Beyond bad luck: Macroeconomic implications of persistent heterogeneity in optimism

Oliver Pfäuti*

Fabian Seyrich[†]

Jonathan Zinman[‡]

May 20, 2025 Link to most recent version

Abstract

Household savings behavior and financial situations have important implications for macroeconomic fluctuations and policy. They are empirically persistent within households and heterogeneous across them. Yet standard modeling practice assumes ex-ante identical households and accounts for heterogeneity only in shock realizations. We consider persistent heterogeneity in a key aspect of consumer decision making—(biased) beliefs about one's own future financial situation—and show that it strongly conditionally correlates with actual financial decisions and conditions. We use novel microdata to quantitatively discipline ex-ante optimism heterogeneity in an otherwise standard HANK model. Optimistic bias drives households to spend instead of accumulating buffer stock saving and thereby produces more empirically realistic HtM, wealth, and MPC distributions, with optimistic households exhibiting higher MPCs than their rational counterparts even away from the borrowing constraint. Accounting for optimism makes targeted transfers less stimulative and incentivizing self-insurance less effective, but also implies that public insurance is less distortionary.

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

Keywords: Household Heterogeneity; Optimism; Household Financial Situation; 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

Early drafts of this paper were titled "Heterogeneity in what? Cognitive Skills, Beliefs and the Liquid Wealth Distribution" and "Bad luck or bad decision? Macroeconomic Implications of Persistent Heterogeneity in Cognitive Skills and Overconfidence". We thank Christian Bayer, Eduardo Dávila, Greg Kaplan, Rohan Kekre, Alexander Kriwoluzky, Ralph Luetticke (discussant), Peter Maxted, Kurt Mitman, Christina Patterson, Matthew Rognlie, Johannes Stroebel, Hannes Twieling, Gianluca Violante, Michael Weber, Nathan Zorzi, and seminar and conference participants at the Bank of Finland, the ifo Conference on Macroeconomics and Survey Data, EEA/ESEM, the University of Zurich, the University of Virginia, the University of Surrey, the Bonn-Berlin PhD Workshop, the Swiss Economists Abroad end-of-year conference, the HeiTueHo Workshop, EWMES, BEAM, UC-Berkeley Psychology and Economics of Poverty Convening, the Macro Mini-Conference at Dartmouth, the 5th Behavioral Macroeconomics Workshop, the Sailing the Macro Conference, the 5th Joint BoC - ECB - NY Fed Conference, and the ASSA meeting for helpful comments and suggestions. Seyrich gratefully acknowledges financial support by the Leibniz Association through the project "Distributional effects of macroeconomic policies in Europe".

1 Introduction

Household heterogeneity in 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

It nevertheless 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 and persistence in household expectations (see e.g. D'Acunto and Weber (2024) for a recent review).

Although macroeconomic work on beliefs has focused on aggregate variables, standard consumption-savings theory predicts that what matters most directly for households' savings and consumption choices is their expectations about their *own* future financial situation. Such beliefs 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 that such forecasts tilt strongly optimistic in the aggregate, across various time periods, datasets, and countries (Souleles, 2004; Claus and Nguyen, 2023; Cocco et al., forthcoming; D'Acunto et al., 2020), with many consumers making optimistic forecast errors.

We build on this work by developing a new approach for quantifying links between exante persistent heterogeneity in optimism, ex-post forecast errors, and consumption-savings decisions. We focus on persistent heterogeneity because quantitative macro models have been struggling to jointly fit the data on key variables that are empirically quite persistent: wealth, MPCs, and hand-to-mouth (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 optimistic forecast errors, not saving, and being HtM.

These results suggest that heterogeneity in optimism can indeed help explain observed heterogeneity in households' savings behavior, and they guide us in augmenting an otherwise standard heterogeneous-agent New Keynesian (HANK) model with empirically-disciplined belief

¹See Kaplan and Violante (2022), McKay and Wolf (2023), and Auclert et al. (2025) for recent review articles on how household heterogeneity shapes macroeconomic fluctuations, policy effectiveness, and design. Aguiar et al. (2024) focuses on the empirical persistence of hand-to-mouth status. Fagereng et al. (2021), Ganong et al. (2024), and Lewis et al. (2024) find that unobserved and persistent individual characteristics are important drivers of heterogeneity in MPCs.

 $^{^{2}}$ A few important exceptions have focused on preferences, as we discuss below.

³Evidence from other fields suggests that optimistic bias is a persistent feature of human decision making, as we document below.

heterogeneity. The model jointly matches key untargeted moments in micro and macro data, even with only one asset, 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 about *ex-ante* forecast biases from observed *expost* forecast errors. 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 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.

Specifically, we quantify the variation in ex-post optimism predicted by a well-established and measurable source of persistent behavioral heterogeneity in economics: *overconfidence*. Overconfidence is unlikely to be driven by bad luck because overconfidence is—like ex-ante optimism—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). We accordingly proxy for ex-ante optimism using a measure of overconfidence about one's own cognitive skills, and specifically about one's relative performance on an intelligence test, that exhibits trait-like temporal stability (Stango and Zinman, 2020, 2024).⁴ Related forms of overconfidence about performance have been documented as persistently prevalent and impactful in high-stakes workplace settings (e.g., Hoffman and Burks (2020); Huffman et al. (2022); Weidmann et al. (2024); Heck et al. (2024)).

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 to 36 percentage points more likely to make optimistic forecast errors about their own future financial situation than households in the 25th percentile.

We then find strong conditional correlations between savings behavior and overconfidence, as well as between actual financial situations (HtM status) and overconfidence. These findings suggest that overconfidence not only produces optimistic forecast errors but also spurs households to save less and become financially constrained.

⁴Another form of overconfidence concerns the precision agents assign to their forecasts, which is different from our definition (Moore and Healy, 2008; Broer and Kohlhas, 2024). Recent empirical work finds that various forms of overconfidence are correlated within-person (Chapman et al., 2023; Stango and Zinman, 2023).

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, as estimated in the micro data.⁵

In our baseline model specification, we microfound differences in ex-ante optimism as differences in beliefs about one's own future idiosyncratic labor productivity (and then show, in extensions, that our model results are qualitatively robust to other microfoundations including ex-ante optimism about aggregates and about expenses). Focusing on labor productivity and hence 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. Our modeling approach also specifically implies that ex-ante optimists will tend to overestimate their *rank* in future income distributions, which further strengthens our mapping to the model from micro data on overconfidence about *relative* performance. Focusing on income also squares with mounting empirical evidence that consumers hold non-FIRE beliefs about it (d'Haultfoeuille et al., 2021), with many making optimistic forecast errors (e.g., Souleles (2004); D'Acunto et al. (2024)).⁶

The model formalizes how optimistic households tend to consume rather than to accumulate a buffer stock, creating heterogeneity in savings behavior across otherwise identical households. Optimists persistently forecast a better future, dampening their subjective valuation of selfinsurance in the form of buffer stock 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 household strives both to avoid the HtM state, and to save its way off the constraint after having reached it.⁷

Accounting for such heterogeneity in beliefs greatly improves the model's fit to several key moments in micro and macro data. From a micro perspective, our model implies MPC

⁵In an extension, we also add households with a pessimistic bias. This actually slightly increases the share of HtM households and the average MPC in general equilibrium, as we detail below. We further show robustness to a specification in which heterogeneity in optimism is persistent but not permanent.

⁶Although we lack the requisite data on income realizations to measure income forecast errors, we do show that households' financial situation forecasts are strongly positively correlated with their income forecasts.

⁷This intuition also provides evidence against an alternative, rational-expectation (RE) interpretation of our empirical findings: overconfidence merely reflects higher risk that drives the higher ex-post optimism. The RE explanation is unlikely for three reasons: First, our regressions control for a wide set of consumer-level characteristics that should absorb much of any variation in risk exposure, including risk aversion, education, cognitive skills, and income. Second, within a standard modeling framework using a log-normal, idiosyncratic productivity process, quantitatively accounting for the observed differences in ex-post optimism would require ex-post optimistic households to face an order of magnitude more income risk than their rational counterparts. This seems implausible for otherwise identical households. Third and most fundamentally, if optimistic households indeed faced higher risk and were rational about it, these households would save more and have a lower propensity to be hand-to-mouth. 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.

heterogeneity across households even conditional on wealth and income, including a relatively flat MPC-income distribution in line with recent findings (see e.g. Fuster et al. (2021), Boehm et al. (2023), Nielsson et al. (2025)).⁸ This then allows 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)). In contrast, rational one-asset HANK models fail to jointly match average wealth and average MPCs and produce counterfactually few HtM households (Auclert et al., 2024b; 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., 2024b).

Accounting for heterogeneity in optimism also generates important and distinct implications for macroeconomic policies.

We first analyze the aggregate consequences of unexpected transfer payments intended to stimulate private consumption. Such interventions often target lower-income households (e.g., U.S. stimulus checks during the Great Recession and the COVID pandemic). 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, and in the data, dictates that income-targeted transfers generate a much weaker output response 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 transfers to the bottom income quartile produce 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 self-insurance incentives, muting the aggregate response of households compared to models in which households are uniformly rational. This makes interventions to incentivize self-insurance less effective, and to provide insurance directly more effective.

Direct insurance, which we model in the form of minimum income benefits, is more effective because optimistic households are much less likely to substitute away from private savings than their rational counterparts. Optimistic households undervalue the insurance benefits because

⁸Also 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.

they underestimate their probability of reaching bad income states. As such they 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 rational models.

Indirect insurance, which we model in the form of government debt issuance (e.g., Woodford (1990), Aiyagari and McGrattan (1998)), is less effective because it is difficult to induce optimistic households to accumulate sufficient buffer stock savings. Higher government debt levels reduce households' self-insurance cost by reducing the price of liquid assets. But the induced increase in 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 instead 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. We find that these contrasting effects can have normative implications as well: the optimal government debt level is substantially lower in our model with heterogeneity in optimism, irrespective of whether we consider the 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 public insurance.

Related literature. We contribute to several literatures.

Our paper connects various literatures considering potential roles for optimism (and/or pessimism), including work on sentiment (Pappa et al., 2023; Bhandari et al., 2024; Kamdar and Ray, 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 (as we

consider in Section 4.4.2). Another input we consider is beliefs about expenses, which could also be quantitatively important for macro given evidence on the prevalence, heterogeneity, and magnitude of expense neglect (Berman et al., 2016; Kaur et al., 2025) and expense shocks (Fulford and Low, 2024). Prior work considering consumer beliefs about their own financial situations has either not analyzed consumer-level forecast errors (Claus and Nguyen, 2023; Kamdar and Ray, 2024) or not quantified the role of ex-ante optimism in producing them (Souleles, 2004; Cocco et al., forthcoming).⁹ Our modeling of optimism about one's financial situation as being driven by ex-ante 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.¹⁰

We also contribute to behavioral macroeconomics more broadly, by accounting for behavioral heterogeneity in decision making. Other work focuses on a representative behavioral agent.¹¹ Behavioral HANK models tend to allow for heterogeneity only in the budget constraint, with a homogeneous behavioral bias or a homogeneous information friction about an aggregate variable only.¹² Kaplan and Violante (2022) do 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 (2024) 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 where agents reason less in familiar states and this can lead to learning traps in which households remain persistently at the borrowing constraint. While they focus on ex-ante identical households, ours are ex-ante heterogeneous.

⁹Cocco et al. (forthcoming) 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).

¹⁰We 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 ex-ante 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 better assess the model's fit to the data and its policy implications.

¹¹See, e.g., Woodford (2013), Gabaix (2014), Woodford (2019), Gabaix (2020), Bordalo et al. (2020), Boutros (2023), and Lian (2023).

¹²See, e.g., Farhi and Werning (2019), Auclert et al. (2020), Angeletos and Huo (2021), Laibson et al. (2025), and Pfäuti and Seyrich (2024). Broer et al. (2022) analyzes heterogeneous expectations in a Krusell and Smith (1998) model without full information, abstracting from nominal rigidities.

We also discipline our behavioral parameter with data, leaving the MPC untargeted, while they use MPC estimates to calibrate their key behavioral parameter. Additionally, our model features nominal rigidities and we take an important step, beyond the crucial one of matching key empirical moments, by applying our model to fiscal and public insurance policy.

A parallel strand of literature considers (persistent) heterogeneity in reduced-form or presumedclassical preferences. Aguiar et al. (2024) find that allowing for heterogeneity in patience and the elasticity of intertemporal substitution (EIS) helps match several empirical facts about HtM households. They suggest that behavioral factors might provide a potential microfoundation for their modeling choices. Krueger et al. (2016), Carroll et al. (2017), and Auclert et al. (2020) introduce permanent heterogeneity in patience to better match wealth inequality data. Kaplan and Violante (2022) show that EIS heterogeneity can produce similar results to discount factor heterogeneity for HtM shares and MPCs. They also show, however, that allowing for such heterogeneity 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 shows that the correlation of patience with HtM status is relatively weak compared to ex-ante optimism, both qualitatively and quantitatively.¹³ And whereas prior work uses the degree of ex-ante heterogeneity as a free parameter to match some empirical targets, we quantitatively discipline our ex-ante 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 its own future financial situation is a key determinant of its savings behavior. The more optimistic the household ex-ante, the less it saves, *ceteris paribus*.

We use novel micro panel data to quantify heterogeneity in ex-ante optimism, and in particular its relationship to ex-post forecast errors. We first construct household-level own-financial situation forecast errors (FCEs), taking data limitations into account in various ways, and find a strong slant towards persistent ex-post optimism. Since ex-post optimism can be due in part to bad luck, we then develop an empirical strategy using persistent overconfidence about one's cognitive skills to help identify how ex-ante optimism drives ex-post optimistic FCEs. We

¹³We do not observe consumers' EIS. We do, however, control for risk aversion, which in many standard macro models is equal to the inverse of the EIS.

further show that ex-ante optimism is indeed strongly conditionally correlated with less saving and HtM status.

We use the American Life Panel (ALP), as it is the only dataset we know of that permits measurement of each of our key ingredients at the household level: ex-post optimism, ex-ante behavioral biases, savings behavior and overall financial condition, and a rich set of demographics and other characteristics (including discounting and risk aversion). The ALP's long-running panel component also enables us to estimate persistence in our key variables. The ALP goes to great lengths to obtain a nationally representative sample of U.S. adults, but we report unweighted as well as sampling probability-weighted estimates, 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 prevalence of ex-post optimism

We start by constructing our key input for measuring ex-post optimism: consumer-level forecast errors about households' 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 2009 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)).¹⁴ 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 too optimistic, accurate, or too pessimistic about their financial situations, ex-post. Table 1 illustrates this classification. This coarse classification of potential forecast error realizations yields two nuanced measurement challenges that prior work has not accounted for, to our knowledge.

The first measurement challenge is that the set of feasibly-observed FCEs depends on the consumer's forecast and/or realization. For example, it is impossible to observe an optimistic

¹⁴Forecasts are highly correlated with expected own-income growth in the ALP surveys that also elicit an income forecast (Appendix Table A1; see also Kamdar and Ray (2024)). Bhandari et al. (2024) and Kamdar and Ray (2024) show that own-condition forecasts also correlate strongly with forecasts about aggregates. 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.

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

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

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. 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 three partially overlapping samples of forecast errors: "potentially symmetric", "potentially optimistic", and "potentially pessimistic".

The second measurement challenge posed by categorical data constraints is difficulty detecting some ex-post forecast errors: for example, some better-better pairs will be misclassified as accurate when they are in fact ex-post optimistic. For example, consider a forecast of 20 percent improvement coupled with a 10 percent improvement realization. This pair would be misclassified as "accurate", because both the forecast and realization are measured as "better". This problem is particularly germane when there is an aggregate slant toward optimism bias in the aggregate, as we document below.¹⁵ The same issue applies to worse-worse pairs. To mitigate this potential measurement bias against detecting ex-post forecast errors, we consider additional "v2" samples that exclude better-better pairs from the "potentially optimistic" sample and worse-worse pairs from the "potentially pessimistic" sample. Appendix Table A2 illustrates our five FCE samples and shows that their mean counts of measurable FCEs per panelist range from 4.2 to 6.2.

2.1.1 Facts about ex-post optimism

We now develop some descriptive statistics about the forecast errors, distilling them into two facts that motivate our focus on optimism in the rest of the section and paper.

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

¹⁵We do not have to worry about the opposite case. For example, a consumer expecting to be 10 percent better off but then being 20 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".

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., forthcoming).¹⁶ As Appendix Table A3 shows, 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, 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, for each of our samples, the mean proportion of a consumer's observed FCEs that are optimistic or pessimistic—note that the denominators include all FCEs, including accurate ones (FCE=0)—and the share of consumers making optimistic/pessimistic FCEs at least half the time. Panel C then reports estimated ratios of sample proportions of optimism and pessimism, or of discretized versions thereof, for each of our sample definitions. Across these 12 estimates, optimism is 1.83 to 2.91 times more prevalent than pessimism.¹⁷

Sample	Potentially opt or pess FCEs				Potentially symmetric		
Sample	Al	All			FCEs		
Estimate	Unwtd	Wtd	Unwtd	Wtd	Unwtd	Wtd	
	(1)	(2)	(3)	(4)	(5)	(6)	
Panel A. Prop of FCEs that are optimistic	0.33	0.33	0.39	0.38	0.33	0.32	
Share consumers with optimistic proportion ≥ 0.5	0.32	0.32	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	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	1.94	1.95	2.00	2.20	2.13	
Shares ≥ 0.5	2.29	2.29	2.35	2.29	2.91	2.91	

 Table 2: Household financial condition FCE proportions

Note: Denominators in Panels A and B include accurate forecasts: see Appendix Table A2 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.

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

¹⁶The one counterexample we know of is Hyytinen and Putkuri (2018)'s evidence of aggregate mean-zero forecast errors from Finland.

¹⁷For comparison, Panel C of Table A4 reports relative optimism estimates that are not adjusted for measurement error created by the categorical data constraints detailed above. Estimates are uniformly higher for each functional form and (un)weighting combination, ranging from 2.33 to 5.00.

the absolute prevalences of optimism and pessimism. In contrast, estimates of relative prevalence 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. Using the potentially symmetric FCE sample to allow for apples-to-apples consideration of persistent pessimism as well as optimism, we 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 A5 illustrates these results in more detail.

2.1.2 The insufficiency of ex-post FCEs for identifying expectation biases

Having established the relative prevalence of optimistic vs. pessimistic forecast errors, we turn towards quantifying the decision-relevant quantity: the extent of any ex-ante optimistic bias. This is challenging because ex-post forecast errors reflect the combination of an ex-ante bias and an unpredictable shock realization, and the lack of a rational benchmark renders ex-ante optimism not directly observable. The challenge is especially pronounced when the belief of interest regards a consumer-specific variable, as in our case.¹⁸ Absent a convincing rational expectations benchmark, even a persistent pattern of ex-post optimistic FCEs could in principle be due to a spell of bad shocks (Souleles, 2004). We develop an empirical strategy that addresses this challenge.

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

We distinguish ex-ante from ex-post optimism by identifying the variation in ex-post optimism predicted by a conceptually-linked, persistent, and trait-like ex-ante behavioral bias. This strategy allows us to develop two key new sets of facts: one on the relationship between ex-ante optimism and ex-post optimism, and one on the relationship between ex-ante optimism, savings behavior and financial constraints.

 $^{^{18}{\}rm For}$ some aggregate beliefs, one can use professional forecasts and/or a long time series to infer the rational expectations benchmark.

2.2.1 Overconfidence

Overconfidence about one's own abilities—that is, overestimating one's own abilities (e.g. Malmendier and Taylor, 2015)—is a priori a promising candidate for identifying ex-ante optimism. Overconfidence and optimism are conceptually kindred manifestations of 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). 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 (e.g., Huffman et al. (2022)), suggesting that overconfident consumers do not fully account for any association between their overconfidence and shock realizations—if they did, overconfident consumers would instead make the same decisions as classically rational consumers and realize the same outcomes.

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)).¹⁹ We then define a consumer's degree of overconfidence as their self-assessed rank minus the actual rank—a higher value thus indicates more overconfidence.²⁰ This measure exhibits a high degree of stability within-panelist over time (Stango and Zinman, 2020, 2024).

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:

Prop. optimistic errors_i =
$$\beta_0 + \beta_1 \cdot oc_i + \Gamma \cdot X_i + u_i$$
. (1)

¹⁹Number series scores correlate strongly with those from other fluid intelligence tests like IQ and Raven's. ²⁰The 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)). 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).

Prop. optimistic $errors_i$ is the proportion of optimistic forecast errors of a consumer *i*, as described in Table 2. oc_i is overconfidence, our measure of ex-ante optimism. 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.

 X_i is a vector of control variables for cognitive skills, income, standard demographics (age, gender, education, race, and ethnicity), and preferences.²¹

As detailed in Stango and Zinman (2023), 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.²²

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)—(our three samples for measuring ex-post optimism from ex-post FCEs) × (unweighted and weighted estimates)—and we find a positive coefficient in all 6 cases. Five of these six specifications have p-values of smaller than 0.1, despite the modest sample sizes.²³

Quantitatively, the estimates in Table 3 are quite substantial. Being one percentile higher in the overconfidence distribution is associated with a 0.24 - 0.72 percentage point increase in the proportion of optimistic forecast errors. This corresponds to someone at the upper end of the interquartile range being 12 to 36 percentage points more optimistic, in terms of the proportion of optimistic FCEs, than someone at the lower end of the IQR. (For benchmarking purposes, the sample mean proportions range from 0.32 to 0.44.) The average across our six estimates in Table 3 corresponds to an interquartile range of 20 percentage points. This steep slope will be

²¹We measure preference heterogeneity using 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)).

²²We 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.

 $^{^{23}}$ If we run the same regressions but with the proportion of pessimistic forecast errors instead of optimistic ones as our dependent variable, we find that the overconfidence coefficient is negative in all six variations of regression (1), as shown in Appendix Table A6. These estimates are relatively noisy due to the fact that we observe relatively few pessimistic forecast errors, but the strong sign pattern lends further support to our prior that overconfidence is in fact directionally related to ex-ante optimism and not just to an overall tendency of having misperceived beliefs.

a key target to discipline our model later.²⁴

LHS = Optimism FCE proportion								
Sample	Po	tentially	Optimist	ic	Potential	ly Symmetric		
Sample	v	1	v2					
Estimate	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)		
Controls?	Yes	Yes	Yes	Yes	Yes	Yes		
Ν	778	778	750	750	742	742		
N Panelists	389	389	375	375	371	371		
Mean (LHS)	0.36	0.32	0.44	0.42	0.38	0.37		

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

Note: Standard errors, clustered on panelist, in parentheses. One regression per panel-column. RHS variables measured using the Stango-Zinman data and merged onto our household financial situation forecast error panel. The results shown represent 100×the change in the LHS variable associated with a 1 percentage point increase in the overconfidence rank. "Controls" include income, standard demographics (education age, gender, race, and ethnicity), patience, two measures of risk aversion, 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.

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

Table 4 shows that ex-ante optimism is also strongly correlated with consumers' savings behavior and HtM status. Specifically, we run regression (1) with a measure of the household's behavior or their actual financial situation as the dependent variable. Our key measure of behavior is an indicator for not saving over the prior year on a flow basis (including debt paydown as saving). Our key measure of HtM status is an indicator for having experienced severe financial distress over the past year.²⁵ Three of our four estimates of the conditional

²⁴For robustness, Panel B of Appendix Table A7 considers our alternative measure 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 these 6 have p-values 0.05, further reassuring us that overconfidence is strongly correlated with ex-post optimism.

²⁵We 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. Various reasons motivate focusing on this measure of HtM status. First, being backward-looking, it avoids a potential mechanical correlation with optimism (e.g., 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). Second, our ALP data has limited coverage of liquid assets.

correlation between overconfidence and these variables have p-values < 0.05, and all four are positive. The magnitudes of overconfidence's conditional correlations with actual financial situations and behavior are substantial, with a one percentile increase in overconfidence associated with a 0.26 or 0.61 percentage point decrease in the likelihood of having saved and a 0.14 or 0.52 increase in the likelihood of being financially constrained. Across the ex-ante optimism interquartile range, these numbers correspond to a 13 or 31 pp increase in the likelihood of not having saved (n.b. the sample mean is 0.57) and a 7 or 26 pp increase in the likelihood of being financially constrained (sample mean of 0.37). These results suggest that beliefs, in the form of ex-ante optimism, meaningfully affect saving decisions and financial outcomes.

Heterogeneity in patience has been offered as an alternative explanation for savings heterogeneity (Krueger et al., 2016; Carroll et al., 2017; Aguiar et al., 2024), but Table 4 favors overconfidence over patience as the key margin of ex-ante heterogeneity. First, note that although Columns 1 and 2 suggest that patience is indeed negatively correlated with our savings indicator—as standard consumption theory, including our own model later on, predicts—the Column 2 estimate is statistically weaker than overconfidence's. Second, the link between patience and HtM is weaker than for overconfidence and HtM, as Columns 3 and 4 show. Both patience estimates have larger p-values than overconfidence's in the same regression and neither is statistically significant at conventional levels.

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

LHS:	1 = (Dic	l not save last year)	1 = (1)	HtM)
Estimate	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)
Other Controls?	Yes	Yes	Yes	Yes
Ν	1355	1355	1358	1358
N Panelists	680	680	680	680
Mean~(LHS)	0.57	0.59	0.37	0.40

 Table 4: Overconfidence strongly conditionally correlates with actual financial situation

Note: Same regression specification as Table 3 but with different LHS variables (please see the main text for details on how we define those).

Potential concern. One might wonder whether the results in Table 3 could be driven by bad luck instead of ex-ante optimism. Specifically, consider a hypothesis that overconfident households are not ex-ante optimistic but instead relatively prone to bad luck.

This hypothesis is not well-supported by our results. Recall that we control for a rich set of consumer-level characteristics, including risk aversion, education, cognitive skills, and income. This implies that any residual correlation between overconfidence and risk realizations is unlikely to be strong enough to materially confound our estimates of the conditional correlation between overconfidence and optimism. More fundamentally, the hypothesis that overconfidence is conditionally correlated with bad shocks, but not with ex-ante optimism, implies counterfactual predictions for Table 4. This is because a classically-rational consumer (someone who is not ex-ante optimistic) facing more risk should save more, not less as we find. We formalize this in Section 4.1, together with the prediction that a non-optimist facing more risk should be less likely to end up in HtM states, not more likely as we find.

In sum, overconfidence as an ex-ante optimistic bias can jointly explain our findings in Tables 3 and 4, whereas overconfidence merely reflecting greater risk exposure for classically rational consumers cannot.

2.3 Taking stock and outlook

This section develops four sets of micro facts that will motivate and help discipline our quantitative model in subsequent sections. First, consumers' forecasts and FCEs about their own future financial situation tilt strongly optimistic on average, both across-households and within them over time. Second, optimistic FCEs are persistent within-consumer, and much more persistent than pessimistic FCEs. Third, optimistic FCEs can be partly explained by ex-ante optimism, and the degree of this ex-ante optimistic bias is heterogeneous across households. Fourth, optimistic bias is strongly conditionally correlated with being less likely to save and more likely to be HtM. Next we augment a standard heterogeneous-agent New Keynesian model to explore some macroeconomic consequences of persistent heterogeneity in ex-ante optimism.

3 Model

We now develop an augmented HANK model that adds persistent belief heterogeneity, disciplined by our findings in the previous section, to HANK's usual heterogeneity through idiosyncratic shock realizations (the "luck" invoked in this paper's title). 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 the standard 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. Building on our motivation and empirical findings above, 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. Accordingly, households permanently belong to different types denoted by g.²⁶ An individual household's productivity is denoted by e_t and 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 beginningof-period asset holdings b_{t-1} , is given by:

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

subject to

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

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

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 where our key innovation enters, and we discuss it next.

Belief heterogeneity. Capturing household beliefs about their own financial situations requires us to make two modeling decisions. First, we must define "financial situation". We equate this with the right-hand side of the budget constraint, $b_{t-1} + (1 - \tau_t)w_t e_t n_t$. Second, we must choose the object of optimism, i.e., about which part of the right-hand-side of the budget constraint some households are ex-ante too optimistic (or pessimistic, in an extension).

In our baseline model, we focus on ex-ante optimism about one's idiosyncratic productivity. This approach is empirically supported by evidence that own-situation forecasts are highly correlated with expected own-income growth in the ALP surveys that also elicit an income

²⁶We emphasize permanent heterogeneity but show that our results hold in a version of the model where beliefs are not fully persistent and households can switch between types (Appendix D.3).

forecast (Appendix Table A1; see also Kamdar and Ray (2024)).²⁷ Focusing on idiosyncratic productivity makes sense conceptually as well, since our proxy for optimism—overconfidence—concerns cognitive skills, and such skills have well-established links to income and workplace performance. Our modeling approach specifically implies that ex-ante optimistic households overestimate their future *rank* in the productivity distribution, which maps well into our empirical measure of overconfidence in *relative* performance. Focusing on productivity moreover makes sense technically and conventionally, as it is usually the only source of uncertainty in standard incomplete markets models.

We present two alternative approaches in Section 4.4.2. One assumes that optimism about one's own financial situation instead stems from optimism about aggregates. The second assumes that optimism about future financial situation stems from expense neglect (i.e., from underestimating future expenses). Each of these alternatives produces similar qualitative results as our baseline model where optimism operates through idiosyncratic productivity.

Specifically, 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 and thus their income. Note that, absent aggregate risk, a household's optimism about future productivity necessarily implies 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 $e_i \in \{e_1, e_2, \dots, e_J\}$ reaches productivity $e_j \in \{e_1, e_2, \dots, e_J\}$ in the following period, and assume that the productivity states are ordered such that $e_1 < e_2 < \dots < e_J$. 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

$$\tilde{p}_{ij} \equiv \begin{cases}
\alpha_g p_{ij}, \text{ if } i < j \\
\frac{1}{\alpha_g} p_{ij}, \text{ if } i > j \\
1 - \sum_{j \neq i} \tilde{p}_{ij}, \text{ if } i = j,
\end{cases}$$
(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 1.²⁸ Note that rational expectations

²⁷There is also empirical evidence showing that households, on average, systematically overestimate their job finding probability <u>Mueller et al.</u> (2021) and underestimate their job losing probability <u>Balleer et al.</u> (2024).

²⁸We 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.

are captured by setting $\alpha_g = 1$ and thus are nested in our setup.²⁹ In Section 4.4, we discuss an extension with pessimistic households, $\alpha_g < 1$.

An immediate implication of equation (4) is that optimistic households will more often make optimistic forecast errors about their own financial situation compared to rational households, consistent with the empirical findings reported in Section 2.2.2. We discipline this degree of belief heterogeneity, that is the group specific values of α_g , by directly targeting Table 3's empirical estimates of how much more frequently ex-ante optimistic households make optimistic forecast errors (see Section 3.1).

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. (2024b), which is based on Erceg et al. (2000)).³⁰ 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 appears additively 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.

²⁹Modelling 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. Appendix D.3.1 provides an alternative modeling approach, where the degree of optimism depends on the distance between the states. Our results are robust to this approach.

³⁰Auclert et al. (2023) and Broer et al. (2020) argue in favor of using sticky wages rather than sticky prices in HANK models.

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

$$Y_t = N_t, (6)$$

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.

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

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

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

$$B_t + T_t = (1 + r_t)B_{t-1}.$$
(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:

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

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

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 idiosyncratic productivity, we follow McKay et al. (2016) in assuming 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 and normalize the average productivity level 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)). We set the elasticity of substitution between unions, ϵ , to 11 as in Bayer et al. (2024a).

To reflect our empirical findings on heterogeneity in optimism about own financial situation, we assume that there are three different groups, indexed by g, that differ only in their degree of ex-ante optimism α_g . Our data provides little guidance on the size of these groups and thus we rely on the estimates in Huffman et al. (2022), where about half of workers in their population are persistently overconfident in high-stakes workplace tournaments. We therefore assume that 50% of households are *rational*, with $\alpha_r = 1$. This group corresponds to the bottom 50% of our empirical ex-ante optimism (overconfidence) distribution. The other 50% of households have optimistic-biased beliefs, to varying degrees.³¹ 25% percent are *mild* optimists, corresponding to the third quartile of our ex-ante optimism distribution, and the remaining 25% are *strong* optimists corresponding to the top quartile.

We calibrate α_g for mild and strong optimists by targeting the quantitative relationship between ex-ante optimism and ex-post optimistic forecast errors implied by Table 3. Concretely, we match the difference in average proportion of optimistic forecast errors between each optimistic group and the rational group, using Table 3's regression models to predict ex-post optimism at the median of each group. So for mild (strong) optimists we predict the proportion of optimistic forecast errors for someone at the 62.5th (87.5th) percentile of overconfidence, subtracting off the predicted value for someone at the 25th percentile. We conservatively use

 $^{^{31}}$ Section 4.4 shows that allowing for an empirically-realistic mass of pessimists does not change our results appreciably. In addition, Appendix D.1 shows that our results are robust if we assume that there are only two groups—one rational and one optimistic—or if we assume that there are less (33%) or more (60%) optimistic households.

the smallest predicted differences produced by our six specifications in Table 3: 10 percentage points for mild optimists relative to their rational counterparts, and 15 pp for strong optimists. Given our income risk process, this results in $\alpha_m = 1.6$ and $\alpha_s = 2.1$, and implies that mild (strong) optimists make forecasts about their income over the next year that are, on average, 10% (18%) higher than rational households'.

Parameter	Description	Value
R	Steady state real rate (annualized)	4%
β	Discount factor	0.985
γ	Risk aversion	2
arphi	Inverse of Frisch elasticity	2
\underline{b}	Borrowing limit	0
$\frac{\bar{B}}{4\bar{Y}}$	Average wealth to average income	4.1
ϵ	Elasticity of substitution between unions	11
Idiosyncratic risk		
$\overline{ ho_e}$	Persistence of idiosyncratic risk	0.966
σ_e^2	Variance of idiosyncratic risk	0.033
Permanent heterogeneity		
$\overline{\mu_g}$	Mass of households	$\{0.5, 0.25, 0.25\}$
α_g	Degree of optimism	$\{1, 1.6, 2.1\}$

 Table 5: Stationary equilibrium calibration

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

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 rational HANK models that abstract from belief heterogeneity.

4.1 Heterogeneity across observationally equivalent households

We first show how heterogeneity in ex-ante optimism affects households' behavior at the microlevel. Belief heterogeneity introduces heterogeneity across observationally equivalent households, that is, across households with the same level of wealth and income. The key mechanism in our model is that optimists, given their expectation of relatively high future income, perceive their saving motives to be weaker than their otherwise identical rational counterparts. Figure 1 illustrates how this affects household behavior.

The upper panel in Figure 1 shows our baseline model's savings policy functions, for each of the three permanently different household groups, along the wealth distribution for three different income states. Each sub-figure plots separate functions for each of the rational, mild optimist, strong optimist groups. These show that, conditional on the idiosyncratic state households' current income and wealth—the net savings of a household (weakly) decreases in her optimism, almost always by a substantial amount. There are even many states in which strong optimists dissave while other households are accumulating a buffer against future shocks. Overall, optimists save (weakly) less in all states of the world. Two implications not directly illustrated here, that we detail later, are that optimists accumulate less wealth over time and are much more likely to enter and persist in a HtM state.

The lower panel in Figure 1 shows the MPCs out of unexpected, one-time income windfalls of \$500. Again, heterogeneity in ex-ante optimism produces heterogeneity across observationally equivalent households. The more optimistic households are, the higher their MPC even conditional on the idiosyncratic state.³² These differences can be quite strong outside of highincome states; 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 a middle-income state, because they expect their future financial situation to be better. Second, even if they are not constrained, more optimistic households perceive a transfer of a given size as relatively small compared to their expected lifetime income. Thus, conditional on their idiosyncratic state, they are in the more concave part of their policy function compared to rational households expecting a lower lifetime income. Overall, our results here can be interpreted as surfacing the empirical latent MPC heterogeneity inferred by Lewis et al. (2024).³³

What if optimists face more risk? We now return to the potential concern raised in Section 2.2.2, here using our model to further illustrate how an alternative, risk-based interpretation of ex-post optimism fails to fit several key pieces of empirical evidence.

Say overconfident households do not hold ex-ante biased beliefs—that is, that they hold rational expectations—but do face higher risk. Due to the log-normality of the productivity shock process, this alternative version of the model can reproduce the same increases in the likelihood of making ex-post optimistic forecast errors we target in our baseline model, but at the cost of producing wildly counterfactual predictions on other dimensions.

Specifically, the rational-and-riskier model of ex-post optimism requires mild (strong) optimists to face 9x (12x) more income volatility, driven by different idiosyncratic shock volatilities.

 $^{^{32}}$ Another way to see this is that the savings policy functions in the upper panel of Figure 1 are steeper for the more optimistic households, implying that the MPCs out of a marginal increase are higher for more optimistic households.

³³Our results are also consistent with Koşar and Melcangi (2025)'s finding that the MPC is increasing in the subjective earnings growth uncertainty of households, because in our model optimists end up having slightly higher subjective uncertainty (measured as their perceived earnings volatility, which is 3% and 9% higher for mild and strong optimists), and they also have higher MPCs as we document above.





Note: Given that we have 11 productivity levels, here we show the third level for low-income, the sixth is the median, and the ninth for high-income. MPCs out of a surprise \$500 stimulus check. For MPCs, we shrink the wealth range on the horizontal axis to focus on the most concave part of the MPC functions. The ordering of the three belief groups remains the same at higher wealth levels.

Even if one accepts these greater volatilities as empirical realistic, the rational-and-riskier model's predictions on saving and HtM status are incorrect by wide margins. This is because a rational household accounts for any added volatility when making savings decisions—and added volatility of course increases the value of precautionary saving. As such, the alternative model predicts that mild (strong) optimists save 1.7x (1.9x) more than the rational group (see Appendix Figure D4 for the savings policy functions). This runs counter to our previously discussed empirical findings that optimists save substantially *less* (Table 4). Even more strikingly perhaps, the alternative model predicts HtM shares of 0% for optimists—that is, that approximately zero optimistic households will be HtM—and 19% for rational households. Here again the alternative model incorrectly signs the direction of optimistic vs. rational heterogeneity, per our previously discussed empirical finding that optimists are much more likely to be HtM compared to their rational counterparts (Table 4).

In short, even if ex-post optimistic households do face substantial additional risk, they must also hold ex-ante optimistic beliefs to fit the data.

4.2 Hand-to-Mouth Shares and Average MPCs

Having considered our model's performance in reproducing key untargeted patterns in the micro data, we now turn to key macro moments, starting with the aggregate HtM share and average marginal propensity to consume (MPC) of households. We compare our model's predictions to a standard "Rational HANK" model that differs from ours only in its lack of any ex-ante optimism. This benchmark model employs the standard HANK specification of rational and ex-ante identical consumers.³⁴

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. (2024b), 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.

	HANK w/ Belief Het.	Rational HANK
	(1)	(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	-

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

Note: MPCs are out of a surprise \$500 stimulus check. "HANK w/ Belief Het." is our baseline model (one-asset, with heterogeneity in optimism). "Rational HANK" denotes its rational counterpart with no heterogeneity in optimism (i.e., with ex-ante identical households).

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%.³⁵ Our predicted share of HtM households, 0.23, is lower than our estimated empirical share of around 0.40 (Table 4),

³⁴When comparing our model to its rational counterpart, we take the standard approach of recalibrating the discount factors 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 each model.

 $^{^{35}}$ 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. (2025)).

but given our one-asset model's standard strict definition of HtM status this discrepancy is not surprising.³⁶

Table 6 also decomposes the aggregate HtM share and MPCs across groups to further illustrate the how optimism affects key aggregates. Following our discussion in the previous section, optimistic households push 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 only 2% of rational households are HtM (for the same reason as in the rational model), 24% of the mild optimists and 68% of strong optimists are HtM. These higher shares contribute to higher MPCs, together with higher MPCs even away from the 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 empirical estimates (Kaplan and Violante (2022), Seidl and Seyrich (2023), Auclert et al. (2024b), Wolf (2025)). This restriction also produces an excessively polarized wealth distribution (Kaplan and Violante, 2022). One way to see this "Missing Middle" problem is that the median wealth to mean annual earnings ratio 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 rational 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 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% in Kaplan and Violante (2022).

Hence, even though we only target the amount of total wealth in the economy, our one-asset model endogenously produces a realistic wealth distribution.

 $^{^{36}}$ We label a household HtM if it holds 0 assets. This implies that all HtM households are "poor HtM" in the parlance of Kaplan et al. (2014).

4.4 Extensions and alternative approaches

We first show that our results are robust to: (1) accounting for pessimistic households, and (2) alternative microfoundations for ex-ante optimism. We then discuss the relationship of belief heterogeneity to discount factor heterogeneity.

4.4.1 Pessimistic households

We now add pessimistic households to our model to check whether this changes our results materially.

Specifically, we now split the group of rational households into 25% pessimistic households and 25% rational households. We set the degree of pessimism to the inverse of mild optimism $(\alpha_p = 1/\alpha_m)$. We 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 in this way actually 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 savings needs 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, qualitatively speaking. But these partial equilibrium effects are quantitatively modest, because the mass of rational households getting replaced by pessimists have very low HtM probability and MPC in our baseline model (as in the rational model). 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 pushing them to the borrowing constraint.

4.4.2 Other sources of ex-ante optimism

We now consider ex-ante optimism about other key components of household financial situation besides idiosyncratic productivity.

Qualitatively, the main takeaway from our baseline model generalizes to any component of future financial situation: more optimism leads to more consumption and less saving as savings decisions are driven by expected future cash-on-hand. Optimistic households are thus more likely to end up being hand-to-mouth, and to have higher MPCs even conditional on wealth and income, irrespective of the object of optimism.

Quantitatively, results might differ across different microfoundations. To explore this, we now assume that all households are fully rational with respect to their idiosyncratic productivity and instead analyze two other potential microfoundations for ex-post optimism. We start with ex-ante optimism about an aggregate variable. We model this concerning the future aggregate wage level, w_t , with optimists expecting a higher wage from t + 1 onwards, $E_t((1 + \alpha_g)w_{t+1})$. This optimism about the future wage level implies that the expectation bias does not grow with the horizon.³⁷ Our analysis here thus also serves as a conservative robustness check of our baseline model, in which ex-ante optimism about idiosyncratic productivity implies that ex-post forecast errors grow with the time horizon (e.g., "a better job probably leads to an even better job down the road").

Targeting our lower bound estimates of ex-post optimism, such that mild (strong) ex-ante optimists about the aggregate variable are 10 (15) percentage points more likely to make optimistic forecast errors about their financial situation than rational households, produces $\alpha_m = 0.075$ and $\alpha_s = 0.135$. These conservative amounts of ex-post optimism heterogeneity produce an average MPC of 7%, which is more than $2\times$ the rational model's but less than half of our baseline model's. If instead we target our average empirical estimates of ex-post optimism heterogeneity, the ex-ante aggregate optimism model predicts an average MPC of 16%, in the consensus range of empirical estimates.

Another potentially meaningful contributor to ex-ante optimism about one's financial situation is *expense neglect*: underestimating future expenses. 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.

To explore expense neglect as an alternative microfoundation for optimism, we add expense shocks to the model. The budget constraint now reads:

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

where $\Xi_t(e_t)$ are expense shocks that do not directly enter the household's utility and are potentially a function of its idiosyncratic productivity.³⁸ While expense shocks occur with probability λ_{Ξ} , households expect them to occur with $\alpha_g \lambda_{\Xi}$, with α_g again being group-specific and $\alpha_g = 1$ capturing rational households. $\alpha_g < 1$ captures expense neglect.

For simplicity, we assume that expense shocks occur in every period ($\lambda_{\Xi} = 1$) and take the value $\Xi_t(e_t) = \xi_1 e_t^{\xi_2}$.³⁹ Following the empirical estimates in Fulford and Low (2024), we target average expense shocks equal to 13% of income. We also assume that expense shocks

 $^{^{37}}$ If households were optimistic about wage *growth*, the expectation bias about future wage levels would grow with the horizon.

³⁸Given our focus on expense neglect, and standard practice in literature on expense shocks, we only consider shocks that increase expenses.

³⁹We abstract from very large (but infrequent) shocks because we abstract from borrowing and potential consumer default in our model. To offset these differences, we target a somewhat higher frequency of more routine shocks than found in Fulford and Low (2024).

increase in income but less than 1-for-1, consistent with the findings in Fulford and Low (2024) that lower-income households are somewhat more exposed to expense shocks. We achieve this by setting $\xi_1 = 0.14$ and $\xi_2 = 0.34$. If we then target our lower-bound empirical estimates of ex-post optimism, ex-ante expense neglect produces $\alpha_m = 0.5$ and $\alpha_s = 0.0$, with the latter implying that strong optimists fully neglect future bad expense shocks. The average MPC in this model increases to 9%, compared to 3% in its rational counterpart. If instead we target our average empirical estimates of ex-post optimism heterogeneity, the average MPC increases; for example, setting $\alpha_m = 0.15$ (and keeping $\alpha_s = 0.0$) yields an average MPC of 12%.

Overall, we conclude that alternative approaches to microfounding households' ex-post optimistic bias about their future financial situations yield similar inferences about how heterogeneity in optimism affects household behavior and the average MPC.

4.4.3 Relationship to discount factor heterogeneity

As illustrated by Krueger et al. (2016), Kaplan and Violante (2022), and Aguiar et al. (2024), 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. (2024) further suggest that behavioral biases could provide a microfoundation for the low β of some households. Yet our empirical evidence in Section 2 points towards optimism having a somewhat stronger connection to HtM status than patience.

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

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. \blacksquare

The intuition is that optimism affects expected marginal utility, which depends on the idiosyncratic 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.⁴⁰

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:

⁴⁰The "endogenous discount factor heterogeneity model" in Appendix C of Kaplan and Violante (2022) a version of the discount factor heterogeneity model in which the heterogeneity in discount factors is not permanent—could be observationally equivalent to our model. But this would require 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.

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 the "missing middle" discussed above. 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 with respect to the optimal public debt level, the two models can differ vastly in their normative implications.

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. (2024b)). This approach seeks to capture illiquid assets 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 now includes capital.

The household's budget constraint now reads:

$$c_t + \frac{b_t}{1+r_t} + k_t = b_{t-1} + (1+r_t^k)k_{t-1} + (1-\tau_t)w_t\bar{e}_g e_t n_t,$$
(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. (2024a) 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:

$$b_t, k_t \ge 0.$$

A representative firm operates a Cobb-Douglas production function using capital (K) and

labor (N) as input factors:

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

where χ denotes the capital share in production.

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

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

Calibration. We maintain the same values for each of the parameters that also appear in our baseline model (except for the discount factor, per standard practice). 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. (2024a), and the liquid asset-to-annual income ratio to 0.2 as in Kaplan and Violante (2022). Given that it is well known that two-asset models can jointly match average wealth and the average MPC (Kaplan and Violante (2022)), we use the per-period capital market participation probability λ and the discount factor β to jointly target the empirical average wealth-to-annual income ratio of 4.1 as in our baseline model (Kaplan and Violante, 2022) and its predicted average MPC of 0.18.

Belief heterogeneity and the liquidity premium. Table 7 Column 1 shows the key stationary equilibrium predictions of our two-asset model. We discuss the return gap and its relationship to the liquidity premium at the end of this sub-section, focusing for now on the predicted aggregate share of HtM households: 0.40. This closely approximates our empirically estimated population share of 37-40%, with the improvement over our baseline model's prediction of 0.23 driven by the same mechanism found in prior work: some "wealthy HtM" households who would not be HtM in a one-asset model now choose to save only in the illiquid asset, due to its higher return (Kaplan et al., 2014, 2018). Again the behavior of rational and optimistic households differs starkly, with only 13% of rational households being HtM as compared to 55% and 80% for mild and strong optimists (see Appendix Table A9).

	Two-asset HANK w/ belief Het.	Rational two-asset HAN			
	(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%		

 Table 7: MPC and return gap across two-asset models.

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, which is the lower end of the consensus range of empirical estimates.

We compare our model to two different calibrations of the rational two-asset HANK model. The first calibration leaves all other parameters the same as in our model (except for recalibrating β to target the mean wealth-to-annual income ratio of 4.1) and the MPC untargeted.⁴¹ Table 7 Column 2 shows that this specification predicts an HtM share of 0.25 and an average quarterly MPC of 0.07, both of which are substantially below our model's predictions and the lower ends of the consensus ranges of empirical estimates. The second calibration instead targets an MPC at the lower end of the consensus range: 0.15.⁴² Column (3) shows that the predicted HtM share remains on the low side, and that the return gap required to clear asset markets rises from 5.7% to of 9.3% (Column 3). This is the liquidity premium required to push the average MPC up from 0.07 to 0.15, by inducing households to hold illiquid assets rather than a liquid buffer stock (recall the well-known result that rational households strive to avoid hitting the borrowing constraint, and to save their way off it when they do). Note that targeting the same MPC as in our baseline model, 18%, would produce an even higher return gap than 9.3%.

Our model produces a much lower return gap of 2.6% because ex-ante optimistic households undervalue self-insurance needs and thus require a much smaller premium on illiquid assets, thereby driving demand for the illiquid asset up and its return down. It may seem at first glance that our two-asset model undershoots substantially, given empirical estimates of the return gap in the ballpark of 5% (see, e.g., Jordà et al. (2019)). But both our model and the rational versions of it abstract from aggregate risk, which Ilut et al. (2024) estimates to be a quantitatively more important contributor to the return gap than illiquidity. Accounting for aggregate risk would thus push our estimated return gap closer to the data and a rational HANK model's farther away from it.

 $^{^{41}\}text{Targeting}$ the same average wealth-to-annual income ratio requires quarterly β = 0.990, as compared to 0.993 in our model.

⁴²In targeting the quarterly average MPC of 0.15 we set $\beta = 0.9805$, $\lambda = 0.15$, and $\delta = 0.00875$.

5 Fiscal Policy Implications of Heterogeneity in Optimism

We now show that heterogeneity in optimism 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 and more empirically realistic MPC-income distribution than rational models, as optimism is a key predictor of HtM status even conditional on income. This has implications for the effectiveness of income-targeted transfer payments. Second, optimistic households are less responsive to changes in self-insurance 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 in steady-state (we consider increasing public debt as an example).

5.1 The distribution of MPCs and targeted transfers

We start by considering an unanticipated, income-targeted transfer policy. Such transfer policies have become a regular part of policymakers' stimulus toolkits in recent recessions. 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 (2025)). 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 one-asset baseline model (blue dashed line with circle points) and compares it to its rational counterparts. The rational, one-asset model (black solid line with square points) produces unrealistically low MPCs throughout the income distribution as rational households strongly tend to precautionary save to avoid hitting the borrowing constraint. As discussed above, a common patch for this poor fit to the data is targeting a more realistic average MPC, at the cost of the model then producing far too little wealth. Our analysis here highlights an additional drawback: compared to empirical estimates (e.g., Boehm et al. (2023)), the MPC in the low-wealth rational model decreases far too sharply with income (orange, dotted line with diamonds). The reason for this steep slope is that the model can only produce high MPCs for households with approximately zero liquidity, which dictates that basically all low-income households must be at the borrowing constraint to match the average MPC. But Fuster et al. (2021), to take another empirical example, shows that households above and below a middle-income threshold have similar average MPCs: 17% and 20%.

Our model produces a flatter and hence more empirically realistic MPC-income gradient. Most strikingly, it is almost flat along the first three income quartiles, consistent with the findings in Fuster et al. (2021). The average MPC only drops substantially in the highest income quartile (but still amounts to almost 10%). This pattern is a consequence of the previouslydiscussed model mechanisms: optimists are more likely to be HtM at any income level than their rational counterparts. Additionally, even conditional on the same asset holdings, optimists exhibit higher MPCs (see Figure 1 and our discussion in Section 4.1). As such, optimism adds another driver of high MPCs to the standard "bad luck" mechanism.

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 0.8,⁴³ 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 only about 30% as strong as 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 0.8.⁴⁴





Note: Panel (a) shows predicted average MPCs, out of an unanticipated \$500 transfer, by income quartile for three different models. "Belief Het." is our baseline model and "Rational" its rational counterpart. "Rational, low wealth" is the rational HANK model with wealth reduced 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.

A second channel further weakens the effectiveness of targeted transfers in our model: muted ⁴³Targeted-transfer stabilization programs typically keep eligibility windows open for a period of time and

we, following standard practice, approximate this with an AR(1)-process. ⁴⁴The low-MPC rational HANK model (solid black line) produces a multiplier of about 0.5, due to its low MPC across all income groups.

relaxation of the self-insurance motive. The persistence in transfers implies some insurance for higher-income households who are not currently eligible but could become so in the event of a future negative shock. This temporarily decreases self-insurance motives, and strongly so for rational households, further increasing spending and total output.⁴⁵ But optimistic households undervalue this insurance because they underestimate the likelihood of being income-eligible in the future. They thus *perceive* their self-insurance motive (which they perceive to be weak to begin with) to be less relaxed than rational households, and as such barely increase their spending through this channel.⁴⁶

5.2 Savings behavior and fiscal insurance policies

Accounting for the muted responsiveness of optimistic households to changes in self-insurance incentives is even more important when evaluating policies focused on insurance provision in steady-state. 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. (2024b), we model MIB as a transfer $tr_{i,t}$ to household *i* at time *t* contingent on the household's pre-tax labor income $w_t n_{i,t} e_{i,t}$ falling short of some threshold level:

$$tr_{i,t} = \max\{0, a_1\bar{y} - a_2w_tn_{i,t}e_{i,t}\},\$$

where \bar{y} is the median income in the stationary equilibrium and $0 \le a_1, a_2 \le 1$. Transfers thus decrease in individual income at the withdrawal rate a_2 and no transfers are paid to households whose labor income satisfies $w_t n_{i,t} e_{i,t} \ge \frac{a_1}{a_2} \bar{y}$. Following Bayer et al. (2024b), 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:

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

⁴⁵See e.g. Bayer et al. (2023)'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 Gehrke (2024), 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 have distortionary effects). Beraja and Zorzi (2024) analyze potential size-dependency for stimulus transfers.

⁴⁶The relaxation of the precautionary savings motive is also an important driver in the rational HANK model with low average MPCs (black-solid lines 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.

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

	HANK w/ Belief Het.		Rat. H	ANK	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%	
Bottom 50W	2.1%	1.5%	12.7%	9.2%	2.3%	0.9%	
Real rate	4%	4.9%	4%	5.5%	4%	6.9%	

 Table 8: Effects of introducing public insurance

Note: MPCs refer to MPCs out of a \$500 stimulus check. "HANK w/ Belief Het." is our baseline model (oneasset, 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. "10W" and "50W" refer to wealth distribution percentiles.

Table 8 shows the stationary equilibrium effects of MIB on the average MPC, the HtM share, and wealth shares across the three models considered previously: our baseline one-asset model (Columns 1a and 1b), its rational one-asset HANK model counterpart (Columns 2a and 2b), and the low-wealth HANK that targets 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 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 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 large increase in the equilibrium real interest rate, from 4% to 5.5% or 6.9%. These higher rates are required to induce enough saving, in aggregate, to hold total wealth constant even as self-insurance motives have been dampened by the policy. The wealth share effects show that these higher interest rates are more effective at getting wealthier households to save and thus further skew the wealth distribution. MIB increases the wealth share owned by the top 10% by 2pp or 4pp, and decreases the share held by the bottom 50% by 3.5pp or 1.4pp.

In our model, savings crowd-out is modest because optimistic 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 0.17,⁴⁷ while the share of HtM households increases

⁴⁷There are two opposing effects of the introduction of MIB on the average MPC: the effective lower income risk reduces households' MPC conditional on their individual state, but there are more households in individual states with higher MPCs due to crowd out of savings. In the rational models, the latter dominates whereas in our baseline model, they roughly cancel out because minimum income benefits only mildly crowd out households'

by 4pp. It follows that the real interest rate increase is substantially smaller than in the rational models, here rising only to 4.9%. The effect on wealth inequality is also muted, with the wealth share held by the bottom 50% only decreasing by 0.6pp and the top 10%'s share only increasing by 1pp. We infer that public insurance policies like MIB have smaller side effects when some households are 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 the self-insurance possibilities for households.

But this increase in liquidity supply has muted effects in our model compared to a 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.⁴⁸

The solid black lines in Figure 3 show that in our model's rational counterpart, liquidity provision is quite effective at driving down the HtM share and increasing the wealth share of the bottom 50%. 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, compared to nearly zero in the rational 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 rational 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 undervaluation of self-insurance motives.

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 each model, using a utilitarian social welfare function

savings.

⁴⁸When 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. For the rational model, we choose the discount factor that gives the same real interest rate as our baseline model at the wealth level of 4.1, as discussed in Section 4.2.



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 rational HANK model, and the blue-dashed lines show the case for our baseline HANK model with belief heterogeneity.

that seeks to maximize the average expected discounted lifetime utility of households.⁴⁹



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 line) 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 -1.

Figure 4 shows that average welfare peaks at a much lower debt level in our model compared to the rational one-asset HANK model: optimal debt is about 210% of annual GDP, compared to about 370% in the rational HANK model. Since optimistic households underestimate their

⁴⁹Such an objective function takes into account aggregate efficiency, risk-sharing, and intertemporal-sharing components (Dávila and Schaab, 2024). 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".

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 rational model. Even though we abstract from many 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.⁵⁰

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 higher than our baseline model's (and hence even higher than in the rational model without ex-ante heterogeneity) because the households who benefit more from government liquidity provision (those with higher discount factors, because they value savings more) also get de facto higher social welfare weights in a utilitarian welfare function (because their future utility is discounted less).

As such, our accounting for the strong theoretical and empirical relationships between heterogeneity in ex-ante optimism, savings behavior, and HtM status can matter greatly for optimal policy.

6 Conclusion

We provide evidence suggesting that persistent heterogeneity in households' beliefs about their own future financial situation is an important driver of heterogeneity in savings behavior and in households' financial situations. Using a novel U.S. micro data set, we identify consumers with an ex-ante, trait-like optimistic bias about their own future financial situation and find that these households are substantially less likely to save and more likely to be financially constrained.

Guided by these findings, we then introduce persistent belief heterogeneity into a HANK model and uncover a new explanation for why many households are persistently HtM: optimistic households undervalue the need for buffer stock savings and thus, become and stay HtM not simply or even primarily due to "bad luck". Rather, they *choose* to consume instead of save in

⁵⁰In 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 crowd out of productive capital. See e.g., Woodford (1990), Aiyagari and McGrattan (1998), Davila et al. (2012), and Angeletos et al. (2023) for analyses of optimal public liquidity provision in rational models and Auclert et al. (2024a)'s characterization of long-run optimal fiscal policy including the transition in rational models.

anticipation of a better future.

Accounting for this additional mechanism—"beyond bad luck"—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 more empirically realistic, liquidity premium than required in other models. Overall, heterogeneity in optimism substantially improves the empirical fit of existing HANK models while being 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 self-insurance 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 accounting more completely for the rich set of mechanisms through which optimism could affect welfare. Our finding that optimism correlates strongly with persistent and severe financial distress suggests that there could be important productive and utility costs—costs that might be amplified by financial stress (Sergeyev et al., 2024). But other literatures suggest that optimism may not be all bad. Some work on investment suggests that optimism can produce productive benefits by countering other biases or frictions against risktaking. Work on motivated beliefs highlights how people can get some direct utility from their optimism (e.g., Brunnermeier and Parker (2005)).

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, may require accounting for heterogeneity in the relationship between beliefs about macroeconomic variables and about one's own financial situation.

References

- AGUIAR, M., M. BILS, AND C. BOAR (2024): "Who are the Hand-to-Mouth?" *Review of Economic Studies*, rdae056.
- 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*.

— (2025): "Fiscal and Monetary Policy with Heterogeneous Agents," Working paper.

- BALLEER, A., G. DUERNECKER, S. FORSTNER, AND J. GOENSCH (2024): "The effects of biased labor market expectations on consumption, wealth inequality, and welfare," *Working Paper*.
- 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 (2024a): "Shocks, frictions, and inequality in US business cycles," *American Economic Review*, 114, 1211–1247.
- BAYER, C., B. BORN, R. LUETTICKE, AND G. J. MÜLLER (2023): "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 (2024b): "A HANK2 model of monetary unions," *Journal of Monetary Economics*, 147.
- BENIGNO, P. AND L. PACIELLO (2014): "Monetary policy, doubts and asset prices," *Journal* of Monetary Economics, 64, 85–98.

⁽²⁰²⁴b): "The intertemporal Keynesian cross," *Journal of Political Economy*, 132, 3919–4234.

- 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.
- BROER, T., A. KOHLHAS, K. MITMAN, AND K. SCHLAFMANN (2022): "Information and Wealth Heterogeneity in the Macroeconomy," *Working Paper*.
- BROER, T. AND A. N. KOHLHAS (2024): "Forecaster (mis-) behavior," *Review of Economics* and Statistics, 106, 1334–1351.
- BRUNNERMEIER, M. K. AND J. A. PARKER (2005): "Optimal expectations," American Economic Review, 95, 1092–1118.
- 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 (forthcoming): "Evidence on expectations of household finances," *Management Science*.

- DÁVILA, E. AND A. SCHAAB (2024): "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 (2024): "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 impatience related to cognitive ability?" *American Economic Review*, 100, 1238–1260.
- D'ACUNTO, F., D. HOANG, AND M. WEBER (2020): "Managing households' expectations with unconventional policies," Tech. rep., National Bureau of Economic Research.
- D'ACUNTO, F. AND M. WEBER (2024): "Why survey-based subjective expectations are meaningful and important," Annual Review of Economics, 16.
- D'ACUNTO, F., M. WEBER, AND X. YIN (2024): "Subjective Income Expectations and Household Debt Choices," .
- 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? Spending responses to gains, losses, news, and loans," *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, Ô., 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.
- KAPLAN, G., G. L. VIOLANTE, AND J. WEIDNER (2014): "The wealthy hand-to-mouth," Brookings Papers on Economic Activity.
- 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 (2025): "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.
- KRUSELL, P. AND A. A. SMITH, JR (1998): "Income and wealth heterogeneity in the macroeconomy," *Journal of Political Economy*, 106, 867–896.
- LAIBSON, D., P. MAXTED, AND B. MOLL (2025): "Present bias amplifies the household balance-sheet channels of macroeconomic policy," *The Quarterly Journal of Economics*, 140, 691–743.
- 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.
- MALMENDIER, U. AND T. TAYLOR (2015): "On the verges of overconfidence," Journal of Economic Perspectives, 29, 3–8.
- 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.
- MOORE, D. A. AND P. J. HEALY (2008): "The trouble with overconfidence." *Psychological* review, 115, 502.
- MUELLER, A. I., J. SPINNEWIJN, AND G. TOPA (2021): "Job seekers' perceptions and employment prospects: Heterogeneity, duration dependence, and bias," *American Economic Review*, 111, 324–363.
- NIELSSON, U., J. RANGVID, F. SAIDI, F. SEYRICH, AND D. STREITZ (2025): "Empirical Monetary-Fiscal Equivalence," *Working Paper*.
- 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 (2024): "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 Working 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. (2025): "Interest Rate Cuts vs. Stimulus Payments: An Equivalence Result," *Journal* of Political Economy, 133, 1235–1275.
- 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 Macroe*conomics Annual, 33, 1–50.

Online Appendix Beyond bad luck: Macroeconomic implications of persistent heterogeneity in optimism

Oliver Pfäuti Fabian Seyrich Jonathan 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

	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)			
Ν	$15,\!047$	$15,\!047$	$15,\!049$	$15,\!049$			
N panelists	3057	3057	3056	3056			

Table A1: Financial situation forecasts are positively correlated with income forecasts

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".

Panel A1	nel A1. Sample with measurable potentially optimistic FCEs, v1								
	Rea	alization at t	+1	Panelist-level FCE counts					
$\underline{FC} \text{ at } t$	Better	Same	Worse						N with
Better	Accurate	Optimist	Optimist	$N \ {\rm FCEs}$	Mean	SD	Min	Max	$\geq 2~{\rm FCEs}$
Same	Pessimist	Accurate	Optimist	18,899	5.5	3.4	0	12	2,928
Worse	Pessimist	Pessimist	Accurate						
D1 49	Q 1	41 1. 1	4 4 11			El-	J. "D.		
Panel A2	. Sample wi Rea	alization at t	+1	optimistic FC	Paneli	st-leve	$\frac{1}{2}$ el FCE	counts	etter pairs
FC at t	Better	Samo	Worse						N with
$\frac{\Gamma \odot at t}{Better}$	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 wit	th measurabl	e potentially	pessimistic F	$\frac{\text{CEs, v1}}{\text{D}}$. 1	1 000		
	Rea	alization at t	+ 1		Paneli	st-leve	el FCE	counts	
$\underline{FC} \text{ at } t$	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 wit	th measurabl	e notentially	nessimistic F	CEs v2·	Evel	ude "W	Vorse-W	orse" pairs
<u> </u>	Rea	alization at t	+1	Panelist-level FCE counts					
FC at t	Better	Same	Worse						N with
$\frac{1 \odot \text{ at } i}{\text{Better}}$	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	<u>measurable</u>	potentially sy	metric FC.	Es D	. 1	1 000		
	Rea	alization at t	+ 1		Paneli	st-leve	ei fCE	counts	
$\underline{FC} \text{ at } t$	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						

Table A2: Measuring ex-post forecast errors (FCEs) with coarse data: Sample definitions

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".

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	Re	alizatio	n this ye	ar		
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	Re	alizatio	n this ye	ar		
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	Re	alization	n this ye	ar		
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		

Table A3: Household financial condition forecasts and forecast errors tilt optimistic

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.

Sample	Potentially opt or pess FCEs			Potentially symmetric		No adjustment for		
Sampio	Al	1	v2		FCEs		measurement error	
Estimate	Unwtd	Wtd	Unwtd	Wtd	Unwtd	Wtd	Unwtd	Wtd
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A. Prop of forecast errors that are optimistic	0.33	0.33	0.39	0.38	0.33	0.32	0.28	0.28
Share consumers with optimistic proportion ≥ 0.5	0.32	0.32	0.40	0.39	0.32	0.32	0.24	0.25
N consumers (≥ 2)	2928	2928	2792	2792	2787	2787	3073	3073
Panel B. Prop of forecast errors that are pessimistic	0.18	0.17	0.20	0.19	0.15	0.15	0.12	0.11
Share consumers with pessimistic proportion ≥ 0.5	0.14	0.14	0.17	0.17	0.11	0.11	0.05	0.05
N consumers (≥ 2)	2542	2542	2454	2454	2787	2787	3073	3073
Panel C. Relative prop or shares: Optimism/Pessimism								
Proportions	1.83	1.94	1.95	2.00	2.20	2.13	2.33	2.55
Shares > 0.5	2.29	2.29	2.35	2.29	2.91	2.91	4.80	5.00

 Table A4: Household financial condition FCE proportions (Columns 1-6 are same as Table 2)

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.

errors	
	Panel A. Unweighted estimates

Table A5: Financial situation forecast errors are persistent: Comparing consecutive forecast

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.

LHS = Pessimism FCE proportion							
Camala	Pot	tentially	Pessimist	Potentially Symmetric			
Sample	v1		v2				
Estimate	Unwtd	Wtd	Unwtd	Wtd	Unwtd	Wtd	
	(1)	(2)	(3)	(4)	(5)	(6)	
Overconfidence	-0.04	-0.04	-0.07	-0.05	-0.09	-0.22	
	(0.09)	(0.15)	(0.10)	(0.18)	(0.07)	(0.18)	
Controls?	Yes	Yes	Yes	Yes	Yes	Yes	
Ν	666	666	638	638	742	742	
N Panelists	333	333	319	319	371	371	
Mean~(LHS)	0.19	0.18	0.22	0.20	0.16	0.15	

Table A6: Overconfidence and its correlation with ex-post pessimism

Note: Standard errors, clustered on panelist, in parentheses. One regression per panel-column. RHS variables measured using the Stango-Zinman data and merged onto our household financial situation forecast error panel. The results shown represent $100 \times$ the change in the LHS variable associated with a 1 percentage point increase in the overconfidence rank. "Controls" include income, standard demographics (education age, gender, race, and ethnicity), patience, two measures of risk aversion, 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:
 Overconfidence strongly conditionally correlates with both smooth and indicator measures of ex-post optimism

Panel A. LHS = Optimism FCE proportion (same as Table 3)						
Commle	Potentially Optimistic				Potentially Symmetric	
Sample	v1		v2			
Estimate	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)
Controls?	Yes	Yes	Yes	Yes	Yes	Yes
Ν	778	778	750	750	742	742
N Panelists	389	389	375	375	371	371
Mean (LHS)	0.36	0.32	0.44	0.42	0.38	0.37
Panel B. LHS	= 1 = (op	timistic	proporti	$\mathrm{ion} \ge 0.5$	5)	
	Pe	otentially	Optimisti	ic	Potentially Sy	$\operatorname{mmetric}$
	v1		v2			
	Unwtd	Wtd	Unwtd	Wtd	Unwtd	Wtd
	(1)	(2)	(3)	(4)	(5)	(6)
Overconfidence	0.42	0.56	0.48	1.20	0.41	1.09
	(0.18)	(0.35)	(0.20)	(0.56)	(0.19)	(0.53)
Controls?	Yes	Yes	Yes	Yes	Yes	Yes
Ν	778	778	750	759	742	742
N Panelists	389	389	375	375	371	371
Mean (LHS)	0.36	0.32	0.45	0.44	0.37	0.37

Note: Panel A is the same as Table 3 in the main text. 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), two measures of risk aversion, 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 .

A.1 Two-Asset Model: Additional Tables

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

Table A8: Calibration of our two-asset model

Note: 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.

	Two-asset HANK w/ belief Het.	Rational two-as	sset 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	-	-	

Table A9: MPC and return gap across two-asset models.

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.

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:

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

where " $\hat{}$ " 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

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

$$\frac{\beta}{\hat{\beta}} = \frac{p_{DU}u'(c_{t+1}^{2,U}) + (1 - p_{DU})u'(c_{t+1}^{2,D})}{\alpha p_{DU}u'(c_{t+1}^{2,U}) + (1 - \alpha p_{DU})u'(c_{t+1}^{2,D})}.$$
(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.⁵¹ 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.

⁵¹A 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.





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 Varying the Shares of Optimists

In our baseline model, we assume that (a) 50% of households are optimists (Huffman et al. (2022)) and (b) we divide our optimists into two groups, mild and strong optimists. We now check the robustness of these two assumptions.

We start by assuming different shares of optimists in total, keeping the half-half split into mild and strong optimists among optimists. To this end, we first assume that only 33% of the population exhibit an optimistic bias while 67% are rational. We first keep the degree the degree of the bias as in our baseline, that is $\alpha_m = 1.6$ and $\alpha_s = 2.1$, to isolate the effects of lower shares of optimists. This results in an average MPC of 14%. If we instead re-calibrate $\alpha_m = 1.65$ and $\alpha_s = 1.9$ to match an increase in the likelihood to make an optimistic forecast error of 12pp and 15pp—the respective lower bound estimates of the increase in likelihood of ex-post optimism of the median households in the two optimism groups—the average MPC is 12%. Bottom line, a lower share of optimists does not change our qualitative results, but quantitatively, our model predicts a lower average MPC compared to our baseline. Yet, if we set $\alpha_m = 2.1$ and $\alpha_s = 3.0$ to match an increase in the likelihood to make an optimistic forecast error of 16pp and 22pp (roughly our middle point estimates), the average MPCs are 23% which are again in the range of empirical estimates.

In contrast, if we calibrate the model to a higher share of optimists, the model predicts a higher average MPC. If we increase the share of optimists to 60% and leave the degree of the optimistic bias the same as in our baseline, the model predicts an average MPC of 20%.

Lastly, we keep the share of optimists constant at 50% as in our baseline but assume that all optimistic households share the same degree of ex-ante optimism, that is $\alpha_s = \alpha_m = 1.85$. This two-group calibration implies an average MPC of 18% exactly as in our baseline two group model. Even if we set $\alpha_s = \alpha_m = 1.6$, that is, if we simply group the strict optimists to become mild optimists, the average MPC would still be 14%. We conclude that our results are robust, both qualitatively and quantitatively with respect to different assumptions about the share of optimists and their grouping.

D.2 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.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]:

$$\begin{bmatrix} \rho_B & (1-\rho_B)/2 & (1-\rho_B)/2 \\ (1-\rho_B) & \rho_B & 0 \\ (1-\rho_B)/2 & 0 & \rho_B. \end{bmatrix}$$
(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).⁵² 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, at the lower bound of 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:

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

$$(21)$$

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.

 $^{^{52}}$ Stango and Zinman (2023, 2020) estimate 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.

We again calibrate $\alpha_m = 1.5$ and $\alpha_s = 3.5$ matching the slope in the increase in the likelihood to make 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 Household behavior with risk heterogeneity



Figure D4: Household behavior with rational optimism and risk heterogeneity

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 to match the degree of ex-post optimism that we find in our micro data, 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.