Beyond bad luck... Macroeconomic implications of persistent heterogeneity in optimism

Oliver Pfäuti[*] Fabian Seyrich[†] Jonathan Zinman[‡]

February 25, 2025

Abstract

What explains heterogeneity in households' savings behavior and in their financial situations? Using a novel micro data set, we identify consumers with an optimistic bias about their own future financial situation and find them to be less likely to save and more likely to be financially constrained. We quantify heterogeneity in the optimistic bias and use these estimates to introduce and discipline heterogeneity in optimism in an otherwise standard HANK model. The model rationalizes that optimistic households spend instead of precautionary save rendering many of these households financially constrained. The model produces empirically realistic wealth and MPC distributions, and optimistic households have higher MPCs than rational households even away from the borrowing constraint. Accounting for heterogeneity in optimism renders targeted transfers less stimulative, reduces crowding out of precautionary savings through minimum income benefits and reduces the optimal public debt level.

JEL Codes: D91; E21; E62; E71; G51

Keywords: Household Heterogeneity; Optimism; Subjective Financial Forecast Errors; Overconfidence; Financial Constraints; Fiscal Policy; HANK

[*]The University of Texas at Austin, Department of Economics, pfaeuti.oliver@gmail.com

[†]Frankfurt School of Finance & Management and DIW Berlin, fabian.seyrich@gmail.com

[‡]Dartmouth College, Department of Economics and NBER, jzinman@dartmouth.edu.

We thank Eduardo Dávila, Greg Kaplan, Ralph Luetticke (discussant), Kurt Mitman, Peter Maxted, Matthew Rognlie, Johannes Stroebel, Hannes Twieling, Gianluca Violante, Michael Weber, Nathan Zorzi and seminar and conference participants at the Macro Mini-Conference at Dartmouth and the ASSA meeting at San Francisco. Pfäuti gratefully acknowledges financial support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 (Project C02) and the Stiftung Geld & Währung. Seyrich gratefully acknowledges financial support by the Leibniz Association through the project "Distributional effects of macroeconomic policies in Europe". Zinman thanks the Nelson A. Rockefeller Center for Public Policy at Dartmouth for funding.

1 Introduction

Heterogeneity in households' savings behavior and financial situations has significant implications for macroeconomic fluctuations and policy design. The empirical magnitude of such heterogeneity is strikingly large and persistent.[1] Yet it remains standard practice in macroeconomic models of fluctuations and stabilization to assume ex-ante identical households and account for savings heterogeneity only in shock realizations: Households are wealthy or poor only because of good luck or bad luck, abstracting from more fundamental dimensions of heterogeneity including beliefs.[2]

Mounting empirical evidence points to substantial dispersion in household expectations (see D'Acunto and Weber (2024) for a recent review). Although work thus far focuses on aggregate variables, standard consumption-savings theory predicts that what matters most directly for households' savings and consumption choices are their expectations about their *own* future financial situation. This suggests that quantifying heterogeneity in beliefs about households' own future financial situations could help us to better understand the drivers of heterogeneity in savings behavior and financial situations, including in hand-to-mouth (HtM) status and marginal propensities to consume (MPCs).

Beliefs about households' own future financial situations have been elicited for decades with questions of the form: "Do you think that a year from now you will be better off financially, or worse off, or just about the same as now?". Prior empirical work documents a strong optimistic slant in this micro belief in the aggregate that holds across various time periods, datasets, and countries (Souleles, 2004; Claus and Nguyen, 2023; Cocco et al., 2021).

We contribute to this empirical work by developing a new approach to quantify persistent optimistic biases at the individual level and its link to consumption-savings decisions, using observed forecast errors while accounting for key data limitations in doing so. We focus on persistent optimism because quantitative macro models have been struggling to jointly fit the data on key variables that are empirically quite persistent: wealth, MPCs, and HtM status.[3] Our micro estimates show that persistent optimistic forecast errors are prevalent in nationally representative data, consumers are more than twice as likely to be persistently optimistic than pessimistic, and ex-ante optimistic bias is strongly correlated with not saving and being HtM.

These results suggest that belief heterogeneity can indeed help explain observed heterogene-

[1] See Kaplan and Violante (2022), McKay and Wolf (2023), and Auclert et al. (2024b) for recent review articles on how household heterogeneity shapes macroeconomic fluctuations, policy effectiveness and design. Aguiar et al. (forthcoming) focuses on the empirical persistence of hand-to-mouth status. Ganong et al. (2024), Fagereng et al. (2021) and Lewis et al. (2024) find that unobserved and persistent individual characteristics are important drivers of heterogeneity in MPCs.

[2] Important exceptions have focused on preferences, as we discuss below.

[3] Evidence from other fields suggests optimistic bias is a persistent decision making feature (Carver et al., 2010).

ity in households' savings behavior, and they guide us in augmenting an otherwise standard heterogeneous-agent New Keynesian (HANK) model with empirically-disciplined belief heterogeneity. The model jointly matches key untargeted moments in micro and macro data, and allows us to derive macroeconomic implications of persistent heterogeneity in optimism including its novel implications for various fiscal policies.

To discipline the degree of persistent belief heterogeneity in the model, we develop a new empirical strategy for conducting inference from observed *ex-post* forecast errors on *ex-ante* forecast biases. In our case, this is key for distinguishing between permanent consumer heterogeneity (optimism as a feature of decision making) and ex-post optimism that is simply due to bad luck. Measuring ex-ante bias directly, using only observed forecast errors, would require benchmarking beliefs against the rational forecast. But the rational forecast is unobserved for a consumer-specific variable like our object of interest: one's overall household financial situation.

Instead, we distinguish ex-ante from ex-post optimism by identifying the variation in ex-post optimism that is predicted by a well-established and observable source of persistent behavioral heterogeneity in economics: *overconfidence*. Overconfidence is unlikely to be driven by bad luck because overconfidence is—like ex-ante optimism—often seen as a trait-like characteristic of consumers. Johnson and Fowler (2011), for example, find that "overconfident populations are evolutionary stable in a wide range of environments". And prior work conceptualizes overconfidence and over-optimism more or less interchangeably, as biases towards overestimating the likelihood that favorable states of the world will be realized (e.g., Carver et al. (2010) Section 1; Johnson and Fowler (2011) footnote 14; Spinnewijn (2015) footnote 2).As a result, a strong correlation between overconfidence an ex-post optimism then implies an important role for ex-ante optimism as driver of ex-post optimism.

We use a measure of overconfidence about one's own cognitive skills, and specifically about one's performance on an intelligence test, that exhibits trait-like temporal stability estimated using various methods (Stango and Zinman, 2020, 2024). Related forms of overconfidence about performance have been documented as persistently prevalent and impactful in high-stakes workplace settings (e.g., Huffman et al. (2022); Weidmann et al. (2024); Heck et al. (2024)).[4]

We find that overconfidence is indeed strongly correlated with persistent ex-post optimism about one's financial situation. Conditional on a rich set of household variables that includes not just standard demographics and income, but also cognitive skills (including financial literacy), patience, and risk aversion, households in the 75th percentile of the overconfidence distribution are about 13-36 percentage points more likely to be persistently optimistic about their own future financial situation than households in the 25th percentile.

[4]Related work finds that various forms of overconfidence are correlated within-person (Chapman et al., 2023; Stango and Zinman, 2023).

We then consider how consumers' decisions and actual financial situations are related to their overconfidence. We find strong conditional correlations between savings behavior and overconfidence, as well as between HtM status and overconfidence; suggesting that overconfidence produces not only an optimistic bias but also leads households not to save and to become financially constrained. And overall, we find that the data favor optimism and overconfidence as explanations for HtM status over patience and risk aversion.[5]

Guided by our empirical findings, we introduce belief heterogeneity into an otherwise standard heterogeneous agent New Keynesian (HANK) model with incomplete markets, idiosyncratic productivity risk, borrowing constraints, and a nominal rigidity in the form of sticky wages. Households differ permanently in their degree of ex-ante optimism, with 50% of households holding rational expectations while 50% are ex-ante overly optimistic to varying degrees as estimated in the micro data.[6]

In our baseline model specification, we micro-found differences in over-optimism as differences in beliefs about one's own future labor productivity (and then show, in extensions, that our model results are robust to other microfoundations including over-optimism about aggregates and about expenses). Focusing on labor income makes sense conceptually, since our measure of ex-ante optimistic bias—overconfidence—concerns cognitive skills, and such skills have well-established links to income and workplace performance. Focusing on income also makes sense empirically, given mounting evidence that consumers hold non-FIRE beliefs about it (d'Haultfoeuille et al., 2021), with many making optimistic forecast errors (e.g., Souleles (2004)). We also show that households' financial situation forecasts are strongly positively correlated with their income forecasts.

The model formalizes how more optimistic households tend to consume rather than precautionary save, creating heterogeneity in savings behavior across otherwise identical households. Optimists persistently forecast a better future, dampening their subjective valuation of self-insurance in the form of precautionary saving and increasing their marginal propensity to consume (MPC) even away from the borrowing constraint. Given their lower buffer stocks, optimists are ex-post more likely to end up borrowing-constrained. In contrast, a rational

[5]Given that through the lens of typically estimated log-normal productivity processes, more risky income processes can increase ex-post optimism, one potential concern might be that overconfidence is associated with higher risk. Yet, this is highly unlikely to drive our results for three reasons: First, we control for a wide set of consumer-level controls, including risk aversion, education, cognitive skills, and income in our regressions. Second, through the lens of a standard model, explaining our empirical magnitudes would require these households to face drastically higher risk, in the order of a magnitude. Third, if these households indeed faced higher risk *and* were rational about it, standard models predict them to save more and to be less likely to be hand-to-mouth. As suggested in the previous paragraph, the micro data strongly rejects that pattern and in fact reveals the opposite—overconfident households are conditionally substantially less likely to save and more likely to be HtM.

[6]In an extension, we also add households with a pessimistic bias. This does not materially change our results qualitatively and tends to strengthen them quantitatively.

household strives both to avoid the HtM state, and to save its way off the constraint after having reached it.

Accounting for such heterogeneity in beliefs greatly improves the model's fit of several key moments in micro and macro data. From a micro perspective, our model implies MPC heterogeneity across households even conditional on wealth and income, including relatively flat MPC-income and MPC-wealth distributions in line with recent findings (see e.g. Boehm et al. (2023)).[7] Those distributions then allow our model to jointly match total wealth in the economy, high HtM prevalence, and an average quarterly marginal propensity to consume (MPC) in the consensus range of 15-25% (e.g., Jappelli and Pistaferri (2010), Havranek and Sokolova (2020)). This is in contrast to rational, one-asset HANK models which fail to jointly match average wealth and average MPCs (Auclert et al., forthcoming; Kaplan and Violante, 2022).

Our model also accounts well for other untargeted wealth statistics, including filling in the "missing middle" of a standard one-asset HANK model. A two-asset version of our model fits the data with a substantially lower and empirically more realistic liquidity premium than usual (Kaplan and Violante, 2022; Auclert et al., forthcoming).

Accounting for heterogeneity in optimism also generates important and distinct implications for macroeconomic policies. We first derive the aggregate consequences of unexpected transfer payments intended to stimulate private consumption. Such interventions often target lower-income households (e.g., COVID stimulus checks, part of the American Recovery and Reinvestment Act of 2009). The effectiveness of such transfers depends on the distribution of MPCs across targeted vs. non-targeted groups, and as such our model's ability to produce an empirically realistic MPC-income distribution matters. The weaker correlation between MPCs and income in our model dictates that income-targeted transfers are much less powerful than in a rational model recalibrated to match the same average MPC. That model counterfactually imposes high MPCs on practically all low-income households and implies that giving transfers to the bottom income quartile produces an impact multiplier of 2.8 in general equilibrium. Our model produces a multiplier of only 0.8.

Our model also changes the efficacy of fiscal policies impacting household self-insurance decisions in steady-state. The key mechanism is that optimistic households react less strongly to changes in precautionary savings incentives which mutes the aggregate response of households to such changes compared to models in which households are uniformly rational.

As such our second policy application considers steady-state social insurance in the form of

[7] Also consistent with our model, Ganong et al. (2024) find that households with persistently low liquid asset holdings still have high MPCs even when they enter states of high liquidity—as observed due to substantial increases in unemployment benefits during the Covid-19 pandemic. They further highlight that these patterns are likely driven by permanent household characteristics, supporting our approach of modeling differences in optimism as a form of permanent heterogeneity.

minimum income benefits. Our model predicts less crowd out of private precautionary savings than its standard, rational counterpart. Optimistic households undervalue the insurance benefits because they underestimate their probability of reaching bad income states and therefore reduce any existing buffer stock only mildly. They are also much less likely to have any existing buffer stock to begin with. Introducing minimum income benefits thus only weakly increases the steady-state share of HtM households and the equilibrium real interest rate in our model, in contrast to the rational models.

We contrast this with an indirect insurance fiscal policy through government debt issuance (e.g., Woodford (1990), Aiyagari and McGrattan (1998)). Higher government debt levels reduce households' self-insurance cost by reducing the price of liquid assets. But the induced increase in precautionary savings is muted in our model because optimistic households undervalue the insurance function of cheaper assets. Thus even at high public debt levels, many optimistic households do not save themselves out of being constrained, the HtM share remains high, and the wealth share of the bottom 50% remains stubbornly low. If all households are rational, low-wealth households save themselves away from the borrowing constraint and increase their savings strongly in response to cheaper liquidity. This drives down the HtM share strongly and increases the wealth share held by the bottom half of the distribution. These contrasting effects have normative implications as well: the optimal government debt level is substantially lower with heterogeneity in optimism, irrespective of whether we consider a model in which households can only save in government bonds or also in productive capital.

Overall, we show that accounting for observed *persistent* differences in consumer beliefs about their future financial situation is crucial for understanding household finances, macroeconomic fluctuations and stabilization, and general equilibrium. Our approach contrasts sharply both with models assuming rational expectations ("RE") and with behavioral models where the only potential deviation from RE regards some aggregate variable. In those classes of models, households become borrowing constrained because they are unlucky, i.e., hit by adverse shocks, and HtM tends to be a relatively transitory state. In our model, households are financially constrained mostly because of ex-ante optimism about their future financial situation. Our setup can accommodate various drivers of optimism, jointly fits key moments in the micro and macro data remarkably well, and has several novel implications for fiscal policy and social insurance.

Related literature. We contribute to several literatures.

Our paper is related to the literature that considers potential roles for optimism (and/or pessimism), including those regarding sentiment (Kamdar and Ray, 2024; Pappa et al., 2023; Bhandari et al., 2024). That work has mainly focused on time-varying beliefs about aggregates, often with homogeneous consumers. We focus on heterogeneity in persistent (optimistic) beliefs about one's own financial situation, to which beliefs about aggregates are an input. Another

input are beliefs about expenses, which could also be quantitatively important for macro given evidence on the prevalence, heterogeneity, and magnitude of expense neglect (Kaur et al., 2025; Berman et al., 2016) and expense shocks (Fulford and Low, 2024). Cocco et al. (2021) focus on financial situation belief updating in response to shocks (in contrast to our focus on the large persistent component of optimistic bias) and on modeling its life-cycle implications in partial equilibrium (in contrast to our focus on modeling macro fluctuations in general equilibrium). Our modeling of optimism about one's financial situation as being driven by over-optimism about income relates to Rozsypal and Schlafmann (2023)'s interest in optimistic and pessimistic income forecast errors, albeit with several key differences in research questions, inference, and modeling.[8]

We also contribute to behavioral macroeconomics more broadly, by accounting for behavioral heterogeneity in decision making. Other work focuses on a representative behavioral agent.[9] Behavioral HANK models tend to allow for heterogeneity only in the budget constraint, with a representative behavioral decision maker or information friction about an aggregate variable only.[10] Kaplan and Violante (2022) allow for heterogeneity in present-biased or temptation preferences but leave those preferences as free parameters. We instead use micro data to quantitatively discipline heterogeneity in optimism. We also find a distinct pattern of results. Kaplan and Violante (2022) show that present bias actually can lower the average MPC, in stark contrast to our results indicating that optimism can generate a realistic average MPC while also matching other key micro and macro data moments. The model with temptation can generate a realistic average MPC but suffers from the "missing middle" problem, unlike ours. Pfäuti and Seyrich (2023) study a case of heterogeneous behavioral biases, but focus on expectations about aggregate variables in that case. Guerreiro (2023) allows for heterogeneous attention about aggregates, but assumes rational expectations about households' idiosyncratic shocks. Ilut and Valchev (2023) develop a theory of imperfect reasoning and introduce this into an Aiyagari (1994) economy. They show agents reason less when they face familiar states which can lead to learning traps in which households remain persistently at the borrowing constraint. We contribute to their work by providing quantitative estimates of the relevant behavioral bias and we focus on the role of ex-ante heterogeneity whereas they focus on households that are

[8] We focus on persistent belief heterogeneity regarding financial situation, whereas Rozsypal and Schlafmann (2023) (RS) observes at most one forecast error per household regarding income. Our empirical strategy helps distinguish optimism from bad luck as drivers of outcomes like HtM status, whereas RS' approach focuses on ex-post errors only. Additionally, RS study the implications of their empirical findings in a partial equilibrium setup, whereas we develop a general equilibrium framework to study the macroeconomic implications of belief heterogeneity.

[9] See, e.g., Woodford (2013), Gabaix (2014), Woodford (2019), Gabaix (2020), Bordalo et al. (2020), Lian (2023), and Boutros (2023).

[10] See, e.g., Farhi and Werning (2019), Auclert et al. (2020), Angeletos and Huo (2021), Laibson et al. (forthcoming), and Pfäuti and Seyrich (2023).

ex-ante identical.[11] Additionally, our model features nominal rigidities and we also take a step beyond the crucial one of matching key empirical moments by demonstrating use cases for our model, in particular, analyzing positive and normative implications of different insurance policies.

A parallel strand of literature considers (persistent) heterogeneity in reduced-form or presumed-classical preferences. Aguiar et al. (forthcoming) find that allowing for heterogeneity in patience and the elasticity of intertemporal substitution helps match several empirical facts about HtM households. They suggest that behavioral factors might provide a potential micro-foundation for their modeling choices. Krueger et al. (2016) and Auclert et al. (2020) introduce permanent heterogeneity in patience to better match wealth inequality data. Kaplan and Violante (2022) show that heterogeneity in risk aversion can produce similar results to heterogeneity in discount factors in terms of HtM shares and MPCs. They also show, however, that allowing for heterogeneity in risk aversion or in discount factors does not solve the standard HANK's "missing middle problem" of producing a wealth distribution that is too polarized. We show that allowing for heterogeneity in optimism, in contrast, fills in the missing middle. Furthermore, our micro data does not favor patience or risk aversion alone as an empirically likely key margin of heterogeneity; e.g., their correlations with HtM status are relatively weak compared to ex-ante optimism, both qualitatively and quantitatively. In addition, all of these papers use the degree of ex-ante heterogeneity as free parameters to match some empirical targets in contrast of our approach to directly discipline the degree of heterogeneity with estimates from micro data.

Outline. We detail our data and empirical findings in Section 2. Section 3 shows how we introduce heterogeneity in ex-ante optimism into HANK models, and Section 4 presents our model's stationary equilibrium results. Section 5 develops fiscal policy implications and Section 6 concludes.

2 Micro Data and Empirical Results

Standard consumption-savings theory predicts that a household's expectation about her own future financial situation is a key determinant of her savings behavior. The more optimistic the household, the less she is expected to save, ceteris paribus. Against this background, we leverage a novel micro panel data set to analyze the question how households perceive their own future financial situation, and whether differences therein are linked to heterogeneity in savings behavior and actual financial situations. We first construct households' financial forecast errors (FCEs) and show that these FCEs display an optimistic bias on average. We then propose

[11] Ilut and Valchev (2023) use MPC estimates to calibrate their key behavioral parameter, whereas we leave the MPC untargeted. We leave the combination of their approach with our focus on permanent heterogeneity for future research.

an empirical strategy to quantify the degree of the optimistic bias that consumers exhibit and the heterogeneity in this bias across consumers. Importantly, we account for the fact that forecast errors may be driven by luck, i.e., by shocks, which implies that forecast errors alone are insufficient to detect and quantify optimistic biases. We then show that optimistic biases are indeed strongly correlated with observed behavior and financial situations, as a more optimistic bias is associated with less savings and a larger likelihood of being financially constrained.

Our empirical strategy requires us to observe both individuals' forecast errors about their own future financial situation as well as behavioral biases. To do so, we use the American Life Panel (ALP), a long-running online panel that goes to great lengths to obtain a nationally representative sample of U.S. adults. In this panel, we observe consumer's financial forecast errors as well as—given the work by Stango and Zinman (2023)—behavioral biases and preferences of the same consumers. Throughout this section, we use both unweighted and sampling probability-weighted estimates to guide inference, given that each approach requires different assumptions to extrapolate from sample-level to population-level inference (Solon et al. (2015)).

2.1 Forecast errors and the strong prevalence of ex-post optimism

We start by constructing individual level forecast errors about consumer's own financial situations. The ALP elicits consumers' forecasts of their own financial situations and subsequent realizations in many of its survey modules. This allows us to build a panel of 21,586 forecast-realization pairs, provided by 3,467 ALP panelists, across fourteen surveys administered in January and July from July 2010 to January 2016. The forecasting question is: "... do you think that a year from now you will be better off financially, or worse off, or just about the same as now?" This question has long been used, by the Michigan Survey of Consumers and many other national household surveys across the world, to help measure consumer sentiment (e.g., Souleles (2004)). Forecasts are highly correlated with expected own-income growth in the ALP surveys that also elicit an income forecast (Appendix Table A3; see also Kamdar and Ray (2024)). Bhandari et al. (2024) and Kamdar and Ray (2024) show that these forecasts also correlate strongly with forecasts about aggregates.[12] We measure realizations a year later with the contemporaneous version of the household financial situation question: "We are interested in how people are getting along financially these days. . . ." to which they can answer "better", "same" or "worse" than a year ago. Comparing consumers' forecasts with their realizations one year later, yields three potential outcomes: consumers can either be ex-post too optimistic, accurate, or too pessimistic about their financial situations. Table 1 illustrates this classification.

[12]Our modeling approach can account for either source of optimism about one's own financial situation—optimism about one's own income, or optimism about aggregates (as well as expense neglect)—as we formalize in Section 4.4.

Table 1: Measuring ex-post forecast errors (FCEs) with coarse data

	Realization at $t+1$		
FC at t	Better	Same	Worse
Better	Accurate	Optimist	Optimist
Same	Pessimist	Accurate	Optimist
Worse	Pessimist	Pessimist	Accurate

Note: The table shows how we classify consumers into three groups based on how their forecasts in period t about period $t+1$ outcomes align with the $t+1$ realizations.

This coarse classification of potential forecast error realizations yields two nuanced challenges that have so far—to the best of our knowledge—not been accounted for in the literature. First, the set of FCEs that are feasible to observe depends on the consumer's forecast and/or realization. For example, it is impossible to observe an optimistic FCE when the original forecast takes the lowest possible value offered as a response option. This is because one would need to observe a realization lower than that value—worse than "worse", in our data—to observe an optimistic FCE. An analogous problem applies for pessimism as it is impossible to observe a pessimistic FCE when the original forecast takes the highest possible value offered as a response option ("better", in our data). Additionally, one can only observe a potentially symmetric FCE—an error that could be either optimistic or pessimistic in direction—if either the forecast or realization takes an interior value ("Same" is the only interior value in our data). To account for this data limitation, we will use two different sample splits ("potentially symmetric", and "potentially optimist/pessimist").

The second nuance is that categorical data constraints can make it especially difficult to detect ex-post optimistic or ex-post pessimistic forecast-realization pairs. Some better-better (some worse-worse) pairs will be misclassified as accurate when they are in fact ex-post optimistic (pessimistic). For example, consider a too optimistic forecast of 20 percent improvement which is then followed by a 10 percent improvement realizations. This would be misclassified as "accurate", because both the forecast and realization are measured as "better". To mitigate this bias towards "accurate" in our sample, we introduce as a third sample split a second version (labeled "v2") of "potentially optimist/pessimists" in which we exclude observations with better-better or worse-worse pairs, respectively.[13] Appendix Table A1 illustrates all of these different sample splits graphically.[14]

[13]Note, that we do not have to worry about the opposite case. For example, a consumer expecting to be 20 percent better off but then only being 10 percent better off will be labeled as "accurate" instead of "pessimist". But we only focus on the distinction between "optimist" vs. "non-optimist" (and similarly for "pessimist" vs. "non-pessimist"), so being accurate or pessimistic will both be classified as "non-optimist".

[14]Appendix Table A1 also shows that the mean count of measurable FCEs per panelist ranges from 4.2 to 6.2 depending on the FCE definition.

2.1.1 Facts about ex-post optimism

We now report some descriptive statistics about the forecast errors. We do so in form of two facts that motivate us to zoom in on optimism in the rest of the section.

Fact 1: Forecasts and forecast errors tilt strongly optimistic on average.
We start by noting that both forecasts and forecast errors for consumers' own financial situation tilt strongly optimistic on average. This mirrors similar findings, from other countries and time periods, in prior work (Souleles, 2004; Claus and Nguyen, 2023; Cocco et al., 2021).[15] Forecasts in our data are more than twice as likely to predict improvement (27 to 30 percent of observations) as deterioration (10 to 14 percent of observations). Using our potentially symmetric sample split, forecast errors are roughly two to three times more likely to be in an optimistic than pessimistic direction.

Optimism is also substantially more prevalent than pessimism *within* consumer. Panels A and B of Table 2 depict the proportion of a consumer's observed FCEs that are optimistic or pessimistic, as well as an indicator that equals one when the consumer makes optimistic/pessimistic FCEs at least half the time, per each of the sample splits. Given the lack of any quantitative FCE measures, we rely on these two measures of optimism/pessimism. Panel C then takes the ratio of sample proportions of optimism and pessimism, or of discretized versions thereof. Across our 12 estimates—(three FCE sample splits) $\times$ (two functional forms of panelist-level measures) $\times$ (weighted vs. unweighted estimates)—optimism is 1.83 to 2.91 times more prevalent than pessimism.[16]

We focus on the relative optimism estimates in Table 2 Panel C more than level estimates of optimism and pessimism in Panels A and B, because the relative estimates are likely more accurate. This is—as discussed above—because coarse data on forecasts and realizations likely misclassifies many modest optimists and modest pessimists as accurate, thereby understating the absolute prevalences of optimism and pessimism. Estimates of relative prevalence nevertheless will be unbiased if measurement error in each level is proportional to its true prevalence.

Fact 2: Optimistic FCEs are persistent within-consumer and more persistent than pessimistic FCEs.
The second fact motivating our focus on optimism, and persistent optimism in particular, is that measured ex-post forecast errors about one's own financial situation are persistent over time, within-consumer. Again using the sample split where both FCEs are potentially symmetric to allow for apples-to-apples consideration of persistent pessimism as well as optimism, we

[15] The one counterexample we know of is Hyytinen and Putkuri (2018)'s evidence of aggregate mean-zero forecast errors from Finland.

[16] For comparison, Columns 7 and 8 in Panel C of Table A5 report relative optimism estimates that are not adjusted for measurement error created by categorical data constraints detailed above. Estimates are uniformly higher for each functional form and (un)weighting combination, ranging from 2.33 to 5.00.

Table 2: Household financial condition FCE persistence, Part 3: Panelist-level measures

	Potentially opt or pess FCEs				Potentially symmetric	
Sample	All		v2		FCEs	
Estimate	Unwtd	Wtd	Unwtd	Wtd	Unwtd	Wtd
Panel A. Prop of FCEs that are optimistic						
Proportion	0.33	0.32	0.39	0.38	0.33	0.32
Share consumers with optimistic proportion ≥ 0.5	0.32	0.28	0.40	0.39	0.32	0.32
N consumers (≥ 2)	2928	2928	2792	2792	2787	2787
Panel B. Prop of FCEs that are pessimistic						
Proportion	0.18	0.17	0.20	0.19	0.15	0.15
Share consumers with pessimistic proportion ≥ 0.5	0.14	0.14	0.17	0.17	0.11	0.11
N consumers (≥ 2)	2542	2542	2454	2454	2787	2787
Panel C. Relative Optimism/Pessimism						
Proportions	1.83	1.88	1.95	2.00	2.20	2.13
Shares ≥ 0.5	2.29	2.00	2.35	2.29	2.91	2.91

Note: N.b. Denominators in Panels A and B include accurate forecasts: see Table 1 for details on sample splits and number of forecast errors per panelist in each sample. Relative proportions simply use the sample estimates in Panel A and B to estimate: (persistent optimism)/(persistent pessimism). Weighted estimates use the mean sampling probability weight across surveys where the panelist provides a financial situation realization.

find that about 60 percent of these consecutive forecast errors are the same (both optimistic, both realistic, or both pessimistic), compared to the zero-persistence benchmark of 33 percent. About 47 percent of panelists who make an optimistic forecast error in the previous period make the same error in the next period. The comparable estimate for pessimism is only 29 percent, suggesting that pessimism is not as persistent. Appendix Table A2 illustrates these results in more detail.

2.1.2 The insufficiency of ex-post forecast errors to detect expectation biases

Having established the much stronger prevalence of optimistic vs. pessimistic forecast errors, we turn towards quantifying the underlying expectation bias at the consumer level. The empirical challenge in doing so is to quantify the degree of optimism that is driven by an ex-ante bias as opposed to ex-post bad luck. The challenge arises because of a lack of a rational benchmark which renders ex-ante optimism unobservable. This is particularly pronounced in our case given that the beliefs of interest regard a consumer-specific variable. In contrast to an aggregate variable, one cannot rely on the use of a professional forecast or a long time series, as a benchmark. In addition, it has long been recognized that the ex-post FCEs which we have considered thus far are not fully dispositive, because in principle even a persistent pattern of ex-post FCEs could be due to a spell of bad shocks (Souleles, 2004). We next propose an empirical strategy to overcome these challenges.

2.2 A strategy for quantifying ex-ante optimism from ex-post FCEs

Ex-post forecast errors reflect the combination of an unobservable ex-ante bias and an unpredictable shock realization. Hence, we can distinguish ex-ante from ex-post optimism by identifying the variation in ex-post optimism that is predicted by an observable behavioral bias. The conditional correlation we establish between the corresponding behavioral bias and ex-post optimism suggests an important role for ex-ante optimism as a driver of ex-post optimism, because, as we argue below, it is highly unlikely that there is a feedback effect from luck, i.e., from shock realizations, to our behavioral bias.

2.2.1 Overconfidence

Overconfidence about one's own abilities—that is overestimating one's own abilities—as elicited, e.g., by Stango and Zinman (2023) or Chapman et al. (2023) is *a priori* a promising candidate for identifying ex-ante optimism. It is reasonable to hypothesize that it meets the requirement because it is conceptually akin to optimism. Both overconfidence and optimism involve overestimating the likelihood of favorable states, and indeed various fields use the two labels interchangeably in describing such a tendency (Johnson and Fowler, 2011; Carver et al., 2010; Spinnewijn, 2015). Importantly, we have fairly granular measures of overconfidence in our data, allowing us to establish a conditional correlation and to quantify heterogeneity in ex-ante optimism. It is also reasonable to hypothesize that overconfidence is not driven by luck, for two reasons. First, it is trait-like. Second, it is prevalent and impactful for decision making in high-stakes settings, suggesting that overconfident consumers do not account for any association between their overconfidence and shock realizations—if they did, overconfident consumers would instead make the same decisions and realize the same outcomes as classically rational consumers.

Our overconfidence measure comes from the Stango and Zinman (2023) (henceforth SZ) ALP modules designed to measure various behavioral biases and their temporal stability. The same SZ surveys were administered in 2014 and 2017, with 845 panelists completing both. About 350 to 400 consumers in our optimism panel complete the SZ surveys as well, providing sufficient sample size for quantifying heterogeneity in ex-ante optimism. We measure overconfidence using the question: "... what you think about your intelligence as it would be measured by a standard test. How do you think your performance would rank, relative to all of the other ALP members who have taken the test?", elicited as an integer percentile. Respondents overestimate their abilities on average, with 70 percent providing a better-than-average percentile. Later in that survey they take a standard 15-question "number series" test of fluid intelligence (McArdle et al. (2007)).[17] We then define a consumer's degree of overconfidence as their self-assessed rank

[17]Number series scores correlate strongly with those from other fluid intelligence tests like IQ and Raven's.

minus the actual rank so that a higher value indicates more overconfidence.[18] This measure exhibits a high degree of stability within-panelist over time, with Stango and Zinman (2023, 2020) estimating serial correlations ranging from 0.51 to 0.96 over three years using various methods and samples. These numbers correspond to persistence estimates of 0.95-0.997 at quarterly frequency.

2.2.2 Estimation

We estimate the component of persistent ex-post optimism that is due to ex-ante optimism using specifications of the form:

$$\text{Prop. optimistic errors}_i = \beta_0 + \beta_1 \cdot oc_i + \Gamma \cdot \text{controls}_i + u_i, . \tag{1}$$

Prop. optimistic errors$_i$ is the proportion of optimistic forecast errors of a consumer as described in Table 2, oc_i is overconfidence, our measure of ex-ante optimism. We discuss a potential concern with using proportions of optimistic forecast errors—due to our lack of quantitative FCE measures—at the end of the section. Despite the warranted strong priors of a positive correlation between Prop. optimistic errors$_i$ and oc_i discussed above, we conduct inference conservatively by using two-sided hypothesis tests with a null of no relationship.

controls$_i$ is a vector of control variables for preferences, cognitive skills, income, and standard demographics (age, gender, education, race, and ethnicity). The controls also facilitate exploration of whether standard preference heterogeneity has a strong empirical association with ex-post optimism (and, importantly, actual financial situations, below).[19]

As detailed in Stango and Zinman (2023, 2020), the overconfidence and the preference variables are likely subject to substantial classical measurement error, as are cognitive skills to a lesser but still potentially meaningful extent, so for each of those variables we use its two elicited measures as instruments for each other (Gillen et al., 2019). Besides accounting for measurement error, this approach has the additional attractive feature of producing estimates based on the stable (persistent) component of each instrumented variable, thereby aligning with our interest in whether and how persistent decision making heterogeneity contributes to consumer outcomes and macro dynamics.[20]

[18] We focus on the relative overconfidence measure in SZ because it is more powerful, both statistically (it is more granular in our data) and conceptually (fluid intelligence is linked more strongly to productivity than numeracy is). The SZ data provides a second measure of (over)confidence about cognitive skills, regarding absolute performance on the numeracy test, that is strongly correlated with our measure of overconfidence in relative performance (Stango and Zinman (2023), Chapman et al. (2023)).

[19] We measure preference heterogeneity using SZ's standard elicitations of patience in the form of financial discounting Andreoni and Sprenger (2012), and risk aversion in the form of quantitative choices over income gambles Barsky et al. (1997) and a financial risk-taking scale Dohmen et al. (2010).

[20] We do not instrument for income, instead using each snapshot as its own uninstrumented control variable, because in our quantitative model, households' income will be subject to idiosyncratic shocks.

Fact 3: Overconfidence, our measure of ex-ante optimism, is strongly correlated with ex-post optimism.

Table 3 confirms that overconfidence is indeed strongly conditionally correlated with ex-post forecast errors about one's own financial situation. The table shows 6 different specifications of equation (1)—(3 sample splits for measuring optimism from ex-post FCEs) × (unweighted and weighted estimates)—and we find a positive coefficient in all 6 cases. Five of these six specifications are significant at conventional levels, despite the modest sample sizes.

Table 3: Overconfidence strongly conditionally correlates with ex-post optimism

LHS = Optimism FCE proportion						
	Potentially Optimistic				Potentially Symmetric	
	v1		v2			
	Unwtd	Wtd	Unwtd	Wtd	Unwtd	Wtd
	(1)	(2)	(3)	(4)	(5)	(6)
Overconfidence	0.25	0.25	0.30	0.72	0.24	0.64
	(0.11)	(0.21)	(0.13)	(0.37)	(0.12)	(0.36)
Patience	-0.26	-0.41	-0.34	-1.32	-0.32	-1.21
	(0.16)	(0.33)	(0.19)	(0.97)	(0.18)	(0.94)
Financial Risk Aversion	0.03	0.14	0.01	0.02	-0.00	0.00
	(0.09)	(0.15)	(0.10)	(0.20)	(0.10)	(0.18)
Income Gamble Risk Aversion	-0.06	-0.72	-0.15	-0.48	-0.17	-0.57
	(0.17)	(0.30)	(0.19)	(0.52)	(0.18)	(0.55)
Other Controls?	Yes	Yes	Yes	Yes	Yes	Yes
N	778	778	750	750	742	742
N Panelists	389	389	375	375	371	371
Mean (LHS)	0.36		0.44		0.38	

Note: Standard errors, clustered on panelist, in parentheses. One regression per panel-column. LHS variable is described in the panel and column headings, and the last row. RHS variables measured using the Stango-Zinman data and merged onto the our household financial situation forecast error panel. The results shown represent the [100] x the change in the LHS variable associated with a 1 percentage point increase in the RHS variable (each RHS variable here is a percentile rank). "Other control variables" include standard demographics (income, education age, gender, race, and ethnicity) and cognitive skills. For RHS variables likely subject to measurement error and where we are interested in their permanent component– overconfidence, patience, risk aversion, and cognitive skills– we use obviously related instrumental variables (ORIV) to account for measurement error .

Quantitatively, the estimates in Table 3 are quite substantial. Being one percentile higher up in the overconfidence distribution is associated with a 0.24 - 0.72 percentage point increase in the likelihood of making an optimistic forecast error. This corresponds to estimates of the interquartile range between 12 and 36 percentage points higher likelihood of making optimistic forecast errors. The average across estimates in Table 3 corresponds to an interquartile range of 20 percentage points. This slope in the increase in making overly optimistic forecast errors will be a key target to discipline our model later.

For robustness, Panel B of Appendix Table A6 considers an alternative definition of ex-post optimism—namely, a dummy variable that equals 1 if the consumer makes optimistic FCEs at least half the time—adding 6 more specifications of (1). Again we find a positive coefficient in

all cases and 5 out of 6 of these additional estimates are significant further reassuring us that overconfidence is significantly and meaningfully correlated with ex-post optimism.

If we run the same regressions but with the proportion of pessimistic forecast errors instead of optimistic ones as our dependent variable (not shown), we find that the overconfidence coefficient is negative in all six variations of regression (1). The estimates, however, are not significant which might be due to the fact that we do not observe that many pessimistic forecast errors. Yet, it renders further support that overconfidence is in fact related to ex-ante optimism and not just to an overall tendency of having misperceived beliefs.

Standard preferences, in contrast, have statistically weaker correlations with ex-post optimism. Patience has zero p-values < 0.05. The two risk aversion measures often have different-signed coefficients, and only one of 18 p-values < 0.05 between them.

We next show that ex-ante optimism influences actual decisions of consumers. To this end, we run regression (1) but with measures of households' behavior and their actual financial situation as dependent variables. In particular, we use an an indicator for not saving over the prior year on a flow basis (including debt paydown as saving), and an indicator for having experienced severe financial distress over the past year, which we label as our measure of HtM status or of being financially constrained.[21]

Fact 4: Ex-ante optimism strongly conditionally correlates with actual saving's behavior and actual financial situations.

Table 4 shows that overconfidence conditionally correlates strongly with our two measures of financial decision making and financial constraint. The magnitude of overconfidence's conditional correlations with actual financial situations is substantial, suggesting that beliefs, in the form of ex-ante optimism, meaningfully affect saving decisions and financial outcomes. A one percentile increase in the overconfidence distribution is associated with a 0.26 - 0.61 percentage point decrease in the likelihood of having saved and a 0.14 - 0.52 increase in the likelihood of being financially constrained. These numbers correspond to interquartile ranges of 13 - 31 percentage points increases in the likelihood of not having saved and 7 - 26 percentage points

[21] We measure saving with a standard question used in the Survey of Consumer Finances and other household finance surveys: "Over the past 12 months, how did your household's spending compare to your household's income? If the total amount of debt you owe decreased, then count yourself as spending less than income. If the total amount of debt you owe increased, then count yourself as spending more than income. [Response options: Spent more than income/Spent same as income/Spent less than income.]" We define severe financial distress as indicating that any of four events happened in the previous 12 months: forced move, late payments, hunger, or foregone medical care. We explicitly focus on this measure of HtM as it is backward looking. Using forward-looking measures, such as answers to "Could you cover an unexpected expense of $\$x$?" might lead to biased estimates because optimistic households may be too optimistic about their ability to cover such expenses. Similarly, using liquidity-based measures are not applicable given that the ALP only elicits such measures using questions like "If you would sell all your assets and repay all your debt, how much money would you have left?". Again, optimistic households are likely to overstate for how much they could sell their assets. Our financial distress measure does not suffer from these issues.

increases in the likelihood of being financially constrained.

Heterogeneity in standard preferences have oftentimes been used as alternative explanations for savings heterogeneity (Krueger et al., 2016; Carroll et al., 2017; Aguiar et al., forthcoming). Interestingly, the overall pattern of results in Table 4 seems to favor overconfidence over standard preferences as the key margin of ex-ante heterogeneity. Three of the four estimates are significant at the 5% significance level. Meanwhile only three of the 12 estimates on standard preferences are significant, with the coefficients on the two risk aversion measures exhibiting opposite signs. For patience, at least the sign is constantly negative (and consistent with theory), but only one out of the four estimates is significant.

Table 4's results are interesting for several, intertwined reasons. First, they further motivate modeling ex-ante heterogeneity in biased beliefs also instead of ex-ante preferences. Second, our quantitative model will produce the key results in Table 5 endogenously, with empirically realistic prevalence of constrained households as untargeted moments. And third, they help us to rule out competing theories of our empirical findings, as we now discuss.

Table 4: Overconfidence, our measure of ex-ante optimism, strongly conditionally correlates with actual financial situation

	1 = (Did not save last year)		1 = (HtM)	
	Unwtd	Wtd	Unwtd	Wtd
	(1)	(2)	(3)	(4)
Overconfidence	0.26	0.61	0.14	0.52
	(0.11)	(0.22)	(0.11)	(0.25)
Patience	-0.51	-0.33	-0.16	-0.36
	(0.18)	(0.24)	(0.17)	(0.27)
Financial Risk Aversion	0.01	0.30	0.12	0.20
	(0.09)	(0.17)	(0.09)	(0.17)
Income Gamble Risk Aversion	-0.16	-0.31	-0.44	-0.22
	(0.18)	(0.36)	(0.17)	(0.41)
Other Controls?	Yes	Yes	Yes	Yes
N	1355	1355	1358	1358
N Panelists	680	680	680	680
Mean (LHS)	0.57		0.37	

Note: Same regression specification as Table A6 but with different LHS variables (please see Section [] for details on how we define those).

Potential concern. A potential concern that arises when using the proportion of optimistic forecast errors as our dependent variable, is that there might be an association between overconfidence and ex-post shock realizations because, for example, more overconfident households face more risk (e.g., by selecting into riskier jobs). This seems empirically unlikely given that we control for a wide set of consumer-level controls, including risk aversion, education, cognitive

skills, and income. But even if our empirically-established link between overconfidence and ex-post optimism is due to overconfident consumers facing different risks, our empirical findings strongly suggest that it is highly unlikely that these consumers are fully rational about this when making their consumption-savings decision. As we formalize in Section 4, the correlation between ex-post optimism and overconfidence requires that overconfident households face *more* risk than the non-overconfident households. If overconfident households face more risk and are fully rational about it, however, they would save more (and be less likely to be HtM) than the non-overconfident households, which is at odds with our empirical findings. In sum, the most likely mechanism to explain our empirical findings jointly is that overconfidence reflects an ex-ante optimistic bias.

2.3 Taking stock and outlook

This section shows that, first, many consumers often and persistently make optimistic forecast errors about their own future financial situation. Second, these forecast errors can be partly explained by ex-ante optimism and the degree of this ex-ante optimistic bias is heterogeneous across households. Third, optimistic bias is associated with being less likely to save and a higher likelihood of consumers being financially constrained. Preference heterogeneity—be it in patience or risk aversion—is, on the other side, not strongly associated with lower savings or financial constraints. To understand the macroeconomic consequences of heterogeneity in optimism, we next build a heterogeneous-agent New Keynesian model augmented by heterogeneity in the degree of optimism.

3 Model

We now develop an augmented HANK model to which we introduce belief heterogeneity disciplined by our findings in the previous Section 2. Besides belief heterogeneity, the model is otherwise standard: it features incomplete markets in the spirit of Bewley (1986), Huggett (1993), and Aiyagari (1994), and nominal rigidities in the form of sticky wages. Time is discrete and denoted by $t = 1, 2, ...$. We first focus on the case in which households can only save in one asset—a liquid bond issued by the government. Later on, we introduce a second asset in the form of illiquid productive capital.

Households. There is a unit mass of households subject to idiosyncratic risk, incomplete markets, and borrowing constraints. We allow for permanent, ex-ante, heterogeneity in how households form their expectations about their own future financial situations but assume that they are otherwise ex-ante identical.[22] An individual household's productivity is denoted by

[22] We assume that belief heterogeneity is permanent given the results in Stango and Zinman (2020). Consistent with that, Hoffman and Burks (2020) also find that workers' over-optimistic beliefs about their productivity

e_t which follows a Markov process with time-invariant transition matrix $\mathcal{P}$. The process for e_t is the same for all households and the mass of households in state e is always equal to the probability of being in state e in the stationary equilibrium, $p(e)$.

The problem of an individual household of type g in idiosyncratic state e_t, with beginning-of-period asset holdings b_{t-1}, is given by:

$$V_{g,t}\left(b_{t-1}, e_t\right) = \max_{c_t, b_t} \left\{ \frac{c_t^{1-\gamma}}{1-\gamma} - \frac{n_t^{1+\varphi}}{1+\varphi} + \beta \tilde{\mathbb{E}}_{g,t} V_{g,t+1}\left(b_t, e_{t+1}\right) \right\}$$

subject to

$$c_t + \frac{b_t}{1+r_t} = b_{t-1} + (1-\tau_t) w_t e_t n_t \tag{2}$$

$$b_t \geq -\underline{b}, \tag{3}$$

where c_t denotes consumption, n_t hours worked, r_t the net real interest rate, w_t the real wage, τ_t the income tax rate, and V the value function. We assume a standard CRRA utility function where the parameters γ, φ, and β denote relative risk aversion, the inverse Frisch elasticity of labor supply, and the time discount factor, respectively. These parameters as well as the exogenous borrowing limit $\underline{b}$ are the same for all households and time-invariant. The expectations operator $\tilde{\mathbb{E}}_{g,t}$ is our key innovation, and we discuss it next.

Belief heterogeneity. We allow for belief heterogeneity to capture our empirical result regarding heterogeneity in households' belief about their own future financial situation. We equate "financial situation" with the right-hand side of the budget constraint, $b_{t-1} + (1-\tau_t) w_t e_t n_t$. However, we still need to make a choice of the object of over-optimism, i.e., about which part of the right-hand-side of the budget constraint (some) households are overly-optimistic about. In our baseline, we focus on over-optimism about one's idiosyncratic productivity. Focusing on idiosyncratic productivity makes sense conceptually, since our proxy for optimism—overconfidence—concerns cognitive skills, such skills have well-established links to income and workplace performance. Focusing on productivity also makes sense empirically, given that forecasts are highly correlated with expected own-income growth in the ALP surveys that also elicit an income forecast (Appendix Table A3; see also Kamdar and Ray (2024)). And third, uncertainty about productivity is usually the only source of uncertainty in standard incomplete markets models. We present two alternative approaches in Section 4.4.2. In the first one, we assume that over-optimism about one's own income is instead stemming from over-optimism about aggregates. Second, we also consider that over-optimism about future financial situation

are very persistent. We nevertheless run a robustness check in which we assume that the behavioral biases of households is not permanent but that households can switch between types. Appendix D.3 shows that our results are robust in this case.

is coming from neglect of expense shocks. In both cases, our results remain robust.

We model optimism about idiosyncratic productivity as biased beliefs about the transition probabilities $p(e_{t+1}|e_t)$. Optimistic households assign too much probability to reaching (or staying in) relatively high-skill states, and too little probability to reaching (or staying in) relatively low-skill states. This implies that, on average, optimistic households overestimate their future productivity. Note, that absent aggregate risk, a household's over-optimism about future productivity necessarily implies over-optimism about her future financial situation as idiosyncratic productivity is the only source of uncertainty.

Let $p_{ij} \equiv p(e_{t+1} = e_j|e_t = e_i)$ denote the probability that a household with current idiosyncratic productivity $e_i \in \{e_1, e_2, ..., e_J\}$ reaches productivity $e_j \in \{e_1, e_2, ..., e_J\}$ in the following period, and assume that the productivity states are ordered such that $e_1 < e_2 < ... < e_J$. To capture optimism with only one additional parameter independent of the number of individual states, we assume that a household's *perceived* transition probabilities $\tilde{p}_{ij}$ are given by

$$\tilde{p}_{ij} \equiv \begin{cases} \alpha_g p_{ij}, \text{ if } i < j \\ \frac{1}{\alpha_g} p_{ij}, \text{ if } i > j \\ 1 - \sum_{j \neq i} \tilde{p}_{ij}, \text{ if } i = j, \end{cases} \tag{4}$$

where the parameter $\alpha_g \geq 1$ captures the degree of the optimistic bias in group g. If $\alpha_g > 1$, a household assigns too much weight to reaching a better state (this is the case $i < j$) and too little weight to reaching a worse state ($i > j$). The perceived probability of staying in the same state ($i = j$) ensures that the probabilities sum to 1.[23] Note that rational expectations are captured by setting $\alpha_g = 1$ and thus nested in our setup.[24] In Section 4.4, we discuss an extension with pessimistic households, $\alpha_g < 1$ and, as mentioned above, alternative sources of optimism, namely optimism about aggregates as well as neglect of expense shocks. Additionally, Appendix D.3.1 provides an alternative modeling approach, where the degree of optimism depends on the distance between the states.

An immediate implication of equation (4) is that optimistic households will more often make overly-optimistic forecast errors about their own financial situation compared to rational households, consistent with the empirical findings reported in Section 2.2.2. In turn, we can discipline the degree in belief heterogeneity, that is the group specific values of α_g, by directly targeting our empirical estimates in Table 3.

[23] We further restrict α_g such that all perceived transition probabilities lie between 0 and 1. Given a standard calibration for the income process, this restriction is never binding.

[24] Modelling optimism as in (4) is similar to the way Caballero and Simsek (2020) model optimism about an aggregate state with two possible realizations. In contrast to them, we focus on idiosyncratic states and allow for an arbitrary number of realizations. McClung and Nighswander (2021) introduce belief heterogeneity about idiosyncratic employment transition probabilities into a life-cycle model, but consider only two possible states.

Unions. We follow the recent HANK literature and assume that hours worked n_t are determined by union labor demand and that wages are sticky whereas prices are flexible (see especially Auclert et al. (forthcoming), which is based on Erceg et al. (2000)).[25] Each worker provides $n_{k,t}$ hours of work to a continuum of unions indexed by $k \in [0,1]$. Each union aggregates efficient units of work into a union-specific task

$$N_{k,t} = \int \bar{e}_i e_{i,t} n_{i,k,t} di,$$

where i here denotes an individual household carrying its permanent type and its idiosyncratic state.

A competitive labor packer then packages these tasks into aggregate employment services according to the CES technology

$$N_t = \left(\int_k N_{k,t}^{\frac{\epsilon-1}{\epsilon}} \right)^{\frac{\epsilon}{\epsilon-1}} \tag{5}$$

and sells these services to firms at price w_t.

We model wage stickiness by imposing a quadratic utility cost $\frac{\psi}{2} \int_k \left(\frac{W_{k,t}}{W_{k,t-1}} - 1 \right)^2 dk$ that shows up in the household's utility function. A union sets a common nominal wage $W_{k,t}$ 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. We assume the union ensures that each household supplies the same amount of hours.

Firms. A representative firm operates an aggregate production function which is linear in labor input N_t

$$Y_t = N_t, \tag{6}$$

to produce total output Y_t. Prices are fully flexible such that the real wage per efficient hour is constant

$$w_t = 1. \tag{7}$$

Profits are zero. Since the nominal wage is given by $W_t \equiv w_t P_t = P_t$, we have

$$1 + \pi_t = 1 + \pi_t^w, \tag{8}$$

where $\pi_t \equiv \frac{P_t}{P_{t-1}} - 1$ denotes goods price inflation, and $\pi_t^w \equiv \frac{W_t}{W_{t-1}} - 1$ wage inflation.

[25] Auclert et al. (2023) and Broer et al. (2020) argue in favor of using sticky wages rather than sticky prices in HANK models.

Fiscal policy. We abstract from government spending and assume that the fiscal authority sets total taxes minus transfers, T_t, following a simple debt feedback rule

$$T_t - \bar{T} = \vartheta \frac{B_t - \bar{B}}{\bar{Y}}, \tag{9}$$

where $\bar{T}$, $\bar{B}$ and $\bar{Y}$ denote the stationary equilibrium values of taxes, government debt and output, respectively. Furthermore, the government budget constraint is given by

$$B_t + T_t = (1 + r_t)B_{t-1}. \tag{10}$$

Monetary policy. The monetary authority directly controls the real rate r_t and we assume that they keep it constant at its steady state value r. This assumption only matters when we consider aggregate shocks, as we do when examining how optimistic consumers change the effectiveness of temporarily increasing fiscal transfers in Section 5.1.

Equilibrium. Absent aggregate shocks, and given an initial price level P_{-1}, initial nominal wage W_{-1}, initial government debt B_{-1}, and an initial distribution of agents $\Psi_{g,0}(b_{-1}, e_0)$ in each fixed group g, a general equilibrium is a path for prices $\{P_t, W_t, \pi_t, \pi_t^w, r_t, i_t\}$, aggregates $\{Y_t, C_t, N_t, B_t, T_t\}$, individual allocation rules $\{c_{g,t}(b_{t-1}, e_t), b_{g,t}(b_{t-1}, e_t)\}$ and joint distributions of agents $\Psi_{g,t}(b_{t-1}, e_t)$ such that households optimize (given their beliefs), all firms optimize, unions optimize, monetary and fiscal policies follow their rules, and the goods and bond markets clear:

$$\sum_{g,e} \mu_g p(e) \int c_t \Psi_{g,t}(b_{t-1}, e_t) = Y_t \tag{11}$$

$$\sum_{g,e} \mu_g p(e) \int b_t \Psi_{g,t}(b_{t-1}, e_t) = B_t, \tag{12}$$

where μ_g denotes the mass of agents of type g.

3.1 Calibration

Table 5 summarizes our baseline calibration. One period in the model corresponds to a quarter. We calibrate the standard parameters to values often used in the literature. For the idiosyncratic productivity process, we follow McKay et al. (2016): We assume that log-productivity follows an AR(1) process with autocorrelation of $\rho_e = 0.966$ and a variance of $\sigma_e^2 = 0.033$. We then discretize this process into an eleven-states Markov chain using the Rouwenhorst (1995) method. The average productivity level is normalized to 1. We set the discount factor, β, to match a steady state real interest rate of 4% (annualized). Risk aversion is set to $\gamma = 2$, the inverse

Frisch elasticity to $\varphi = 2$, and the borrowing limit to $\underline{b} = 0$ (as, e.g., in McKay et al. (2016)). We set the average wealth to average annual income ratio to its empirical counterpart of 4.1 (Kaplan and Violante (2022)).

To capture our empirical findings on heterogeneity in optimism, we assume that there are three different groups which only differ in their degree of optimism measured by α_g. We assume that 50% of households are *r*ational households and that 50% of households have optimistic-biased beliefs.[26] We subdivide the optimistic households into *m*ild optimists and *s*trong optimists to allow for heterogeneity among optimistic households. This means that the 50% rational households in the model correspond to the bottom 50% of the overconfidence distribution in our data, the mild optimists to the consumers in the 50-75th percentile range and the strong optimists to the 25% most overconfident consumers. Setting $\alpha_r = 1$ and, thus, to rational beliefs, leaves us with calibrating the degree of optimism for mild and strong optimists. We do so by targeting the slope of how much the likelihood of making optimistic forecast errors about their financial situation one-year ahead increases along the optimism-distribution (Table 3). Concretely, we match the difference of the average proportion of overly-optimistic forecast errors of the mild (strong) optimistic households minus the average of rational households in the model to the difference of the groups medians in the data, that is, the proportion of overly-optimistic forecast errors of the 62.5th percentile for mild optimists (87.5th for strong optimists) minus the 25th percentile. For the sake of a conservative calibration we target the lower bound of these estimates, that is, a 10 and a 15 percentage point increase in that likelihood. Given our income risk process, this results in $\alpha_m = 1.6$ and $\alpha_s = 2.1$. This calibration implies that mild optimists have on average a 10% higher expectation and strong optimists an 18% higher expectations about their income within the next year than rational households. While we do not have data to explicitly check these quantitative magnitudes of over-optimism, we think that these numbers are fairly reasonable.

4 Belief Heterogeneity and Fit to the Data

We now show our model's ability to fit various key moments from macro and micro data, in contrast to the rational HANK model that abstracts from belief heterogeneity.

4.1 Heterogeneity across observationally equivalent households

We first show how belief heterogeneity affects households' behavior at the micro-level. Belief heterogeneity introduces heterogeneity across observationally equivalent households, that is, households with the same level of wealth and income. The upper panel in Figure 1 shows

[26] Given that our empirical estimates strongly tilt towards optimism, we do not include pessimists in our baseline. We however run an extension in which we also include pessimists, see Section 4.4.

Table 5: Stationary equilibrium calibration

Parameter	Description	Value
R	Steady state real rate (annualized)	4%
γ	Risk aversion	2
φ	Inverse of Frisch elasticity	2
$\underline{b}$	Borrowing limit	0
$\frac{\bar{B}}{4\bar{Y}}$	Average wealth to average income	4.1
Idiosyncratic risk		
ρ_e	Persistence of idiosyncratic risk	0.966
σ_e^2	Variance of idiosyncratic risk	0.033
Permanent heterogeneity		
μ_g	Mass of households	$\{0.5, 0.25, 0.25\}$
α_g	Degree of optimism	{1, 1.6, 2.1}

Note: Calibration summary for our baseline model using three groups to capture permanent heterogeneity.

the savings policy functions of the three permanently different household groups along the wealth distribution for three different income state levels. It shows that conditional on the idiosyncratic state—households' current income and wealth—the net savings of a household (weakly) decreases in her optimism. Most strikingly, there are income and wealth levels at which rational households accumulate more savings as a buffer stock for future shocks while mild or strong optimists already dis-save, thereby tapping in their buffer stocks. Optimistic households have a higher expectation about their future financial situation compared to rational households due to their over-estimation of their expected income—the more optimistic they are the more so. Consequently, they perceive their insurance needs to be lower than rational households do. Over time, they thus accumulate less precautionary savings than rational households although they face the same income process. As a result, more optimistic households are ex-post more likely to hit the borrowing constraint, in line with our empirical findings. We quantify the differences in HtM status across the groups in the next section.

The lower panel shows the MPCs of households along the wealth distribution for the same three income states. Again, our model produces heterogeneity across observationally equivalent households in their marginal propensity to consume. The more optimistic households are, the higher their MPC even conditional on the idiosyncratic state. These differences can be quite strong: For example, at the median income level, optimistic households can have highly elevated marginal propensities to consume, while observationally equivalent rational households have an MPC of close to zero even if they do not hold any assets. The reasons are twofold: first, more optimistic households can be off their Euler equation even in relatively high income states as they still expect their future financial situation to be better. But even if they are not constrained, more optimistic households perceive a transfer of a given size as relatively smaller

compared to their expected life-time income. Thus, conditional on their idiosyncratic state, they are in the more concave part of their policy function compared to rational households who expect a lower life-time income. This "latent" MPC heterogeneity reflects the empirical findings in Lewis et al. (2024) which find that a lot of MPC heterogeneity is likely due to unobserved, maybe behavioral, variables.[27]

Alternative explanation through risk. We are now able to tackle a potential concern raised in the empirical section. The concern is that in theory, an alternative explanation why overconfident households are more likely to do overly-optimistic forecast errors is because they face higher risk instead of being ex-ante overly-optimistic. In principle, given the log-normality of the shock process, it is possible to target the same increase in the likelihood of doing overly-optimistic forecast errors driven by different idiosyncratic shock volatilities. If these households are rational and understand their higher volatility in their income fluctuations, however, the model predicts that these households save *more* and are less likely to be HtM. For example, by increasing the income volatility by the factor 9 and 12, we increase the likelihood of making optimist forecast errors by 10 and 13pp. In this case, "mild optimist" have 1.7 times and "strong optimists" 1.9 times as much savings compared to rational households on average (see Appendix Figure D4 for the savings policy functions).[28] In addition, the share of mild and strong optimist households that are HtM are both 0%, while 19% of rational households would be HtM in such an economy. Empirically, however, we find that more optimistic households are instead less likely to save and more likely to be HtM (see Table 4). Hence, we conclude that even if these households faced higher risk, they are probably not rational but overly-optimistic about it.

4.2 Hand-to-Mouth Shares and Average MPCs

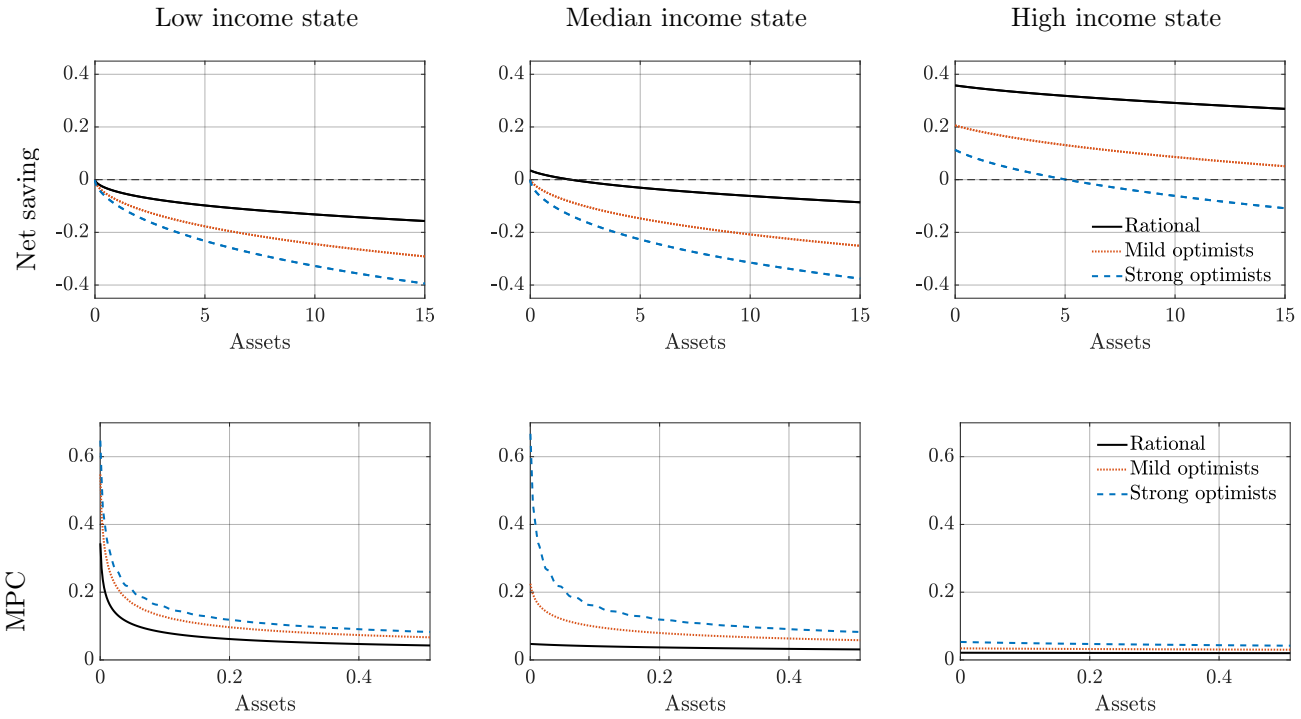
We now analyze the effect that the heterogeneous behavior of households has on the aggregate share of Hand-to-Mouth (HtM) households and the implied average marginal propensity to consume (MPC) of households. We do so by comparing the HtM shares and the average MPC of our baseline model with its rational counterpart which assumes that all households are rational but which is otherwise identical.[29]

[27] Koşar and Melcangi (2024) find that the marginal propensity to consume is increasing in the subjective earnings growth uncertainty of households, at odds with most workhorse models. However this is also consistent with our model given that we find that mild optimists have a 3% higher subjective uncertainty (measured as their perceived earnings volatility) and strong optimists a 9% higher subjective uncertainty and optimism is positively correlated with MPCs, as we discuss in more detail later.

[28] Note, that "optimists" here are the ones facing higher risk and they make more overly-optimistic forecast errors because of the log-normality, not because they hold biased beliefs.

[29] When comparing our model to its rational counterpart, we take the standard approach of recalibrating the discount factor such that both models have the same asset supply and the same steady-state real interest rate (see, e.g., Kaplan and Violante (2022)). The rest of the calibration is the same for both models.

Figure 1: Households' savings policy function and marginal propensities to consume



Note: The upper panels show the savings policy functions for the three different groups—rational households, mild optimists, and strong optimists—along the asset distribution exemplary for three different income (productivity) states. Given that we have 11 productivity levels, we choose the third for low, the sixth is the median and the ninth for high income. The lower panel shows the same for the MPC. The horizontal axis is different than in the upper panel to focus on the highly concave part of the MPC functions but even at higher wealth levels, the ordering of the three groups remains the same.

Table 6 compares the two models. Column 2 reproduces the well-documented finding that rational one-asset HANK models calibrated to match average wealth produce an average MPC and aggregate HtM share that are both far below consensus estimates (Auclert et al. (forthcoming), Kaplan and Violante (2022)). The reason is that rational households have a strong incentive to self-insure themselves against their idiosyncratic risk by accumulating wealth. Thus, with a high enough supply of wealth in the economy, almost no households end up at the borrowing constraint.

Table 6: MPCs and shares of HtM households across the models.

	HANK w/ Belief Het. (1)	Rational HANK (2)
HtM Share	0.23	0.02
Avg. MPC	0.18	0.04
HtM rational HHs	0.02	0.02
Avg. MPC rat. HHs	0.02	0.04
HtM mild optimists	0.24	-
Avg. MPC mild optimists	0.21	-
HtM strong optimists	0.68	-
Avg. MPC strong optimists	0.47	-

Note: MPCs refer to MPCs out of a $500 dollar stimulus check. "HANK w/ Belief Het." is our baseline model (one-asset, with heterogeneity in optimism); "Rational HANK" denotes its rational counterpart that abstracts from heterogeneity in optimism.

In contrast, our model with belief heterogeneity (Column 1) produces an average MPC and a HtM share that are both multiple times larger than in the rational model. Our predictions align well with consensus estimates, albeit more obviously so for the MPC. For example, Jappelli and Pistaferri (2010) and Havranek and Sokolova (2020) report average MPC estimates in the range of 15-25% over a quarterly time horizon, as compared to our 18%.[30] Our predicted share of HtM households, 0.23, is lower than our estimated empirical share of 0.37 (Table 4). However, given the strict definition of HtM status in the model as well as our focus on "poor HtM", for now, this discrepancy is not surprising. Decomposing the aggregate HtM share and MPCs shows that, in line with our discussion in the previous section, it is the mildly optimistic households and especially the strong optimists who pushes up the HtM share and the average MPC: the more optimistic households are, the more likely they end up being HtM, in line with our empirical findings in Section 2. While almost no rational households are HtM (2%), 24% of the mild optimists are HtM and even 68% of the strong optimists are.

The higher share of HtM households also results in a higher average MPC for the group of optimistic households which is further pushed up by their higher MPCs even away from the

[30]Figure D3 in the Appendix shows the intertemporal MPCs in our baseline model. In line with the recent literature, MPCs are large on impact and stay elevated for some time (see for example Auclert et al. (2024b)).

borrowing constraint. While rational households have an MPC of 2%, mild optimists have an average MPC of 21% and strong optimists of 47%.

4.3 "Missing Middle Problem" and wealth shares

Rational one-asset HANK models can generate a high average MPC by restricting wealth to be substantially lower than consensus estimates (Wolf (forthcoming), Kaplan and Violante (2022), Seidl and Seyrich (2023), Auclert et al. (forthcoming)). This restriction also produces an excessively polarized wealth distribution (Kaplan and Violante, 2022). One way to see this "Missing Middle" problem is that median wealth to mean annual earnings is about an order of magnitude smaller in rational HANK models than in the data. We offer further confirmation of this finding by recalibrating the standard HANK model used in Table 6 Column 2 to match the average MPC produced by our one-asset model with belief heterogeneity. Matching the average MPC requires setting total wealth to income to 0.6 instead of 4.1, and delivers a median wealth-to-average annual income ratio of 0.1 whereas empirical estimates are around 1.5 (Kaplan and Violante, 2022).

Our one-asset model with belief heterogeneity overcomes this shortcoming and fills in the missing middle: it predicts a median wealth-to-average annual income ratio of 1.2 that is much closer to its empirical counterpart of 1.5. Rational households that have experienced several periods of relatively low productivity make up most of the middle of our wealth distribution. Optimistic households tend to be HtM and thus account for most of the bottom of the distribution, as discussed above. Rational households that have not experienced long spells of bad productivity shocks populate the top of the distribution. Although not targeted, our model predicts that the top 10% of households hold 46% of wealth, as compared to the empirical estimate of 49% (Kaplan and Violante, 2022). We conclude that, overall, our model produces a wealth distribution that matches the data well without explicitly targeting it.

4.4 Extensions and alternative approaches

We first show that our results are robust to: (i) accounting for pessimistic household, and (ii) for alternative micro-foundations of over-optimism. We then discuss the relationship of belief heterogeneity to discount factor heterogeneity.

4.4.1 Pessimistic households

In our baseline, we assume that households are either rational or optimistic motivated by the fact that forecasts and forecast errors strongly tilt towards optimism. We now add pessimistic households to our model to check whether this would change our results. In particular, we now split the group of rational households into 25% pessimistic households and 25% rational

households. We set $\alpha_p = 1/\alpha_m$ and leave the other parameters unchanged (except for the discount factor which we recalibrate to match the same steady state real interest rate of 4% annually).

Incorporating pessimistic households increases the overall HtM share from 23% to 28% and the average MPC from 18% to 21%. In partial equilibrium, one would see effects in the opposite direction, because pessimistic households overestimate their precautionary savings motive compared to rational households. This pushes pessimistic households to save more than the rational households they are replacing in the model. The pessimists are then less likely to end up HtM. Yet, HtM probability is quite low for rational households to begin with. In general equilibrium, the added savings demand from pessimistic households pushes up the asset price, crowding out savings from the larger mass of households close to the borrowing constraint and, consequently, pushing these households to the borrowing constraint.

4.4.2 Other sources of over-optimism

Given the strong correlation between households' optimism about their future financial situation and expectations about their future income, we assume as our baseline approach that the optimistic bias is rooted in over-optimism about future income. Another potentially meaningful contributor to over-optimism about one's financial situation could be expense neglect. Berman et al. (2016) and Kaur et al. (2025) find that expense neglect is prevalent, heterogeneous, and affects decision making. Fulford and Low (2024) find that expense shocks are prevalent, heterogeneous, and likely quantitatively important for key objects in macro.

Accounting for expense shocks, the budget constraint reads:

$$c_t + \frac{b_t}{1+r_t} = b_{t-1} + (1-\tau_t) w_t e_t n_t - \Xi_t(e_t), \tag{13}$$

where $\Xi_t(e_t)$ are expense shocks that do not enter the households' utility. We assume that expense shocks are a function of the productivity of households to account for the relationship that higher income households have potentially higher expenses.

To focus on the role of expense neglect, we now assume that all households are fully rational with respect to their idiosyncratic productivity, but that households' expectations about expense shocks are potentially biased. Expense shocks are i.i.d. and occur with probability λ_Ξ. To capture belief heterogeneity, we assume that households expect them to occur with probability $\alpha_g \lambda_\Xi$ where, again, α_g is group specific.

For simplicity, we assume that expense shocks occur in every period ($\lambda_\Xi = 1$) and take the value $\Xi_t(e_t) = \xi e_t$. We set $\xi = 0.25$.[31] We set $\alpha_m = 0.8$ and $\alpha_s = 0.6$, which implies an increase

[31] We abstract from very large (but infrequent) shocks because we do not allow households to default in our model. To account for these differences, we therefore target somewhat higher values than found in Fulford and

in the likelihood to make an overly-optimistic forecast error of 10pp and 15pp, respectively.

The main take-aways from the model with expense neglect are the same as for our baseline model: The more optimistic households are about their future financial situation, the more they consume and the less they precautionary save. As a result, overly-optimistic households are more likely to end up being hand-to-mouth and they have higher MPCs even conditional on wealth and income. This makes sense as, in principle, it does not matter what part of "future financial situation" households are overly-optimistic about as households only care about their future cash-on-hand when deciding on their precautionary savings. But different underlying micro-foundations of the source of the uncertainty about future cash-on-hand are extrapolated differently to expectation biases about cash-on-hand further out than a year from now when targeting the same bias in forecast errors about one year from now. In case of over-optimism about future idiosyncratic productivity, given the persistency in the underlying productivity process, households' expectation bias grows with the horizon ("a better job might lead to an even better job down the road"). For expense neglect, this channel is not present. In contrast, the lack of persistency in the expense shock implies a very conservative extrapolation of the optimism bias leading to a lower average MPC conditional on matching the same optimistic bias about one-year ahead forecasts. Because of this, the average MPC in the model with expense neglect is 4%, which is twice as large as in the rational model with expense shocks,[32] but lower than in our baseline model. However, if we target our average empirical estimates that mild optimists are 15pp and strong optimists are 25pp more likely to do overly optimistic forecast errors, the model produces an average MPC of 15%, close to our baseline MPC.

A third source of over-optimism might be that optimists are overly-optimistic about the future aggregate wage, w_t. To model this, we assume that optimists expect a higher wage from $t+1$ onwards, $E_t((1+\alpha_g)w_{t+1})$. We calibrate $\alpha_m = 0.075$ and $\alpha_s = 0.135$ such that again mild (strong) optimists are 10 (15) percentage points more likely to make an overly-optimistic forecast errors. In this case, a similar line of argument holds: the average MPC is 7% and, thus, more than twice as high as in the rational model but lower compared to our baseline model. If we assume that mild (strong) optimists are 15pp (25pp) more likely to do overly-optimistic forecast errors—our average estimates rather than the lower bounds—the model predicts an average MPC of 16%.

We conclude that alternative approaches to micro-found the optimistic bias all yield similar conclusions about how heterogeneity in optimism affects household behavior and average MPCs. This is reassuring given that in reality, all of these channels are likely to be present simultaneously.

Low (2024).

[32]The existence of the expense shocks decreases the average MPC in the rational model by 2pp.

4.4.3 Relationship to discount factor heterogeneity

As illustrated by Krueger et al. (2016), Aguiar et al. (forthcoming), or Kaplan and Violante (2022), ex-ante heterogeneity in discount factors β can help the rational model account for some of the MPC patterns observed in the data. Aguiar et al. (forthcoming) suggest that behavioral frictions could provide a microfoundation for the low β of some households. Yet our empirical evidence in Section 2 points towards optimism and not impatience as having a strong connection to HtM status and to observed differences in savings behavior.

Besides the empirical evidence, there are also important distinctions from a modeling perspective between heterogeneity in optimism and heterogeneity in discount factors. Note first that they are not equivalent, as the following Lemma states.

Lemma 1. *Unless marginal utility is constant across individual states, the model with heterogeneity in optimism and the model with heterogeneity in patience are not equivalent.*

Proof. See Appendix B. ■

The intuition is that optimism affects expected marginal utility, which depends on the individual state of a household. In contrast, impatient households have the same lower discount factor independent of their current state. Thus, at the household level, these two models cannot be the same.[33]

At the macro level, it is nevertheless technically possible to produce the same average MPC predicted by our baseline model in a model with discount factor heterogeneity. As we show in Appendix C, this model differs in at least four important dimensions from our baseline model: first, it requires using the discount factor of the mildly and strongly impatient households as free parameters to match their average MPC. Second, it produces a wealth distribution that is too polarized and suffers from a "missing middle" as already discussed in Kaplan and Violante (2022). Third, the differences at the micro-level—reflected in different savings policy functions— can become quite substantial. And fourth, as we discuss later in Section 5.2.2, the two models differ vastly in their normative implications, highlighting that it matters for the optimal debt level *why* households differ in their savings behavior and HtM status.

4.5 Heterogeneity in optimism in a Two-Asset Model

Rational HANK models often introduce a second, illiquid asset to match the average MPC while simultaneously matching total wealth in the economy (Kaplan et al. (2018), Kaplan and Violante (2022), Auclert et al. (forthcoming)). This approach seeks to capture illiquid assets

[33]The "endogeneous discount factor heterogeneity model" in Appendix C of Kaplan and Violante (2022) could be observationally equivalent to our model, if it has as many degrees of freedom as individual wealth $\times$ income states. In this case, the endogenous discount factor can be modeled to replicate the degree of undersaving of optimists in each income and wealth state.

that are good long-run savings vehicles but ill-suited for self-insurance purposes. But in order to match high average MPCs, two-asset HANK models typically require a liquidity premium—a return difference between liquid and illiquid assets—that is arguably substantially higher than in the data (Kaplan and Violante (2022)). We now show that the two-asset version of our model can fit the MPC and wealth data with a substantially lower liquidity premium than required by rational two-asset HANK models.

Model. Per standard practice, adding an illiquid asset requires enriching the model in two ways. First, households can now save in two assets: a liquid but low-return bond, and illiquid but high-return productive capital. Second, the production function includes capital.

The household's budget constraint now reads:

$$c_t + \frac{b_t}{1+r_t} + k_t = b_{t-1} + (1+r_t^k)k_{t-1} + (1-\tau_t)w_t\bar{e}_g e_t n_t, \tag{14}$$

where k denotes the illiquid asset of the household and r^k is its net return. Capital depreciates at rate δ and depreciated capital has to be replaced for maintenance. We follow Bayer et al. (forthcoming) and assume that households make their savings and portfolio choices between liquid bonds and illiquid capital in light of a capital market friction: participation in the capital market is random and i.i.d. in the sense that only a fraction λ of households can adjust their capital holdings in a given period. Households not participating in the capital market in a given period ($k_t = k_{t-1}$) still obtain the return on their illiquid asset holdings and can adjust their bond holdings. We further assume that holdings of both assets must be non-negative:

$$b_t, k_t \geq 0.$$

A representative firm operates a Cobb-Douglas production function using capital (K) and labor (N) as input factors:

$$Y_t = K_{t-1}^{\chi} N_t^{1-\chi}, \tag{15}$$

where χ denotes the capital share in production.

In addition to the equilibrium conditions in Section 3, now the capital market must clear:

$$\sum_{g,e} \mu_g p(e) \int k_t \Psi_{g,t}\left(k_{t-1}, e_t\right) = K_t. \tag{16}$$

Calibration. We maintain the same values for each of the parameters that also appear in our baseline model (except for the discount factor). Appendix Table A8 shows our calibration of the additional parameters and the discount factor. We set the capital share to $\chi = 0.318$ and the quarterly depreciation rate to $\delta = 0.0175$ as in Bayer et al. (forthcoming) and the

liquid asset-to-annual income ration to 0.2 as in Kaplan and Violante (2022). We then use the per-period capital market participation probability λ and the discount factor β to jointly target the average wealth-to-annual income ratio of 4.1 and the same average MPC as in our baseline model of 0.18.

Belief heterogeneity and the liquidity spread. Table 7 shows the stationary equilibrium prediction of our two-asset model and compares its empirical fit to rational two-asset models. The share of HtM households is now 0.40, as compared to 0.23 in our baseline model, because the illiquid asset's higher return induces some savers to substitute from the liquid asset. This is mostly driven by "wealthy HtM" households who would not be HtM in a one-asset model and now choose to save only in the illiquid asset (Kaplan et al. (2018)). Given that the average MPC is the same as in our baseline model, in the two-asset model, on average, constrained households are not as far off their Euler equation and thus spend less out of a $500 windfall. Also in the two-asset model, there is a stark difference between the behavior of rational and optimistic households. Optimistic households remain much more likely to be HtM and the more optimistic the more so (80% strong optimists vs. 55% mild optimists vs. 13% rational households are HtM) because they foresee little value in accumulating a liquid buffer stock (see Appendix Table A9).

Table 7: MPCs and liquidity spread across two-asset models.

	Two-asset HANK w/ belief Het.	Rational two-asset HANK	
	(1)	(2)	(3)
		Calibrated as (1)	Recalibrated
HtM	0.40	0.25	0.27
Avg. MPC	0.18	0.07	0.15
return gap (annualized)	2.6%	5.7%	9.3%

Note: MPCs refer to MPCs out of a stimulus check of $500. The model in Column 3 is recalibrated to produce an average MPC of 0.15.

Column 2 presents the rational two-asset HANK model for comparison, keeping all the parameters the same as in our model except for recalibrating β to target the mean wealth-to-annual income ratio of 4.1.[34] Unlike our model, this produces a HtM share of 0.25 and an average quarterly MPC of 0.07 which is substantially below the lower end of the consensus ranges of empirical estimates.[35]

[34]This requires quarterly $\beta = 0.990$, as compared to 0.993 in our model.

[35]Compared to our model, the average MPC in the rational model is also lower conditional on the share of HtM households, because optimistic households have higher MPCs conditional on their current state given their higher expected life income.

If we recalibrate the rational model to target an average MPC of 0.15, the model requires an annualized return gap of 9.3% (Column 3).[36] Our model produces a much lower return gap of 2.6% because overconfident households underestimate precautionary savings needs and thus require a much smaller premium on illiquid assets, thereby driving demand for the illiquid asset up and its return down.

Given empirical estimates of the return gap in the ballpark of 5% (see, e.g., Jordà et al. (2019)), it may seem at first glance that our two-asset model undershoots substantially. But both our model and the rational versions of it abstract from aggregate risk. Accounting for aggregate risk would likely push our estimated risk premium closer to the data and a standard HANK model's estimate even farther away from it, in the case where standard HANK targets an empirically realistic average MPC as in Column 3. Ilut et al. (2024) argue that the larger part of return gaps are indeed driven by aggregate risk instead of a liquidity premium.

5 Fiscal Policy Implications of Heterogeneity in Optimism

Heterogeneity in optimism affects households' precautionary savings behavior which helps the model to move closer to the data. We now show that this also matters for the design and effectiveness of fiscal policy tools that seek to stimulate and insure consumption. There are two key mechanisms. First, our model produces a flatter MPC-income distribution, as optimism is a key predictor of HtM status even conditional on income. This is in line with empirical evidence and has implications for the effectiveness of income-targeted transfer payments. Second, optimistic households are less responsive to changes in precautionary savings incentives. This dampens crowdout when the government provides insurance (we consider a minimum income benefit as an example), but makes it more difficult to induce households close to the borrowing constraint to self-insure when the government provides liquidity through higher public debt levels.

5.1 The distribution of MPCs and targeted transfers

We start by considering an income-targeted transfer that is not anticipated by consumers. Given the difficulty of directly identifying empirical evidence on the general-equilibrium effects of transfer policies, they are generally evaluated using models that match the observed average MPC (see e.g., Kaplan and Violante (2014), Wolf (forthcoming)). But for targeted transfers it is the MPC of transfer recipients that matters most.

Panel (a) of Figure 2 shows the average MPC along the income distribution in our baseline model (blue dots) and compares it to its rational counterpart (black dots). In addition, it compares it to the rational model recalibrated to match the same average MPC (orange dots). In our model, the MPCs are quite elevated along the whole income distribution. Most strikingly,

[36] In targeting the quarterly average MPC of 0.15 we set $\beta = 0.9805$, $\lambda = 0.15$, and $\delta = 0.00875$.

the MPCs are almost flat along the first three income quartiles. Only in the highest income quartile the MPC drops substantially (but still amounts to almost 10%). The strongest predictor of high MPCs in our model is the degree of optimism, as optimistic households can have large MPCs even if they are in the middle of the income distribution (or above) (see Figure 1 and our discussion in Section 4.1). Only at the very top of the income distribution the MPCs flatten out also for the strong optimists. Flat MPC-income distributions have been found empirically by for example Boehm et al. (2023) or Fuster et al. (2021).

This contrasts strongly with the rational models: first, in the rational counterpart of our model, MPCs are low along the whole income distribution as in this model almost all households are well insured and, thus, rarely HtM. Turning to the rational model with low wealth, while it shares the same average MPC with our model, it predicts a completely different MPC-income distribution. In particular, it predicts a strong MPC-income gradient: MPCs are high for the first quintile, but MPCs then sharply decrease in income. The reason is that the rational model can only produce a high enough average MPC by drastically reducing liquidity, which pushes up basically all low-income households to the borrowing constraint. A strong negative gradient in income is the result of that.

Our model's more realistic depiction of the MPC-income gradient implies that transfers targeted to low-income households are less effective at stimulating consumption than the rational HANK model with the same average MPC would imply, because income is a much weaker predictor of MPCs in our model. Consider a surprise lump-sum transfer to each household in the bottom income quartile, in an aggregate amount of 1 percent of steady-state output on impact, following an AR(1) process with a persistence parameter of 0.8, and financed in the short-run by higher debt which is then slowly repaid with higher taxes. Panel (b) of Figure 2 shows the output response, in terms of percentage deviations from steady-state output. Our model (dashed blue line) predicts an output response that is less than half as strong than in the rational HANK model with the same average MPC (dotted orange line): the rational HANK model implies a transfer multiplier of about 2.8 on impact whereas our model implies one of 0.8.[37]

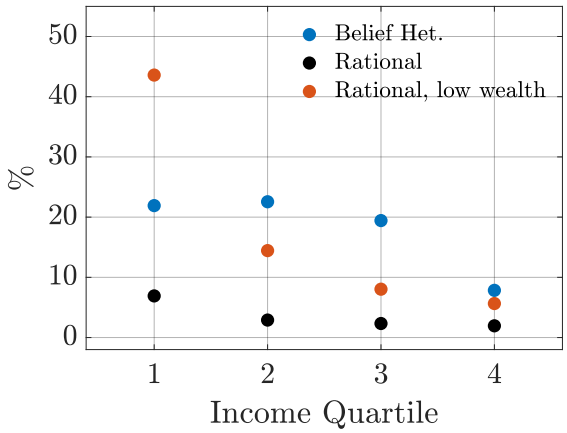
A second channel further weakens the effectiveness of targeted transfers in our model: muted relaxation of the precautionary saving motive. The persistence in the targeted transfers provides some insurance for households' income risk, temporarily decreasing households' precautionary savings motive. This motive is prevalent and strong in the rational model with classically rational households, who fully understand the relaxation of their precautionary savings motive, further increasing spending and total output.[38] But the more optimistic households are the

[37]The low-MPC standard HANK model (solid black line) produces a multiplier of about 0.5, due to its low MPC across all income groups.

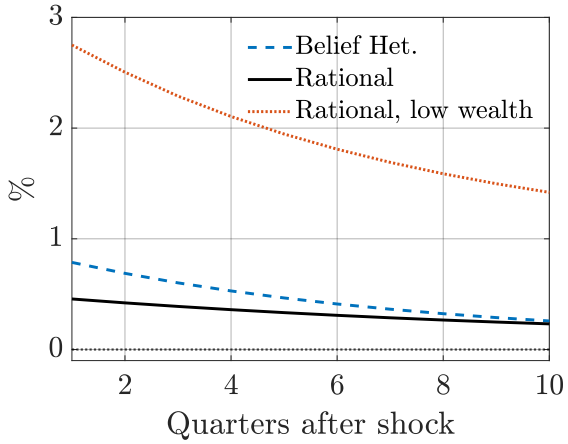
[38]See e.g. Bayer et al. (2023a)'s analysis of targeted transfers in a rational HANK model where the relaxation of households' precautionary savings is an important contributor to high multipliers. Kekre (2023), Dengler and

Figure 2: Income-MPC distribution and targeted transfer shocks

(a) MPCs along the income distribution



(b) Output response



Note: Panel (a) shows the average MPCs for the different income quartiles for three different models. "Belief Het." is our baseline model and "Rational" its rational counterpart. "Rational, low wealth" is the rational HANK model recalibrated to match the average MPC of our baseline model. The vertical axis depicts the MPCs in percent. Panel (b) shows the effects of a positive income-targeted transfer shock on total output expressed in percentage deviations from steady state output.

more they undervalue the insurance because they underestimate the likelihood of being income-eligible to exercise the insurance option in the future. They thus *perceive* their precautionary savings motive to be less relaxed than rational households, and as such do not increase their spending as much.[39]

5.2 Precautionary savings behavior and fiscal insurance policies

Accounting for the muted responsiveness of optimistic households to changes in precautionary savings incentives is is even more important when evaluating policies focused on insurance provision. We now consider two such policies: minimum income benefits as a form of public insurance, and government liquidity provision that reduces the cost of private insurance.

5.2.1 Minimum income benefits as public insurance

We start by analyzing the effects of introducing minimum income benefits (MIB) that provide public insurance against households' income risk. Following Bayer et al. (2023b), we model MIB as a transfer $tr_{i,t}$ to household i contingent on the household's pre-tax labor income $w_t n_{i,t} e_{i,t}$

Gehrke (2023), and Broer et al. (2024) find similar results for temporary increases in unemployment benefits and "short-term work", both of which can be understood as targeted transfers although they are not lump-sum and thus additionally have distortionary effects. Beraja and Zorzi (2024) analyze potential size-dependency for stimulus transfers.

[39]The relaxation of the precautionary savings motive is also an important driver in the rational HANK model with low average MPCs (black-solid line in Figure 2). But the MPCs are so low in that model, across all income quartiles, that it still predicts a smaller effect on aggregate output than our model.

falling short of some threshold level:

$$tr_{i,t} = \max\{0, a_1\bar{y} - a_2 w_t n_{i,t} e_{i,t}\},$$

where $\bar{y}$ is the median income in the stationary equilibrium and $0 \leq a_1, a_2 \leq 1$. Transfers thus decrease in individual income at the withdrawal rate a_2 and no transfers are paid to households whose labor income satisfies $w_t n_{i,t} e_{i,t} \geq \frac{a_1}{a_2}\bar{y}$. Following Bayer et al. (2023b), we set $a_1 = 0.5$ and $a_2 = 0.8$ and assume for simplicity that these transfers do not distort labor supply.

Total government transfer payments are then:

$$Tr_t = \mathbb{E}_t tr_{it},$$

where the expectation operator is the cross-sectional average. These transfers are financed via labor-income taxes.

Table 8: Effects of introducing public insurance

	HANK w/ Belief Het.		Rat HANK		RHANK, low w	
	No MIB	MIB	No MIB	MIB	No MIB	MIB
	(1a)	(1b)	(2a)	(2b)	(3a)	(3b)
HtM Share	0.23	0.27	0.02	0.09	0.21	0.28
Avg. MPC	0.18	0.17	0.04	0.06	0.18	0.28
Top10W	46%	47%	36%	38%	56%	60%
Bottom50W	2.1%	1.5%	12.7%	9.2%	2.3%	0.9%
Real rate	4%	4.9%	4%	5.5%	4%	6.9%

Note: MPCs refer to MPCs out of a $500 dollar stimulus check. "HANK w/ Belief Het." is our baseline model (one-asset, with heterogeneity in beliefs), "Rat HANK" denotes its rational counterpart, "RHANK, low w" is the same rational HANK model but with restricted wealth to match the average MPC of "HANK w Belief Het.". "MIB" refers to the stationary equilibrium in the models with public insurance via minimum income benefits (MIB).

Table 8 shows the stationary equilibrium effects of MIB on the average MPC, the HtM share as well as wealth shares in our baseline model (Columns 1a and 1b) and compares this to effects of MIBs in a rational one-asset HANK model (Columns 2a and 2b). We also consider a rational HANK model in which we reduce the amount of wealth such that it produces the same average MPC as our baseline model in the absence of MIB (Columns 3a and 3b).

In the two rational models, targeted transfers crowd-out self-insurance precautionary savings quite strongly. Households correctly forecast the probability of a bad productivity draw and thus internalize the insurance value of receiving a transfer in that state, reducing their precautionary savings accordingly. This increases the average MPC by more than 50% in either rational model, and the HtM share also increases substantially (by 7pp in both models). This crowd-out also causes a strong increase in the equilibrium real interest rate from 4% to 5.5% and even to

6.9% in the model with low wealth. This higher real rate is required to induce households to hold the same amount of wealth, when their precautionary savings risk has been dampened by the policy. Looking at the wealth shares at the top shows that the higher interest rate predominantly incentives the wealthy households to increase their wealth, as the share held by the top 10 increases by 2pp and 4pp, respectively, whereas the share held by the bottom 50% decreases by 3.5pp and 1.4pp, respectively.

In our model, these effects are dampened because overconfident households underpredict the probability of reaching a low-productivity state in which they receive a transfer. The average MPC even slightly decreases from 0.18 to 0.17,[40] while the share of HtM households only mildly increases by 4pp. The real interest rate increase is also substantially smaller, rising only to 4.9%. Also the effect on inequality is smaller as the wealth share of bottom 50% only decreases by 0.5pp and the wealth share of the top 10% only increases by 1pp. We conclude that direct insurance policies like minimum-income benefits have smaller side effects when some households are overly-optimistic.

5.2.2 Liquidity Provision and the Optimal Public Debt Level

Instead of directly insuring (part) of households' income risk, fiscal policy can also facilitate private insurance, by issuing more government debt (e.g., Woodford (1990)). More debt increases the supply of assets and thus of self-insurance possibilities for households. But this increase in liquidity supply has muted effects in our model compared to the rational HANK model. Figure 3(a) shows the share of HtM households, and 3(b) the share of wealth held by the poorest 50% of households, as a function of the government debt level in steady state.[41]

The solid black lines in Figure 3 show that in the rational HANK model, the provision of liquidity drives down the share of HtM, and increases the wealth share of the bottom 50% quite effectively. Households at or near the borrowing constraint have the strongest incentive to self-insure by saving in liquid assets and respond strongly as the price of the asset falls. This drives down their HtM likelihood such that for relatively high public debt levels, almost no households are borrowing constrained.

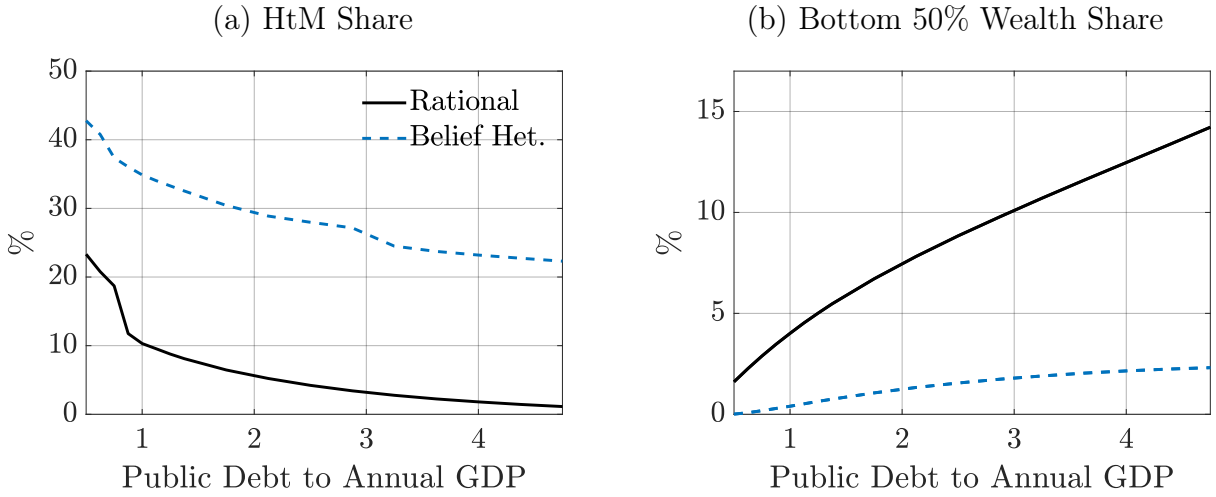
The dashed blue lines in Figure 3 illustrate the much weaker household response to liquidity provision in our model. The share of HtM households has a relatively flat slope with respect to debt supply, and it plateaus well above zero; e.g., it is about 0.23 at a debt-to-GDP ratio of 4,

[40] There are two opposing effects of the introduction of MIB on the average MPC: on the one hand, the effective lower income risk reduces households' MPC conditional on their individual state. On the other hand, there are more households in individual states with higher MPCs as MIB crowd out precautionary savings. In the rational models, the latter dominates whereas in our baseline model, the first effect dominates because minimum income benefits only mildly crowd out households' precautionary savings.

[41] When varying the supply of government debt, we fix the discount factor β as calibrated in Table 5 and let the interest rate adjust to clear the bond market.

compared to nearly zero in the standard model. The bottom 50% wealth slope is remarkably flat, reaching only about a 2% share at a debt-to-GDP ratio of 4 compared to about 13% in the standard model. Even when liquidity is abundant, optimistic households do not tend to save themselves out of being liquidity constrained because they still perceive the asset price as too high compared to their underestimated precautionary savings motive.

Figure 3: Implications of higher government debt



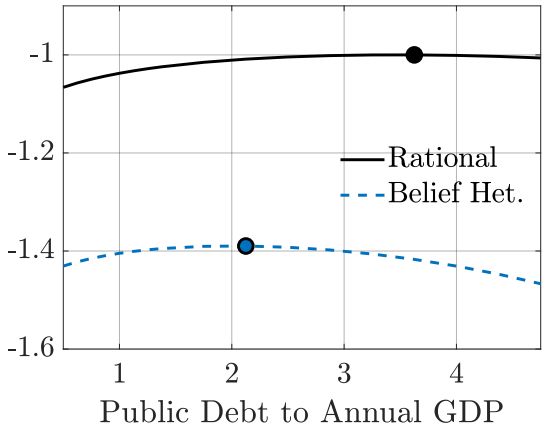
Note: This figure shows the share of HtM households in panel (a) and the wealth share of the bottom 50% of households in panel (b) for varying degrees of average government debt to average earnings ratios (horizontal axis). The black-solid lines show the case for the one-asset standard HANK model that abstracts from permanent heterogeneity in cognitive skills and overconfidence, and the blue-dashed lines show the case for our baseline HANK model featuring permanent heterogeneity in cognitive skills and overconfidence.

The relative unresponsiveness of households at or close to the borrowing constraint in our model also has implications for the optimal amount of government debt. A social planner weighs the benefits of smoother household consumption (from cheaper self-insurance) vs. the costs of the distortionary taxes required to finance the government's additional interest rate payments. We evaluate this trade-off in both models using a utilitarian social welfare function that seeks to maximize the average expected discounted lifetime utility of households.[42]

Figure 4 shows that average welfare peaks at a much lower debt level in our model compared to the standard one-asset HANK model: optimal debt is about 210% of annual GDP, compared to about 370% in the standard HANK model. Since optimistic households underestimate their insurance needs and therefore have a dampened response to the liquidity supply increase even when they are at or close to the borrowing constraint, the very households that the social planner would like to save more are the least responsive ones. This diminishes the social benefit of higher government debt compared to the standard model. Even though we abstract from many

[42]Such an objective function takes into account aggregate efficiency, risk-sharing, and intertemporal-sharing (Dávila and Schaab, 2023). The expectations over the individual lifetime utilities in the social welfare function are assumed to be rational, in the spirit of what Benigno and Paciello (2014) call "paternalistic".

Figure 4: Public Debt and Social Welfare



Note: This figure shows average welfare, defined as average expected discounted lifetime utility, as a function of government debt. Dots show the welfare-maximizing amount of government debt for our baseline model (blue-dashed lines) and its rational counterpart (black-solid line). The y-axis shows (normalized) average expected lifetime utility, and the x-axis shows (Public debt outstanding)/(Annual GDP), $\frac{B}{4Y}$. For readability, we normalize welfare such that the highest level of welfare in the model with rational expectations is normalized to -1.

important channels here—and therefore, our quantitative estimates should be interpreted with caution—the mechanism through which heterogeneity in optimism reduces the optimal debt level likely holds in richer models as well.[43]

Analyzing the optimal debt level also highlights the importance of accounting for *why* households differ in their savings behavior and HtM status. For example, our model and a model with heterogeneity in discount factors produce very different optimal debt levels, even when we consider the discount factor heterogeneity model that produces the same average MPC at our baseline wealth-to-income ratio of 4.1 (see Appendix C). In the model with discount factor heterogeneity, the optimal debt level is 2.5 times as high as in our baseline model and even 1.5 times higher than in the model without ex-ante heterogeneity because the households who benefit more from government liquidity provision (those with higher discount factors, because they value precautionary savings more) also get de facto higher social welfare weights in a utilitarian welfare function (because their future utility is discounted less). As such, accounting for the strong empirical relationships between optimism, savings behavior and HtM status in Table 4, rather than relying on heterogeneity in patience (and its weaker empirical links to HtM status), can matter greatly for optimal policy.

[43]In a robustness exercise, we analyze the optimal debt level in our two-asset model and its rational counterpart. Accounting for heterogeneity in optimism again reduces the optimal debt level significantly, although for both models the level of optimal debt is lower than in the respective one-asset models due to crowdout of productive capital. See e.g., Aiyagari and McGrattan (1998), Davila et al. (2012), Angeletos et al. (2023), or Woodford (1990) for analyses of optimal public liquidity provision in rational models and Auclert et al. (2024a) for characterizing long-run optimal fiscal policy including the transition in rational models.

6 Conclusion

In this paper, we provide evidence suggesting that heterogeneity in households' beliefs about their own future financial situation is an important driver of heterogeneity in savings behavior and in households' financial situation. Using a novel U.S. micro data set, we identify consumers with an optimistic bias about their own future financial situation and find that these households are less likely to save and more likely to be financially constraint.

Guided by these findings, we then introduce persistent belief heterogeneity into a HANK model and uncover a systematic reason why many households are persistently HtM: overly-optimistic households hold too little precautionary savings and, thus, they are not hand-to-mouth simply because of "bad luck" but because they persistently under-save. Accounting for this reason resolves heretofore seemingly intrinsic tensions in HANK models. Unlike other models, our one-asset HANK model can simultaneously match consensus estimates of both the average MPC and the average wealth level. Our two-asset HANK model matches the data with a lower, and perhaps more empirically realistic, liquidity premium than required in other models. Importantly, heterogeneity in optimism substantially improve the empirical fit of existing HANK models while at the same time it can be empirically disciplined by our micro data.

We also show that accounting for the underlying reason why some households are persistently financially constrained matters greatly for fiscal policies. This is particularly pronounced for policies that affect the precautionary savings incentives of households, because optimistic households undervalue insurance and thus have muted responses to changes in such incentives.

One consideration for future work on normative questions—we mostly consider positive ones in this paper—is whether optimism may not be all bad, from a welfare perspective. If it is not all bad, quantitative welfare modeling might seek to account for the benefits. Regardless, our finding that optimism correlates strongly with persistent and severe financial distress suggests important costs—costs that might be amplified by financial stress (Sergeyev et al. (2024)).

We also stop short of examining different combinations of macroeconomic policies in the presence of permanent heterogeneity across households—but our model provides a framework for doing so going forward. Consideration of monetary policy, and fuller consideration of fiscal policy, likely will require accounting for an additional source of heterogeneity: beliefs about macroeconomic variables.

References

Aguiar, M. A., M. Bils, and C. Boar (forthcoming): "Who are the Hand-to-Mouth?" *Review of Economic Studies*.

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

Aiyagari, S. R. and E. R. McGrattan (1998): "The optimum quantity of debt," *Journal of Monetary Economics*, 42, 447–469.

Andreoni, J. and C. Sprenger (2012): "Estimating time preferences from convex budgets," *American Economic Review*, 102, 3333–3356.

Angeletos, G.-M., F. Collard, and H. Dellas (2023): "Public debt as private liquidity: Optimal policy," *Journal of Political Economy*, 131, 3233–3264.

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

Auclert, A., B. Bardóczy, and M. Rognlie (2023): "MPCs, MPEs, and multipliers: A trilemma for New Keynesian models," *Review of Economics and Statistics*, 105, 700–712.

Auclert, A., M. Cai, M. Rognlie, and L. Straub (2024a): "Optimal long-run fiscal policy with heterogeneous agents," .

Auclert, A., M. Rognlie, and L. Straub (2020): "Micro jumps, macro humps: Monetary policy and business cycles in an estimated HANK model." *Working Paper*.

——— (2024b): "Fiscal and Monetary Policy with Heterogeneous Agents," *Working paper*.

——— (forthcoming): "The intertemporal Keynesian cross," *Journal of Political Economy*.

Barsky, R. B., F. T. Juster, M. S. Kimball, and M. D. Shapiro (1997): "Preference parameters and behavioral heterogeneity: An experimental approach in the health and retirement study," *The Quarterly Journal of Economics*, 112, 537–579.

Bayer, C., B. Born, and R. Luetticke (forthcoming): "Shocks, frictions, and inequality in US business cycles," *American Economic Review*.

Bayer, C., B. Born, R. Luetticke, and G. J. Müller (2023a): "The coronavirus stimulus package: How large is the transfer multiplier," *The Economic Journal*, 133, 1318–1347.

Bayer, C., A. Kriwoluzky, G. J. Müller, and F. Seyrich (2023b): "A HANK2 model of monetary unions," *Working Paper*.

Benigno, P. and L. Paciello (2014): "Monetary policy, doubts and asset prices," *Journal of Monetary Economics*, 64, 85–98.

Beraja, M. and N. Zorzi (2024): "Durables and Size-Dependence in the Marginal Propensity to Spend," *Working Paper*.

Berman, J. Z., A. T. Tran, J. G. Lynch Jr, and G. Zauberman (2016): "Expense neglect in forecasting personal finances," *Journal of Marketing Research*, 53, 535–550.

Bewley, T. (1986): "Stationary monetary equilibrium with a continuum of independently fluctuating consumers," in Hildenbrand, W. and Mas-Collel, A. (eds) *Contributions to mathematical economics in honor of Gérard Debreu (Amsterdam: North-Holland)*, 79.

Bhandari, A., J. Borovička, and P. Ho (2024): "Survey data and subjective beliefs in business cycle models," *Review of Economic Studies*.

Boehm, J., E. Fize, and X. Jaravel (2023): "Five facts about MPCs: Evidence from a randomized experiment," .

Bordalo, P., N. Gennaioli, Y. Ma, and A. Shleifer (2020): "Overreaction in macroeconomic expectations," *American Economic Review*, 110, 2748–82.

Boutros, M. (2023): "Windfall Income Shocks with Finite Planning Horizons," *Working Paper*.

Broer, T., J. Druedahl, K. Harmenberg, and E. Oberg (2024): "Stimulus effects of common fiscal policies," Tech. rep., Working Paper.

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

Caballero, R. J. and A. Simsek (2020): "Prudential monetary policy," *Working Paper*.

Carroll, C., J. Slacalek, K. Tokuoka, and M. N. White (2017): "The distribution of wealth and the marginal propensity to consume," *Quantitative Economics*, 8, 977–1020.

Carver, C., M. Scheier, and S. Segerstrom (2010): "Optimism." *Clinical Psychological Review*, 30, 879–889.

Chapman, J., M. Dean, P. Ortoleva, E. Snowberg, and C. Camerer (2023): "Econographics," *Journal of Political Economy Microeconomics*, 1, 115–161.

Claus, E. and V. H. Nguyen (2023): "Biased expectations," *European Economic Review*, 154, 104424.

Cocco, J. F., F. Gomes, and P. Lopes (2021): "Evidence on expectations of household finances," *Working Paper*.

Dávila, E. and A. Schaab (2023): "Welfare Assessments with Heterogeneous Individuals," *Working Paper*.

Davila, J., J. H. Hong, P. Krusell, and J.-V. Ríos-Rull (2012): "Constrained efficiency in the neoclassical growth model with uninsurable idiosyncratic shocks," *Econometrica*, 80, 2431–2467.

Dengler, T. and B. Gehrke (2023): "Short-time work and precautionary savings," *Working Paper*.

d'Haultfoeuille, X., C. Gaillac, and A. Maurel (2021): "Rationalizing rational expectations: Characterizations and tests," *Quantitative Economics*, 12, 817–842.

Dohmen, T., A. Falk, D. Huffman, and U. Sunde (2010): "Are risk aversion and impa-

tience related to cognitive ability?" *American Economic Review*, 100, 1238–1260.

D'Acunto, F. and M. Weber (2024): "Why survey-based subjective expectations are meaningful and important," *Annual Review of Economics*, 16.

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

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

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

Fulford, S. and D. Low (2024): "Expense shocks matter," *Consumer Financial Protection Bureau Office of Research Working Paper*.

Fuster, A., G. Kaplan, and B. Zafar (2021): "What would you do with 500?*Spendingresponsestogains*, *losses*, *news*, *andloans*," *The Review of Economic Studies*, 88, 1760−−1795.

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

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

Ganong, P., F. Greig, P. Noel, D. M. Sullivan, and J. Vavra (2024): "Spending and job-finding impacts of expanded unemployment benefits: Evidence from administrative micro data," *American Economic Review*, 114, 2898–2939.

Gillen, B., E. Snowberg, and L. Yariv (2019): "Experimenting with measurement error: Techniques with applications to the Caltech cohort study," *Journal of Political Economy*, 127, 1826–1863.

Guerreiro, J. (2023): "Belief Disagreement and Business Cycles," *Working Paper*.

Havranek, T. and A. Sokolova (2020): "Do consumers really follow a rule of thumb? Three thousand estimates from 144 studies say "probably not"," *Review of Economic Dynamics*, 35, 97–122.

Heck, P. R., D. J. Benjamin, D. J. Simons, and C. F. Chabris (2024): "Overconfidence persists despite years of accurate, precise, public, and continuous feedback: Two studies of tournament chess players," .

Hoffman, M. and S. V. Burks (2020): "Worker overconfidence: Field evidence and implications for employee turnover and firm profits," *Quantitative Economics*, 11, 315–348.

Huffman, D., C. Raymond, and J. Shvets (2022): "Persistent overconfidence and biased memory: Evidence from managers," *American Economic Review*, 112, 3141–75.

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

Hyytinen, A. and H. Putkuri (2018): "Household optimism and overborrowing," *Journal of Money, Credit and Banking*, 50, 55–76.

Ilut, C. and R. Valchev (2023): "Economic agents as imperfect problem solvers," *The Quarterly Journal of Economics*, 138, 313–362.

Ilut, C. L., R. Luetticke, and M. Schneider (2024): "HANK's Response to Aggregate Uncertainty in an Estimated Business Cycle Model," .

Jappelli, T. and L. Pistaferri (2010): "The consumption response to income changes," *Annual Review of Economics*, 2, 479–506.

Johnson, D. and J. Fowler (2011): "The evolution of overconfidence," *Nature*, 477, 317–320.

Jordà, Ò., K. Knoll, D. Kuvshinov, M. Schularick, and A. M. Taylor (2019): "The rate of return on everything, 1870–2015," *The Quarterly Journal of Economics*, 134, 1225–1298.

Kamdar, R. and W. Ray (2024): "Attention-Driven Sentiment and the Business Cycle," *Working Paper*.

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

Kaplan, G. and G. L. Violante (2014): "A model of the consumption response to fiscal stimulus payments," *Econometrica*, 82, 1199–1239.

——— (2022): "The marginal propensity to consume in heterogeneous agent models," *Annual Review of Economics*, 14, 747–775.

Kaur, S., S. Mullainathan, S. Oh, and F. Schilbach (2025): "Do financial concerns make workers less productive?" *The Quarterly Journal of Economics*, 140, 635–689.

Kekre, R. (2023): "Unemployment insurance in macroeconomic stabilization," *Review of Economic Studies*, 90, 2439–2480.

Koşar, G. and D. Melcangi (2024): "Subjective Uncertainty and the Marginal Propensity to Consume," *Working Paper*.

Krueger, D., K. Mitman, and F. Perri (2016): "Macroeconomics and household heterogeneity," in John Taylor and Harald Uhlig (eds) *Handbook of Macroeconomics*, Elsevier, vol. 2, 843–921.

Laibson, D., P. Maxted, and B. Moll (forthcoming): "Present bias amplifies the household balance-sheet channels of macroeconomic policy," *Quarterly Journal of Economics*.

Lewis, D., D. Melcangi, and L. Pilossoph (2024): "Latent heterogeneity in the marginal propensity to consume," Tech. rep., National Bureau of Economic Research.

Lian, C. (2023): "Mistakes in future consumption, high MPCs now," *American Economic Review: Insights*, 5, 563–581.

McArdle, J. J., G. G. Fisher, and K. M. Kadlec (2007): "Latent variable analyses of age trends of cognition in the Health and Retirement Study, 1992-2004." *Psychology and Aging*, 22, 525.

McClung, N. and T. Nighswander (2021): "Belief Heterogeneity and the Income-Wealth Relationship," *Working Paper*.

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

McKay, A. and C. K. Wolf (2023): "Monetary policy and inequality," *Journal of Economic Perspectives*, 37, 121–144.

Pappa, E., M. O. Ravn, and V. Sterk (2023): "Expectations and incomplete markets," in *Handbook of Economic Expectations*, Elsevier, 569–611.

Pfäuti, O. and F. Seyrich (2023): "A behavioral heterogeneous agent new keynesian model," *Working Paper*.

Rouwenhorst, K. G. (1995): "Asset Pricing Implications of Equilibrium Business Cycle Models," in Thomas F. Cooley (ed) *Frontiers of Business Cycle Research*, Princeton University Press, 294–330.

Rozsypal, F. and K. Schlafmann (2023): "Overpersistence bias in individual income expectations and its aggregate implications," *American Economic Journal: Macroeconomics*, 15, 331–371.

Seidl, H. and F. Seyrich (2023): "Unconventional Fiscal Policy in a Heterogeneous-Agent New Keynesian Model," *Journal of Political Economy Macroeconomics*, 1, 633–664.

Sergeyev, D., C. Lian, and Y. Gorodnichenko (2024): "The Economics of Financial Stress," *Working Paper*.

Solon, G., S. J. Haider, and J. M. Wooldridge (2015): "What are we weighting for?" *Journal of Human Resources*, 50, 301–316.

Souleles, N. S. (2004): "Expectations, heterogeneous forecast errors, and consumption: Micro evidence from the Michigan consumer sentiment surveys," *Journal of Money, Credit and Banking*, 39–72, Vol. 36(1).

Spinnewijn, J. (2015): "Unemployed but optimistic: Optimal insurance design with biased beliefs," *Journal of the European Economic Association*, 13, 130–167.

Stango, V. and J. Zinman (2020): "Behavioral biases are temporally stable," *National Bureau of Economic Research Workin Paper*.

——— (2023): "We Are All Behavioural, More, or Less: A Taxonomy of Consumer Decision-Making," *The Review of Economic Studies*, 90, 1470–1498.

——— (2024): "Behavioral biases are temporally stable," *Working Paper*.

Weidmann, B., J. Vecci, F. Said, D. J. Deming, and S. R. Bhalotra (2024): "How Do You Find a Good Manager?" Tech. rep., National Bureau of Economic Research.

Wolf, C. (forthcoming): "Interest Rate Cuts vs. Stimulus Payments: An Equivalence Result," *Journal of Political Economy*.

Woodford, M. (1990): "Public debt as private liquidity," *The American Economic Review*, 80,

382–388.

——— (2013): "Macroeconomic analysis without the rational expectations hypothesis," *Annual Review of Economics*, 5, 303–346.

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

Online Appendix

Beyond bad luck... Macroeconomic implications of persistent heterogeneity in optimism

Oliver Pfäuti Fabian Seyrich Jonatahn Zinman

Summary of the Online Appendix

Appendix A contains supplementary tables.

Appendix B contains proofs.

Appendix C contains results of a quantitative model of discount factor heterogeneity.

Appendix D contains additional model results and robustness analysis.

A Additional Empirical Results

Table A1: Measuring ex-post forecast errors (FCEs) with coarse data: Sample splits

Panel A1. Sample with measureable potentially optimistic FCEs, v1

	Realization at $t+1$			Panelist-level FCE counts					
FC at t	Better	Same	Worse						N with
Better	Accurate	Optimist	Optimist	N FCEs	Mean	SD	Min	Max	≥ 2 FCEs
Same	Pessimist	Accurate	Optimist	18,899	5.5	3.4	0	12	2,928
Worse	Pessimist	Pessimist	Accurate						

Panel A2. Sample with measureable potentially optimistic FCEs, v2: Exclude "Better-Better" pairs

	Realization at $t+1$			Panelist-level FCE counts					
FC at t	Better	Same	Worse						N with
Better	Accurate?	Optimist	Optimist	N FCEs	Mean	SD	Min	Max	≥ 2 FCEs
Same	Pessimist	Accurate	Optimist	16,841	4.9	3.3	0	12	2,792
Worse	Pessimist	Pessimist	Accurate						

Panel B1. Sample with measureable potentially pessimistic FCEs, v1

	Realization at $t+1$			Panelist-level FCE counts					
FC at t	Better	Same	Worse						N with
Better	Accurate	Optimist	Optimist	N FCEs	Mean	SD	Min	Max	≥ 2 FCEs
Same	Pessimist	Accurate	Optimist	15,836	4.6	3.6	0	12	2,542
Worse	Pessimist	Pessimist	Accurate						

Panel B2. Sample with measureable potentially pessimistic FCEs, v2: Exclude "Worse-Worse" pairs

	Realization at $t+1$			Panelist-level FCE counts					
FC at t	Better	Same	Worse						N with
Better	Accurate	Optimist	Optimist	N FCEs	Mean	SD	Min	Max	≥ 2 FCEs
Same	Pessimist	Accurate	Optimist	14,406	4.2	3.4	0	12	2,454
Worse	Pessimist	Pessimist	Accurate?						

Panel C. Sample with measureable potentially symmetric FCEs

	Realization at $t+1$			Panelist-level FCE counts					
FC at t	Better	Same	Worse						N with
Better	Accurate	Optimist	Optimist	N FCEs	Mean	SD	Min	Max	≥ 2 FCEs
Same	Pessimist	Accurate	Optimist	17,076	4.9	3.4	0	12	2,787
Worse	Pessimist	Pessimist	Accurate						

Note: Panel A1: Someone forecasting "worse" cannot be measured as having an optimistic forecast error based on their subsequent observed realization, because "worse" is the lowest realization one can observe in data. Panel A2: See Panel A1, with the difference here that we drop Better-Better pairs because, if there is an optimistic bias, this pair will misclassify many optimistic FCEs as "Accurate". Panel B1: Someone forecasting "better" cannot be measured as having an pessimistic forecast error based on their subsequent observed realization, because "better" is the highest realization one can observe in data. Panel C: To observe a forecast error that could potentially be in either direction, i.e., either optimistic or pessimistic, at least one of the inputs to the forecast-realization pair must take the middle value of "Same".

Table A2: Financial situation forecast errors are persistent, Part 1: Comparing consecutive forecast errors

Panel A. Unweighted estimates

	FCE this survey			
FCE previous survey	Optimist	Accurate	Pessimist	Total
Optimist	0.13	0.11	0.04	0.28
Accurate	0.09	0.43	0.06	0.58
Pessimist	0.04	0.06	0.04	0.14
Total	0.26	0.60	0.14	1

Panel B. Weighted estimates

	FCE this survey			
FCE previous survey	Optimist	Accurate	Pessimist	Total
Optimist	0.13	0.10	0.04	0.27
Accurate	0.09	0.45	0.06	0.60
Pessimist	0.04	0.06	0.04	0.14
Total	0.26	0.61	0.14	1

Note: FCE = panelist-level ex-post forecast error observed in our data, which spans 2009–2016. Each cell reports an estimate of a sample share. Sample is 13,810 potentially symmetric forecast-realization pairs from 2,787 panelists with ≥ 2 such pairs. If a column or row does not sum to exactly 1, that is due to rounding. Weighted estimates use sampling probability weight from the realization survey.

Table A3: Financial situation forecasts are positively correlated with income forecasts

	Forecasted probability of increase in:			
	Nominal income		Real income	
	Unweighted	Weighted	Unweighted	Weighted
	(1)	(2)	(3)	(4)
1= Optimistic forecast of sfc	0.00487	0.00484	0.00576	0.00546
s.e.	(0.00015)	(0.00020)	(0.00018)	(0.00024)
N	15,047	15,047	15,049	15,049
N panelists	3057	3057	3056	3056

Notes: Each column presents results from a single OLS regression of the row variable on the column variable and a constant. Standard errors, clustered on panelist, in parentheses. Weighted estimates use the ALP sampling probability weight for each observation. Income forecasts in percentage point units, so e.g., a point estimate of 0.005 indicates a 1/2 percentage point increase in sfc optimism per 1 pp increase in the probability of an income increase. SFC forecast optimism is indicated by responding to the question "Now looking ahead - do you think that a year from now you will be better off financially, or worse off, or just about the same as now?" with "Will be better off".

Table A4: Household financial condition forecasts and forecast errors tilt optimistic

Panel A. All forecasts, unweighted	Realization this year			
Forecast last year	Better	Same	Worse	Total
Better	0.10	0.13	0.04	0.27
Same	0.06	0.45	0.10	0.61
Worse	0.01	0.05	0.07	0.12
Total	0.16	0.63	0.21	1
Panel B. July 2009 & 2010, unweighted	Realization this year			
Forecast last year	Better	Same	Worse	Total
Better	0.06	0.16	0.05	0.28
Same	0.05	0.40	0.15	0.60
Worse	0.01	0.05	0.07	0.12
Total	0.12	0.61	0.27	1
Panel C. July 2009 & 2010, weighted	Realization this year			
Forecast last year	Better	Same	Worse	Total
Better	0.07	0.18	0.05	0.30
Same	0.04	0.38	0.14	0.56
Worse	0.01	0.07	0.06	0.14
Total	0.12	0.63	0.25	1
Panel D. January 2015 & 2016, unweighted	Realization this year			
Forecast last year	Better	Same	Worse	Total
Better	0.10	0.14	0.04	0.28
Same	0.06	0.47	0.08	0.61
Worse	0.01	0.05	0.06	0.12
Total	0.17	0.66	0.18	1
Panel E. January 2015 & 2016, weighted	Realization this year			
Forecast last year	Better	Same	Worse	Total
Better	0.11	0.13	0.03	0.27
Same	0.05	0.50	0.08	0.63
Worse	0.01	0.04	0.05	0.10
Total	0.17	0.67	0.16	1

Note: Cells report sample proportions. Forecasts: "Now looking ahead - do you think that a year from now you will be better off financially, or worse off, or just about the same as now?" Response options: Will be better off/About the same/Will be worse off. Realizations: "We are interested in how people are getting along financially these days. Would you say that you are better off or worse off financially than you were a year ago?" Response options: Better off/About the same/Worse off. Weighted estimates use sampling probabilities from the realization survey, which are correlated 0.90 and 0.93 with the weight from the paired forecast survey. Sample size is 21,586 in Panel A, 1,679 in Panels B and C, and 1,882 in Panels D and E.

Table A5: Household financial condition FCE persistence, Part 3: Panelist-level measures

Sample	Potentially opt or pess FCEs: All		Potentially opt or pess FCEs: v2		Potentially symmetric FCEs		No adjustment for measurement error	
Estimate	Unwtd	Wtd	Unwtd	Wtd	Unwtd	Wtd	Unwtd	Wtd
Panel A. Prop of forecast errors that are optimistic								
Proportion	0.33	0.32	0.39	0.38	0.33	0.32	0.28	0.28
SD	0.30	0.28	0.33	0.34	0.31	0.31	0.26	0.27
Share consumers with optimistic proportion ≥ 0.5	0.32	0.28	0.40	0.39	0.32	0.32	0.24	0.25
Share consumers with optimistic proportion >0.5	0.22	0.20	0.29	0.28	0.21	0.20	0.16	0.16
N consumers (≥ 2)	2928	2928	2792	2792	2787	2787	3073	3073
Panel B. Prop of forecast errors that are pessimistic								
Proportion	0.18	0.17	0.20	0.19	0.15	0.15	0.12	0.11
SD	0.23	0.23	0.25	0.25	0.21	0.21	0.17	0.17
Share consumers with pessimistic proportion ≥ 0.5	0.14	0.14	0.17	0.17	0.11	0.11	0.05	0.05
Share consumers with pessimistic proportion >0.5	0.07	0.06	0.09	0.08	0.05	0.04	0.02	0.02
N consumers (≥ 2)	2542	2542	2454	2454	2787	2787	3073	3073
Panel C. Relative prop or shares: Optimism/Pessimism								
Proportions	1.83	1.88	1.95	2.00	2.20	2.13	2.33	2.55
Shares ≥ 0.5	2.29	2.00	2.35	2.29	2.91	2.91	4.80	5.00
Shares >0.5	3.14	3.33	3.22	3.50	4.20	5.00	8.00	8.00

Note: N.b. Denominators in Panels A and B include accurate forecasts: see Table 1 for details on sample splits and number of forecast errors per panelist in each sample. Relative proportions simply use the sample estimates in Panel A and B to estimate: (persistent optimism)/(persistent pessimism). Weighted estimates use the mean sampling probability weight across surveys where the panelist provides a financial situation realization.

Table A6: Overconfidence strongly conditionally correlates with ex-post optimism

Panel A. LHS = Optimism FCE proportion						
	Potentially Optimistic v1		Potentially Optimistic v2		Potentially Symmetric	
	Unwtd	Wtd	Unwtd	Wtd	Unwtd	Wtd
	(1)	(2)	(3)	(4)	(5)	(6)
Overconfidence	0.25	0.25	0.30	0.72	0.24	0.64
	(0.11)	(0.21)	(0.13)	(0.37)	(0.12)	(0.36)
Patience	-0.26	-0.41	-0.34	-1.32	-0.32	-1.21
	(0.16)	(0.33)	(0.19)	(0.97)	(0.18)	(0.94)
Financial Risk Aversion	0.03	0.14	0.01	0.02	-0.00	0.00
	(0.09)	(0.15)	(0.10)	(0.20)	(0.10)	(0.18)
Income Gamble Risk Aversion	-0.06	-0.72	-0.15	-0.48	-0.17	-0.57
	(0.17)	(0.30)	(0.19)	(0.52)	(0.18)	(0.55)
Other Controls?	Yes	Yes	Yes	Yes	Yes	Yes
N	778	778	750	750	742	742
N Panelists	389	389	375	375	371	371
Mean (LHS)	0.3644		0.4387		0.3811	
Panel B. LHS = 1=(optimistic proportion $\geq$ 0.5)						
	Potentially Optimistic v1		Potentially Optimistic v2		Potentially Symmetric	
	Unwtd	Wtd	Unwtd	Wtd	Unwtd	Wtd
	(1)	(2)	(3)	(4)	(5)	(6)
Overconfidence	0.35	0.27	0.41	0.69	0.32	0.57
	(0.12)	(0.24)	(0.14)	(0.38)	(0.13)	(0.37)
Patience	-0.29	-0.50	-0.39	-1.43	-0.37	-1.31
	(0.17)	(0.35)	(0.20)	(0.99)	(0.19)	(0.97)
Financial Risk Aversion	0.05	0.17	0.03	0.04	0.01	0.01
	(0.09)	(0.16)	(0.11)	(0.21)	(0.11)	(0.19)
Income Gamble Risk Aversion	-0.08	-0.75	-0.18	-0.52	-0.20	-0.60
	(0.18)	(0.32)	(0.20)	(0.54)	(0.19)	(0.56)
Other Controls?	Yes	Yes	Yes	Yes	Yes	Yes
N	780	780	752	752	744	744
N Panelists	390	390	376	376	372	372
Mean (LHS)	0.37		0.45		0.39	

Note: Standard errors, clustered on panelist, in parentheses. One regression per panel-column. LHS variable is described in the panel and column headings, and the last row. RHS variables measured using the Stango-Zinman data and merged onto the our household financial situation forecast error panel. The results shown represent the [100] x the change in the LHS variable associated with a 1 percentage point increase in the RHS variable (each RHS variable here is a percentile rank). "Other control variables" include standard demographics (income, education age, gender, race, and ethnicity) and cognitive skills. For RHS variables likely subject to measurement error and where we are interested in their permanent component– overconfidence, patience, risk aversion, and cognitive skills– we use obviously related instrumental variables (ORIV) to account for measurement error .

Table A7: Household financial condition forecast errors are persistent

	Forecast error this survey			
FCE previous survey	Optimist	Realist	Pessimist	Total
Optimist	0.10	0.09	0.01	0.19
Realist	0.08	0.61	0.04	0.73
Pessimist	0.01	0.04	0.03	0.08
Total	0.18	0.74	0.07	1

Note: Sample is 10,546 forecast error pairs from 2,469 panelists. Sample is smaller here than in Appendix Table A4 because here we require ≥ 2 forecast-realization pairs per panelist and only include realizations of "about the same", to allow for the sharpest feasible test of persistence, by holding realizations constant and allowing for forecast errors in either direction (thereby minimizing measurement error from censoring).

B Proofs

Proof of Lemma 1. Lemma 1 says that unless marginal utility is constant across income states, heterogeneity in overconfidence and heterogeneity in patience are not equivalent. To see this, consider a simple counterexample. Focus on two households, $i \in \{1, 2\}$, and two possible future states, which we denote by U and D (e.g., for Up and Down). We focus on the equivalence of overconfident households and relatively impatient households with a discount factor $\hat{\beta} < \beta$. If overconfidence and patience heterogeneity are equivalent, it has to hold that the Euler equations of unconstrained households have to be identical. Imposing that household 1 has the same marginal utility in both economies in the current period implies that the expected discounted future marginal utility has to be identical, too:

$$\beta \tilde{E}_t \left[u'(c^1_{t+1}) \right] = \hat{\beta} E_t \left[u'(\hat{c}^1_{t+1}) \right], \tag{17}$$

where "$\hat{\cdot}$" denotes the economy with heterogeneity in patience. Similarly, for household 2. Assuming, without loss of generality, that household 1 starts in the U state and denoting the probability of moving to the D state by p_{UD}, equation (17) implies

$$\frac{\beta}{\hat{\beta}} = \frac{p_{UD} u'(c^{1,D}_{t+1}) + (1 - p_{UD}) u'(c^{1,U}_{t+1})}{\frac{1}{\alpha} p_{UD} u'(c^{1,D}_{t+1}) + (1 - \frac{1}{\alpha} p_{UD}) u'(c^{1,U}_{t+1})}. \tag{18}$$

(Implicitly, but without loss of generality, we assume here that consumption in the U state is higher than in the D state). Similarly, for household 2, who starts in state D

$$\frac{\beta}{\hat{\beta}} = \frac{p_{DU} u'(c^{2,U}_{t+1}) + (1 - p_{DU}) u'(c^{2,D}_{t+1})}{\alpha p_{DU} u'(c^{2,U}_{t+1}) + (1 - \alpha p_{DU}) u'(c^{2,D}_{t+1})}. \tag{19}$$

Thus, for given transition probabilities, degree of overconfidence α, discount factor in the economy with overconfidence β, and marginal utilities across states, we have one free parameter, $\hat{\beta}$, but two equations that need to hold.[44] Thus, the two economies are in general not identical (it becomes even less likely that the two economies are identical when we allow for more states and households). The only case in which the two are identical is when marginal utility is constant across states, that is when households can perfectly insure themselves against income shocks. Given our incomplete-markets setup, however, that is generally not the case, and therefore, heterogeneity in overconfidence is not equivalent to heterogeneity in patience. ■

[44] A simple numerical example illustrates this. Assume $p_{UD} = p_{DU} = 0.5$, $\alpha = 2$, $u'(c^D) = 1$ and $u'(c^U) = 2 > 1$. It follows that equation (17) implies a discount factor ratio of 0.86 whereas equation (17) implies a ratio of 0.75.

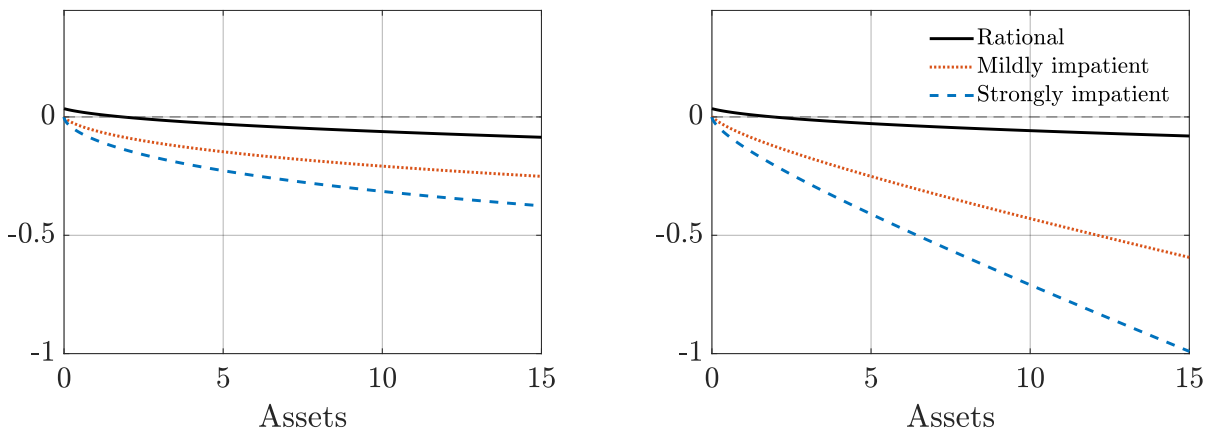
C Quantitative Discount Factor Heterogeneity Model

Lemma 1 states that the discount factor heterogeneity model is not equivalent to our baseline model. Nevertheless, a model with three groups of permanent heterogeneity in discount factors can in principle match average MPCs of the three different groups, if one uses the discount factors of the three groups as a free parameter to target these average MPCs. This requires setting $\beta_1 = 0.9845$ (the discount factor, that all households have in our baseline model), $\beta_2 = 0.9365$, and $\beta_3 = 0.8945$. Given that we target the average MPCs within groups also the aggregate MPC is the same as in our baseline model.

How different are the two models then when it comes to untargeted moments? The discount factor heterogeneity model predicts a HtM share of 20% which is 3pp lower than in our baseline model. Regarding the wealth distribution, this model predicts a wealth shares of 47% by the top 10 and 1% by the bottom 50% both of which are close to our baseline model (47% and 2%, respectively). However, it produces and median wealth-to-average-income ratio of only 0.7 which is much further away from its empirical counterpart of 1.5 than our baseline model predicts (1.2). This highlights the "missing middle problem" discussed in Kaplan and Violante (2022).

When it comes to the non-equivalence on the micro-level, the difference are even more stark. Figure C1 plots the savings policy functions of the respective three groups in our baseline model (left panel) and in the model with discount factor heterogeneity (right panel) for the median income state. It shows that the behavior of the households differ quite drastically along the asset distribution. This reflects Lemma 1: in the discount factor model, the degree of impatience is state-independent, whereas in our model, the degree of optimism is state-dependent.

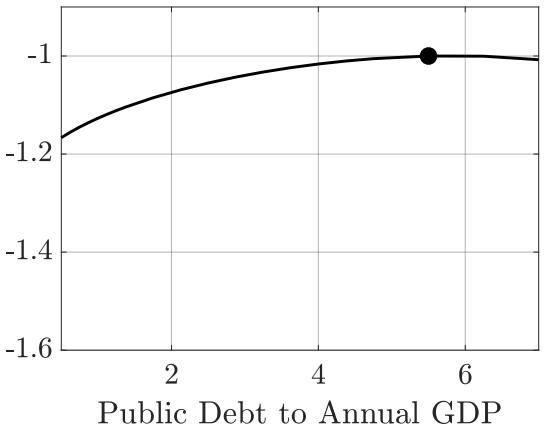
Figure C1: Savings policy in our baseline model and the discount factor heterogeneity model



Note: This figure shows the savings policy functions along the wealth distribution for the median income state in our baseline model with belief heterogeneity "HANK w Belief Het." and in the discount factor heterogeneity model "HANK w β Het." which is calibrated to match the same average MPCs in the three different groups.

The model with discount factor heterogeneity also differs drastically when it comes to its normative implications. Figure C2 shows the average welfare as function of public debt. It shows that the welfare maximizing debt in this model is around 550% of annual GDP which is not only more than double as high as in our baseline model but even much higher than in the rational HANK model without ex-ante heterogeneity (see Figure 4 and Section 5.2.2). This implies that for the normative implications, the reason *why* households differ in their savings behavior matters strongly.

Figure C2: Public debt and social welfare in HANK w β Het.

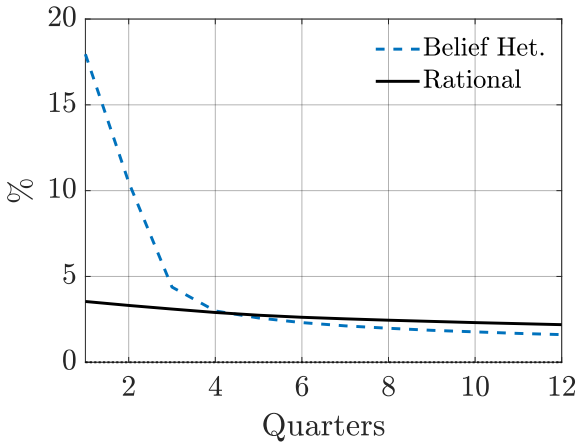


Note: This figure shows average welfare in the model with discount factor heterogeneity, defined as average expected discounted lifetime utility, as a function of government debt. Dots show the welfare-maximizing amount of government debt . The y-axis shows (normalized) average expected lifetime utility, and the x-axis shows (Public debt outstanding)/(Annual GDP), $\frac{B}{4Y}$. For readability, we normalize welfare such that the highest level of welfare is normalized to -1.

D Robustness and Additional Model Results

D.1 iMPCs

Figure D3: Intertemporal MPCs



Note: This figure shows the intertemporal MPCs (iMPCs) in our baseline model ("Belief Het.") as well as in its rational counterpart ("Rational"), that is, the dynamic spending out of an unexpected $500 income shock.

D.2 Two-Asset Model: Additional Tables

Table A8: Calibration two-asset model

Parameter	Description	Value
χ	Capital share	0.318
δ	Depreciation rate	0.0175
λ	Capital market participation rate	0.335
β	Discount factor	0.993

Note: The table shows the values for the additional parameters in our two-asset model and the discount factor. All other parameters stay the same as in our baseline model.

Table A9: MPCs and liquidity spread across two-asset models.

	Two-asset HANK w/ belief Het.	Rational two-asset HANK	
	(1)	(2)	(3)
		Calibrated as (1)	Recalibrated
HtM	0.40	0.25	0.27
Avg. MPC	0.18	0.07	0.15
return gap (annualized)	2.6%	5.7%	9.3%
HtM rat. HHs	0.13	0.25	0.27
Avg. MPC rat. HHs	0.04	0.07	0.15
HtM mild OC HHs	0.55	-	-
Avg. MPC mild OC HHs	0.20	-	-
HtM strong OC HHs	0.80	-	-
Avg. MPC strong OC HHs	0.43	-	-

Note: MPCs refer to MPCs out of a stimulus check of \$500. The model in Column 3 is recalibrated to produce an average MPC of 0.15.

D.3 Non-permanent belief biases

In our baseline model, we assume that the behavioral biases of households are permanent. While this was motivated by the highly persistent estimates of overconfidence in our sample as well as by the fact that the psychology literature traits overconfidence as permanent character trait, we now relax this assumption and assume that households can move between behavioral states. In particular, we assume the following transition probabilities between the three belief states [rational,mild optimists,strong optimists]:

$$\begin{bmatrix} \rho_B & (1-\rho_B)/2 & (1-\rho_B)/2 \\ (1-\rho_B) & \rho_B & 0 \\ (1-\rho_B)/2 & 0 & \rho_B. \end{bmatrix} \tag{20}$$

In addition, we assume that households do not take this switching probabilities into consideration. We set $\rho_B = 0.97$ as the empirical estimates in Stango and Zinman (2020). In this case, the average MPC of households are 0.10, and thus 3 times as high as in the rational model. If we set the optimistic biases of mild and strong optimists to the middle of the range of our empirical estimates (instead to the lower bounds as in our baseline), our model predicts an average MPC of 17%. If instead setting $\rho_B = 0.95$ and, thus, below the empirical estimates in Stango and Zinman (2020), the average MPC is 15%. We, thus, conclude that our results are robust to modelling the behavioral biases as not perfectly permanent but time-varying.

D.3.1 Alternative way of modeling optimism

In our baseline specification of optimism (equation (4)), the degree of optimism is the same for all households within the group of optimism, independent of their current productivity state. We now allow for dependence of the following form:

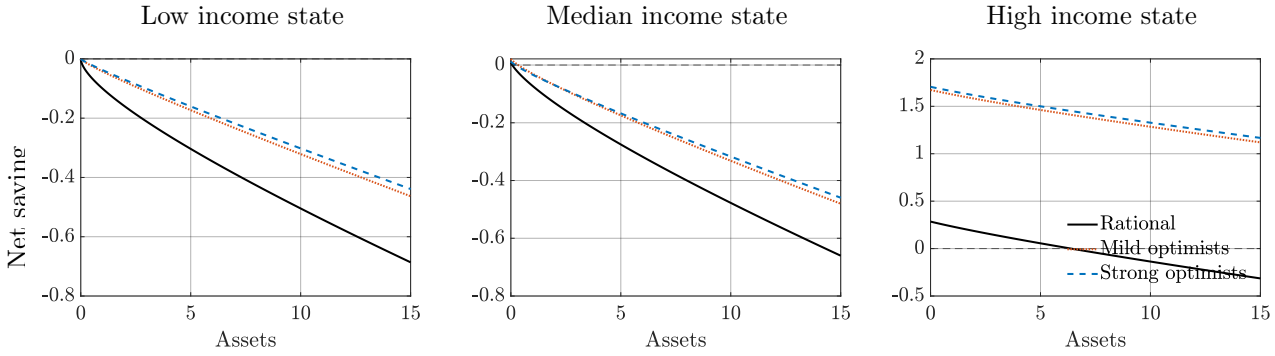
$$\tilde{p}_{ij} \equiv \begin{cases} \alpha_g^{(e_j - e_i)} p_{ij}, \text{ if } i \neq j \\ 1 - \sum_{j \neq i} \tilde{p}_{ij}, \text{ if } i = j. \end{cases} \tag{21}$$

As in our baseline specification, when $\alpha_g > 1$, the transition probabilities of moving upwards ($e_i < e_j$) are overweighted and the probabilities of moving downward are underweighted. Here we posit that these probability distortions are larger for states that are further away from each other. This specification may arise if households' beliefs are more distorted for less frequent events, such as large changes in their idiosyncratic productivity, than for more frequent events.

We again calibrate $\alpha_m = 1.5$ and $\alpha_s = 3.5$ matching the slope in the increase in the likelihood to make overly-optimistic forecast errors about their future financial situation. The predicted average MPC is 0.16 and thus largely unchanged from our baseline estimate of 0.18. The predicted HtM share is now about 3 percentage points higher, at 26%.

D.4 Households' behavior with higher risk

Figure D4: Households' savings policy function with higher risk



Note: The panels show the savings policy functions for the three different groups—rational households, mild optimists, and strong optimists—along the asset distribution exemplary for three different income (productivity) states. In this figure, mild and strong optimists face 9 and 12 times higher risk, but are rational about this. Given that we have 11 productivity levels, we choose the third for low, the sixth is the median and the ninth for high income.