inflation needs to increase more strongly. Therefore, a lower ρ_π can exacerbate attention traps when they occur.

Interestingly, the relationship between the optimal target and $\bar{\pi}$ is non-monotonic in the level of attention. While, for example, at $\gamma = 0.2$, the optimal target is highest at $\bar{\pi} = 0\%$, it is highest at $\bar{\pi} = 4\%$ when $\gamma = 0.05$. To understand this, we can write the unconditional average inflation expectations as

$$\pi^e = \frac{(1-\rho_\pi)\bar{\pi} + \rho_\pi \gamma \pi}{1-\rho_\pi(1-\gamma)}.$$

The following Lemma sheds light on how $\bar{\pi}$ matters for average inflation expectations and how this depends on the level of attention, γ.

Lemma 14. *For the case $\rho_\pi = 1$, average inflation expectations move one-for-one with average inflation, independent of γ:*

$$\pi^e = \pi.$$

For the case $0 < \rho_\pi < 1$, average inflation expectations move less than one-for-one with average inflation

$$0 < \frac{\partial \pi^e}{\partial \pi} = \frac{\rho_\pi \gamma}{1-\rho_\pi(1-\gamma)} < 1,$$

and the strength of this dependency increases with γ

$$\frac{\partial^2 \pi^e}{\partial \pi \partial \gamma} > 0.$$

Average inflation expectations move less than one-for-one with $\bar{\pi}$

$$0 < \frac{\partial \pi^e}{\partial \bar{\pi}} = \frac{(1-\rho_\pi)}{1-\rho_\pi(1-\gamma)} < 1,$$

and the strength of this dependency decreases with γ

$$\frac{\partial^2 \pi^e}{\partial \bar{\pi} \partial \gamma} < 0.$$

So, as attention falls, there are several opposing forces at work. On the one hand, the effect of $\bar{\pi}$ on average inflation expectations becomes stronger and thus, also exerts more pressure on actual inflation via the Phillips Curve. On the other hand, increasing the inflation target—average inflation—has a smaller effect on average inflation expectations at low levels of attention. Thus, to increase inflation expectations in this case, the inflation target needs to increase more strongly, which is of course costly. Comparing the optimal inflation targets in Figure C.9, we see that at low levels of attention the first effect dominates. If $\bar{\pi}$ is relatively high, the inflation target is high.

C.6.5 Figures to Section 3.4.4

Figure C.10: Full Attention, $\gamma = 1$

(a) Optimal Inflation Target

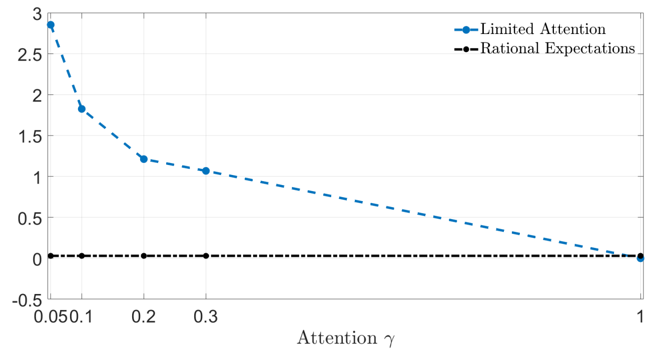

(b) Inflation Volatility

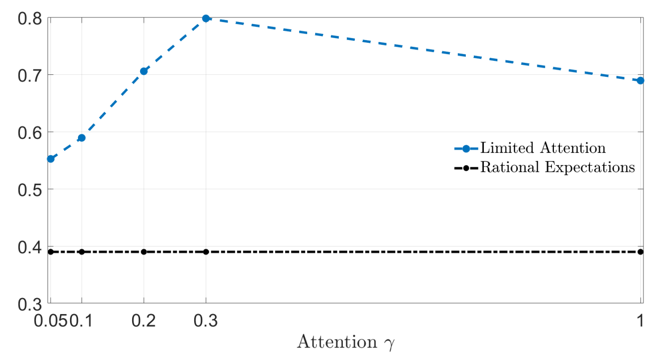

(c) Welfare

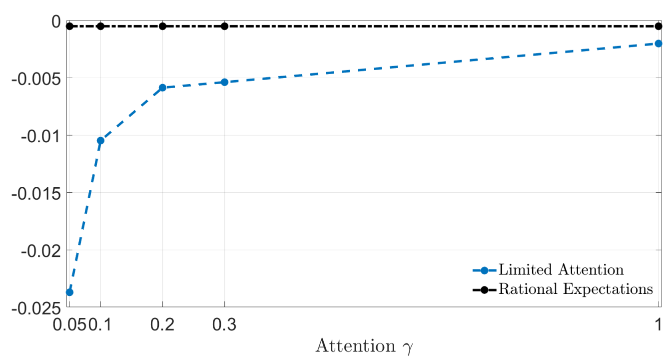

Notes: This figure shows the optimal inflation target (panel (a)), inflation volatility (panel (b)) and welfare (panel (c)) under Ramsey optimal policy for different levels of attention, including full attention, i.e., $\gamma = 1$ and compares it to the full-information rational expectations counterparts (black-dashed-dotted lines).

Appendix D

Bibliography

S. Acharya and K. Dogra. Understanding hank: Insights from a prank. *Econometrica*, 88(3): 1113–1158, 2020.

K. Adam. Optimal monetary policy with imperfect common knowledge. *Journal of Monetary Economics*, 54(2):267–301, 2007.

K. Adam and R. M. Billi. Optimal monetary policy under commitment with a zero bound on nominal interest rates. *Journal of Money, Credit and Banking*, 38(7):1877–1905, 2006a.

K. Adam and R. M. Billi. Optimal monetary policy under commitment with a zero bound on nominal interest rates. *Journal of Money, Credit and Banking*, pages 1877–1905, 2006b.

K. Adam and A. Marcet. Internal rationality, imperfect market knowledge and asset prices. *Journal of Economic Theory*, 146:1224–1252, 2011.

K. Adam and S. Nagel. Expectations data in asset pricing. In *Handbook of Economic Expectations*. Elsevier, 2022.

K. Adam and M. Woodford. Robustly optimal monetary policy in a new keynesian model with housing. *Journal of Economic Theory*, 198:Article 105352, Dec. 2021.

K. Adam, A. Marcet, and J. P. Nicolini. Stock market volatility and learning. *Journal of Finance*, 71(1):33–82, 2016.

K. Adam, A. Marcet, and J. Beutel. Stock price booms and expected capital gains. *Amercian Economic Review*, Vol. 107 No. 8:2352–2408, 2017.

K. Adam, D. Matveev, and S. Nagel. Do survey expectations of stock returns reflect risk-adjustments? *Journal of Monetary Economics*, 117:723–740, 2021.

K. Adam, O. Pfäuti, and T. Reinelt. Subjective housing price expectations, falling natural rates and the optimal inflation target. 2022.

H. Afrouzi and C. Yang. Dynamic rational inattention and the phillips curve. 2021.

H. Afrouzi, S. Y. Kwon, A. Landier, Y. Ma, and D. Thesmar. Overreaction in expectations: Evidence and theory. *Quarterly Journal of Economics (forthcoming)*, 2023.

S. R. Aiyagari. Uninsured idiosyncratic risk and aggregate saving. *The Quarterly Journal of*

Economics, 109(3):659–684, 1994.

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

P. Andrade, J. Galí, H. Le Bihan, and J. Matheron. The optimal inflation target and the natural rate of interest. *Brookings Papers on Economic Activity*, 2019(2):173–255, 2019.

P. Andrade, J. Galí, H. Le Bihan, and J. Matheron. Should the ecb adjust its strategy in the face of lower r*? *Journal of Economic Dynamics and Control*, 2021.

P. Andrade, G. Gaballo, E. Mengus, and B. Mojon. Forward guidance and heterogeneous beliefs. *American Economic Journal: Macroeconomics*, 11(3):1–29, 2006.12019.

G.-M. Angeletos and Z. Huo. Myopia and anchoring. *American Economic Review*, 111(4):1166–1200, 2021.

G.-M. Angeletos and J. La'o. Noisy business cycles. *NBER Macroeconomics Annual*, 24(1):319–378, 2010.

G.-M. Angeletos and C. Lian. Forward guidance without common knowledge. *American Economic Review*, 108(9):2477–2512, 2018.

G.-M. Angeletos and C. Lian. Dampening general equilibrium: incomplete information and bounded rationality. In *Handbook of Economic Expectations*, pages 613–645. Elsevier, 2023.

G.-M. Angeletos, F. Collard, and H. Dellas. Business-cycle anatomy. *American Economic Review*, 110(10):3030–70, 2020.

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

L. Armona, A. Fuster, and B. Zafar. Home price expectations and behaviour: Evidence from a randomized information experiment. *The Review of Economic Studies*, 86(4):1371–1410, 07 2019. ISSN 0034-6527.

A. Auclert. Monetary policy and the redistribution channel. *American Economic Review*, 109(6): 2333–67, 2019.

A. Auclert, M. Rognlie, and L. Straub. The intertemporal keynesian cross. 2018.

A. Auclert, M. Rognlie, and L. Straub. Micro jumps, macro humps: Monetary policy and business cycles in an estimated hank model. Technical report, National Bureau of Economic Research, 2020.

A. Auclert, B. Bardóczy, and M. Rognlie. Mpcs, mpes, and multipliers: A trilemma for new keynesian models. *The Review of Economics and Statistics*, pages 1–41, 2021.

L. Ball, N. G. Mankiw, and R. Reis. Monetary policy for inattentive economies. *Journal of Monetary Economics*, 52(4):703–725, 2005.

N. Barberis, R. Greenwood, L. Jin, and A. Shleifer. X-capm: An extrapolative capital asset pricing model. *Journal of Financial Economics*, 115(1):1–24, 2015.

R. B. Barsky, C. L. House, and M. S. Kimball. Sticky-price models and durable goods. *American Economic Review*, 97(3):984–998, June 2007.

L. Benati. Investigating inflation persistence across monetary regimes. *The Quarterly Journal of Economics*, 123(3):1005–1060, 2008.

J. Benchimol and L. Bounader. *Optimal monetary policy under bounded rationality*. 2019.

P. Benigno and L. Paciello. Monetary policy, doubts and asset prices. *Journal of Monetary Economics*, 64:85–98, 2014.

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

A. Bhandari, J. Borovička, and P. Ho. Survey data and subjective beliefs in business cycle models. 2022.

F. O. Bilbiie. Limited asset markets participation, monetary policy and (inverted) aggregate demand logic. *Journal of economic theory*, 140(1):162–196, 2008.

F. O. Bilbiie. 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(1):221–251, 2011.

F. O. Bilbiie. The new keynesian cross. *Journal of Monetary Economics*, 114:90–108, 2020.

F. O. Bilbiie. Monetary policy and heterogeneity: An analytical framework. 2021.

F. O. Bilbiie, D. R. Känzig, and P. Surico. Capital and income inequality: An aggregate-demand complementarity. *Journal of Monetary Economics*, 126:154–169, 2022.

R. Blundell and S. Bond. Initial conditions and moment restrictions in dynamic panel data models. *Journal of Econometrics*, 87(1):115–143, 1998.

D. Bonciani and J. Oh. Optimal monetary policy mix at the zero lower bound. 2022.

P. Bordalo, N. Gennaioli, Y. Ma, and A. Shleifer. Overreaction to macroeconomic expectations. *American Economic Review*, 110(0):2, 2020.

B. Born, Z. Enders, M. Menkhoff, G. J. Müller, and K. Niemann. Firm expectations and news: Micro v macro. 2022.

M. Boutros. Windfall income shocks with finite planning horizons. 2022.

L. Brandão-Marques, M. Casiraghi, G. Gelos, G. Kamber, and R. Meeks. Negative interest rate policies: A survey. 2021.

T. Broer, N.-J. Harbo Hansen, P. Krusell, and E. Öberg. The new keynesian transmission mechanism: A heterogeneous-agent perspective. *The Review of Economic Studies*, 87(1):77–101, 2020.

T. Broer, A. Kohlhas, K. Mitman, and K. Schlafmann. Information and wealth heterogeneity in the macroeconomy. 2021a.

T. Broer, P. Krusell, and E. Öberg. Fiscal multipliers: A heterogenous-agent perspective. 2021b.

C. Caines and F. Winkler. Asset price beliefs and optimal monetary policy. *Journal of Monetary Economics*, 123:53–67, 2021.

G. A. Calvo. Staggered contracts in a utility-maximizing framework. *Journal of Monetary Economics*, 12:383–398, 1983.

B. Candia, O. Coibion, and Y. Gorodnichenko. The inflation expectations of us firms: Evidence from a new survey. 2021.

F. Canova. G-7 inflation forecasts: Random walk, phillips curve or what else? *Macroeconomic Dynamics*, 11(1):1, 2007.

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

J. A. Carlson and M. Parkin. Inflation expectations. *Economica*, 42(166):123–138, 1975.

C. D. Carroll. Macroeconomic expectations of households and professional forecasters. *the Quarterly Journal of economics*, 118(1):269–298, 2003.

K. E. Case, R. J. Shiller, and A. K. Thompson. What have they been thinking? homebuyer behavior in hot and cold markets. *Brookings Papers on Economic Activity*, pages 265–315, 2012.

G. Chodorow-Reich. Geographic cross-sectional fiscal spending multipliers: What have we learned? *American Economic Journal: Economic Policy*, 11(2):1–34, 2019.

R. Clarida, J. Galí, and M. Gertler. The science of monetary policy: Evidence and some theory. *Journal of Economic Literature*, 37:1661–1707, 1999.

J. H. Cochrane. Michelson-morley, fisher, and occam: The radical implications of stable quiet inflation at the zero bound. *NBER Macroeconomics Annual*, 32(1):113–226, 2018.

O. Coibion and Y. Gorodnichenko. What can survey forecasts tell us about information rigidities? *Journal of Political Economy*, 120(1):116–159, 2012.

O. Coibion and Y. Gorodnichenko. Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review*, 105(8):2644–78, 2015a.

O. Coibion and Y. Gorodnichenko. Is the phillips curve alive and well after all? inflation expectations and the missing disinflation. *American Economic Journal: Macroeconomics*, 7(1):197–232, 2015b.

O. Coibion, Y. Gorodnichenko, and J. Wieland. The optimal inflation rate in new keynesian models: should central banks raise their inflation targets in light of the zero lower bound? *Review of Economic Studies*, 79(4):1371–1406, 2012.

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

O. Coibion, Y. Gorodnichenko, E. Knotek, and R. Schoenle. Average inflation targeting and household expectations. *Journal of Political Economy: Macroeconomics (forthcoming)*, 2022a.

O. Coibion, Y. Gorodnichenko, and M. Weber. Monetary policy communications and their effects on household inflation expectations. *Journal of Political Economy*, 130(6):1537–1584, 2022b.

A. Colciago. Rule-of-thumb consumers meet sticky wages. *Journal of money, credit and banking*, 43(2-3):325–353, 2011.

V. Constancio. Understanding inflation dynamics and monetary policy. In *Speech at the Jackson Hole Economic Policy Symposium*, volume 29, 2015.

D. Debortoli and J. Galí. Monetary policy with heterogeneous agents: Insights from tank models. 2018.

D. Debortoli, J. Galí, and L. Gambetti. On the empirical (ir) relevance of the zero lower bound constraint. *NBER Macroeconomics Annual*, 34(1):141–170, 2020.

M. Del Negro, M. P. Giannoni, and C. Patterson. The forward guidance puzzle. 2012.

M. Del Negro, M. Lenza, G. E. Primiceri, and A. Tambalotti. What's up with the phillips curve? 2020.

J. di Giovanni, Kalemli-Özcan, A. Silva, and M. A. Yildirim. Global supply chain pressures, international trade, and inflation. 2022.

B. Dupor, M. Karabarbounis, M. Kudlyak, and M. Mehkari. Regional consumption responses and the aggregate fiscal multiplier. *Review of Economic Studies (forthcoming)*, 2022.

F. D'Acunto, D. Hoang, and M. Weber. Managing households' expectations with unconventional policies. *The Review of Financial Studies*, 35(4):1597–1642, 2022.

G. Eggertsson and S. Singh. Log-linear approximation versus an exact solution at the zlb in the new keynesian model. *Journal of Economic Dynamics and Control*, 2019.

G. Eggertsson and M. Woodford. The zero interest-rate bound and optimal monetary policy. *Brookings Papers on Economic Activity*, (1):139–211, 2003.

C. J. Erceg, D. W. Henderson, and A. T. Levin. Optimal monetary policy with staggered wage and price contracts. *Journal of monetary Economics*, 46(2):281–313, 2000.

G. W. Evans and S. Honkapohja. Expectations and the stability problem for optimal monetary policies. *The Review of Economic Studies*, 70(4):807–824, 2003.

A. Fagereng, M. B. Holm, and G. J. Natvik. Mpc heterogeneity and household balance sheets. *American Economic Journal: Macroeconomics*, 13(4):1–54, 2021.

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

J. Faust and J. H. Wright. Forecasting inflation. In *Handbook of economic forecasting*, volume 2, pages 2–56. Elsevier, 2013.

A. Ferriere and G. Navarro. The heterogeneous effects of government spending: It's all about taxes. 2022.

J. C. Fuhrer and G. D. Rudebusch. Estimating the euler equation for output. *Journal of Monetary Economics*, 51(6):1133–1153, 2004.

S. Fujiwara, Y. Iwasaki, I. Muto, K. Nishizaki, and N. Sudo. Developments in the natural rate of interest in japan. *Bank of Japan Review*, 2016.

C. Fulton and K. Hubrich. Forecasting us inflation in real time. *Econometrics*, 9(4):36, 2021.

X. Gabaix. A sparsity-based model of bounded rationality. *The Quarterly Journal of Economics*, 129(4):1661–1710, 2014.

X. Gabaix. Behavioral macroeconomics via sparse dynamic programming. 2016.

X. Gabaix. Behavioral inattention. In *Handbook of Behavioral Economics: Applications and*

Foundations 1, volume 2, pages 261–343. Elsevier, 2019.

X. Gabaix. A behavioral new keynesian model. *American Economic Review*, 110(8):2271–2327, 2020.

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

J. Galí, J. D. López-Salido, and J. Vallés. Understanding the effects of government spending on consumption. *Journal of the european economic association*, 5(1):227–270, 2007.

J.-E. Gallegos. Hank beyond fire. 2023.

L. Gáti. Monetary policy & anchored expectations an endogenous gain learning model. 2022.

Y. Gorodnichenko and M. Weber. Are sticky prices costly? evidence from the stock market. *American Economic Review*, 106(1):165–99, 2016.

R. Greenwood and A. Shleifer. Expectations of returns and expected returns. *Review of Financial Studies*, 27 (3), 2014.

R. S. Gürkaynak, B. Sack, and E. Swanson. The sensitivity of long-term interest rates to economic news: Evidence and implications for macroeconomic models. *American economic review*, 95(1): 425–436, 2005.

M. Hagedorn. Prices and inflation when government bonds are net wealth. 2018.

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

M. Hagedorn et al. *A demand theory of the price level.* 2021.

T. Hintermaier and W. Koeniger. Differences in euro-area household finances and their relevance for monetary-policy transmission. 2019.

M. B. Holm, P. Paul, and A. Tischbirek. The transmission of monetary policy under the microscope. *Journal of Political Economy*, 129(10):2861–2904, 2021.

K. Holston, T. Laubach, and J. C. Williams. Measuring the natural rate of interest: International trends and determinants. *Journal of International Economics*, 108:S59 – S75, 2017. ISSN 0022-1996.

M. Huggett. The risk-free rate in heterogeneous-agent incomplete-insurance economies. *Journal of economic Dynamics and Control*, 17(5):953–969, 1993.

O. Jorda. Estimation and inference of impulse responses by local projections. *American Economic Review*, 95(1):161–182, March 2005.

P. L. Jørgensen and K. J. Lansing. Anchored inflation expectations and the slope of the phillips curve. 2021.

G. Kaplan and G. L. Violante. The marginal propensity to consume in heterogeneous agent models. *Annual Review of Economics*, 14:747–775, 2022.

G. Kaplan, B. Moll, and G. L. Violante. Monetary policy according to hank. *American Economic Review*, 108:697–743, 2018.

G. Kaplan, K. Mitman, and G. L. Violante. The housing boom and bust: Model meets evidence.

Journal of Political Economy, 128(9):3285–3345, 2020.

A. N. Kohlhas and A. Walther. Asymmetric attention. *American Economic Review*, 119(1): 2879–2925, 2021.

P. Krusell, T. Mukoyama, and A. A. Smith Jr. Asset prices in a huggett economy. *Journal of Economic Theory*, 146(3):812–844, 2011.

T. Kuchler and B. Zafar. Personal experiences and expectations aboutaggregate outcomes. *Journal of Finance*, VOL. LXXIV, Oct. 2019.

S. Kučinskas and F. S. Peters. Measuring under-and overreaction in expectation formation. *The Review of Economics and Statistics*, pages 1–45, 2022.

M. Kuhn, M. Schularick, and U. I. Steins. Income and wealth inequality in america, 1949–2016. *Journal of Political Economy*, 128(9):3469–3519, 2020.

D. Laibson, P. Maxted, and B. Moll. Present bias amplifies the household balance-sheet channels of macroeconomic policy. 2021.

M. J. Lamla and S. M. Lein. The role of media for consumers' inflation expectation formation. *Journal of Economic Behavior & Organization*, 106:62–77, 2014.

C. Lian. Mistakes in future consumption, high mpcs now. 2023.

R. Luetticke. Transmission of monetary policy with heterogeneity in household portfolios. *American Economic Journal: Macroeconomics*, 13(2):1–25, 2021.

C. Ma. Momentum and reversion to fundamentals: Are they captured by subjective expectations of house prices? *Journal of Housing Economics*, 49:101687, 2020. ISSN 1051-1377.

B. Mackowiak and M. Wiederholt. Optimal sticky prices under rational inattention. *American Economic Review*, 99(3):769–803, 2009.

B. Maćkowiak, F. Matějka, and M. Wiederholt. Dynamic rational inattention: Analytical results. *Journal of Economic Theory*, 176:650–692, 2018.

B. Mackowiak, F. Matejka, and M. Wiederholt. Rational inattention: A review. *Journal of Economic Literature (forthcoming)*, 2021.

U. Malmendier and S. Nagel. Depression babies: Do macroeconomic experiences affect risk taking. *Quarterly Journal of Economics*, 126:373–416, 2011.

U. Malmendier and S. Nagel. Learning from inflation experiences. *Quarterly Journal of Economics*, 131(1):53–87, 2015.

N. G. Mankiw. The savers-spenders theory of fiscal policy. *American economic review*, 90(2): 120–125, 2000.

N. G. Mankiw and R. Reis. Sticky information versus sticky prices: a proposal to replace the new keynesian phillips curve. *The Quarterly Journal of Economics*, 117(4):1295–1328, 2002.

A. Marcet and R. Marimon. Recursive contracts. *Econometrica*, 87(5):1589–1631, 2019.

F. Matějka and A. McKay. Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review*, 105(1):272–98, 2015.

A. McKay and J. F. Wieland. Forward guidance and durable goods demand. *American Economic*

Review: Insights, 4(1):106–22, 2022.

A. McKay, E. Nakamura, and J. Steinsson. The power of forward guidance revisited. *American Economic Review*, 106(10):3133–58, 2016.

A. McKay, E. Nakamura, and J. Steinsson. The discounted euler equation: A note. *Economica*, 84(336):820–831, 2017.

A. Mele, K. Molnar, and S. Santoro. On the perils of stabilizing prices when agents are learning. *Journal of Monetary Economics*, 115:339–353, 2020. ISSN 0304-3932.

P. Michaillat and E. Saez. Resolving new keynesian anomalies with wealth in the utility function. *Review of Economics and Statistics*, 103(2):197–215, 2021.

F. Milani. Expectations, learning and macroeconomic persistence. *Journal of monetary Economics*, 54(7):2065–2082, 2007.

K. Molnar and S. Santoro. Optimal monetary policy when agents are learning. *European Economic Review*, 66:39–62, 2014. ISSN 0014-2921.

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

S. Nagel and Z. Xu. Asset pricing with fading memory. *The Review of Financial Studies*, 35(5): 2190–2245, 2022.

E. Nakamura and J. Steinsson. Fiscal stimulus in a monetary union: Evidence from us regions. *American Economic Review*, 104(3):753–92, 2014.

W. K. Newey and K. D. West. A simple, positive semi-definite, heteroskedasticity and autocorrelation. *Econometrica*, 55(3):703–708, 1987.

L. Paciello and M. Wiederholt. Exogenous information, endogenous information, and optimal monetary policy. *Review of Economic Studies*, 81(1):356–388, 2014.

C. Patterson. The matching multiplier and the amplification of recessions. *American Economic Review (forthcoming)*, 2022.

D. Pfajfar and E. Santoro. News on inflation and the epidemiology of inflation expectations. *Journal of Money, Credit and Banking*, 45(6):1045–1067, 2013.

M. Piazzesi and M. Schneider. *Equilibrium Yield Curves*, volume 21 of *NBER Macoreconomics Annual*, chapter 6, pages 389–442. University of Chicago Press, 2006.

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

M. O. Ravn and V. Sterk. Macroeconomic fluctuations with hank & sam: An analytical approach. *Journal of the European Economic Association*, 19(2):1162–1202, 2021.

C. Roth, M. Wiederholt, and J. Wohlfart. The effects of forward guidance: Theory with measured expectations. 2023.

A. Samarina and A. D. Nguyen. Does monetary policy affect income inequality in the euro area? *Journal of Money, Credit, and Banking*, 2023.

C. A. Sims. Implications of rational inattention. *Journal of Monetary Economics*, 50(3):665–690, 2003.

J. Slacalek, O. Tristani, and G. L. Violante. 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, 2020.

R. F. Stambaugh. Predictive regressions. *Journal of Financial Economics*, 54:375–421, 1999.

L. E. O. Svensson. *Festschrift in Honour of Vitor Constancio*, chapter The Future of Monetary Policy and Macroprudential Policy,, pages 69–123. European Central Bank, 2018.

N. Vellekoop and M. Wiederholt. Inflation expectations and choices of households. 2019.

I. Werning. Incomplete markets and aggregate demand. 2015.

M. Wiederholt. Empirical properties of inflation expectations and the zero lower bound. 2015.

M. Woodford. Interest and prices. 2003.

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

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

T. Yun. Nominal price rigidity, money supply endogeneity, and business cycles. *Journal of monetary Economics*, 37(2):345–370, 1996.

Eidesstattliche Versicherung gemäß §8 Absatz 2 Buchstabe a) der Promotionsordnung der Universität Mannheim zur Erlangung des Doktorgrades der Volkswirtschaftslehre (Dr. rer. pol.)

1. Bei der eingereichten Dissertation mit dem Titel „Three Essays on Macroeconomics„ handelt es sich um mein eigenständig erstelltes Werk, das den Regeln guter wissenschaftlicher Praxis entspricht.

2. Ich habe nur die angegebenen Quellen und Hilfsmittel benutzt und mich keiner unzulässigen Hilfe Dritter bedient. Insbesondere habe ich wörtliche und nicht wörtliche Zitate aus anderen Werken als solche kenntlich gemacht.

3. Die Arbeit oder Teile davon habe ich bislang nicht an einer Hochschule des In- oder Auslands als Bestandteil einer Prüfungs- oder Qualifikationsleistung vorgelegt.

4. Die Richtigkeit der vorstehenden Erklärung bestätige ich.

5. Die Bedeutung der eidesstattlichen Versicherung und die strafrechtlichen Folgen einer unrichtigen oder unvollständigen eidesstattlichen Versicherung sind mir bekannt. Ich versichere an Eides statt, dass ich nach bestem Wissen die reine Wahrheit erklärt und nichts verschwiegen habe.

   Die eingereichten Dissertationsexemplare sowie der Datenträger gehen in das Eigentum der Universität über.

Mannheim

Oliver Pfäuti

Curriculum Vitae

2018 - 2023	University of Mannheim Center for Doctoral Studies in Economics Ph.D. in Economics (Dr. rer. pol.)
2021 - 2022	Yale University Visiting Ph.D. Student
2016 - 2018	Universities of Bern and Basel Master of Science in International and Monetary Economics
2013 - 2016	University of Bern Bachelor of Science in Economics