

Bad luck or bad decisions?

Macroeconomic implications of persistent heterogeneity in cognitive skills and overconfidence

Oliver Pfäuti*

Fabian Seyrich†

Jonathan Zinman‡

March 2024

[Link to most recent version](#)

Abstract

Business cycle models often abstract from persistent household heterogeneity, despite its potentially significant implications for macroeconomic fluctuations and policy. We show empirically that the likelihood of being persistently financially constrained decreases with cognitive skills and increases with overconfidence thereon. Guided by this and other micro evidence, we add persistent heterogeneity in cognitive skills and overconfidence to an otherwise standard HANK model. Overconfidence proves to be the key innovation, driving households to spend instead of precautionary save and producing empirically realistic wealth distributions and hand-to-mouth shares and MPCs across the income distribution. We highlight implications for various fiscal policies.

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

Keywords: Household Heterogeneity; Cognitive Skills; Overconfidence; Financial Constraints; Fiscal Policy; HANK

*The University of Texas at Austin, Department of Economics, pfauti.oliver@gmail.com

†Berlin School of Economics, DIW Berlin, and Freie Universität Berlin, fabian.seyrich@gmail.com

‡Dartmouth College, Department of Economics, jzinman@dartmouth.edu.

An early draft of this paper was titled "Heterogeneity in what? Cognitive Skills, Beliefs and the Liquid Wealth Distribution". We thank Eduardo Dávila, Greg Kaplan, Ralph Luetticke (discussant), Kurt Mitman, Peter Maxted, Matthew Rognlie, Johannes Stroebel, Hannes Tzieling, Gianluca Violante, Michael Weber, Nathan Zorzi, and seminar and conference participants at the Bank of Finland, the 13th ifo Conference on Macroeconomics and Survey Data, EEA/ESEM 2023, the University of Zurich, the University of Virginia, the Bonn-Berlin PhD Workshop, the Swiss Economists Abroad end-of-year conference, the HeiTueHo Workshop, EWMES 2023, and the 5th Behavioral Macroeconomics Workshop for helpful comments and suggestions. Pfäuti gratefully acknowledges financial support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 (Project C02) and the Stiftung Geld & Währung. Seyrich gratefully acknowledges financial support by the Leibniz Association through the project "Distributional effects of macroeconomic policies in Europe". Zinman thanks the Nelson A. Rockefeller Center for Public Policy at Dartmouth for funding.

1 Introduction

Heterogeneity in households’ savings behavior and financial situations has significant implications for macroeconomic fluctuations and policy design.¹ Yet it remains standard practice in macro modeling to assume ex-ante identical households and account for heterogeneity only in shock realizations: Households are wealthy or poor only because of good luck or bad luck, abstracting from choices linked to fundamental and persistent dimensions of heterogeneity across households.²

One important dimension of fundamental heterogeneity is cognitive skills, which has been linked empirically to: differences in economic growth across space and time (Hanushek and Woessmann (2008)), households’ inflation expectations (D’Acunto et al. (2019, 2023b)), responses to changes in incentives (D’Acunto et al. (2023a)), financial mistakes (Agarwal and Mazumder (2013)), and strong negative relationships between behavioral biases and income (e.g., Stango and Zinman (2023), Chapman et al. (2023)). Links between cognitive skills and savings behavior are less well understood, particularly for the sorts of behaviors and outcomes featured in macro modeling (e.g., hand-to-mouth a.k.a. HtM status). And the macroeconomic implications of any such link between cognitive skills heterogeneity and HtM heterogeneity are largely unexplored.

We start by using micro data, from a nationally representative sample of working-age U.S. consumers, to develop several new facts about how cognitive skills, beliefs, and household financial situations are related. We find that the likelihood of being persistently HtM, measured in various ways, decreases sharply with cognitive skills. But allowing for cognitive skills heterogeneity alone is unlikely to help macro models fit the data, as we later formalize, because permanently low-productivity households will still tend to save their way out of HtM status if they are classically rational. This motivates considering beliefs as well, starting with how consumers perceive their own cognitive skills.

We show that persistent overestimation of one’s own skills is prevalent (as in the high-stakes setting of Huffman et al. (2022)) and differs across consumers: overconfidence correlates strongly and negatively with cognitive skills. Overconfident consumers are also about 1.2 times as likely than their well-calibrated counterparts to be persistently overly-optimistic about their future financial situations (measured using standard consumer sentiment forecasts and realizations), suggesting that lower-skilled consumers may be HtM at least in part because of their overconfidence. Consistent with this conjecture, we find strong correlations between persistent overconfidence and our measures of persistent HtM status.

Guided by these micro findings, we add persistent heterogeneity in cognitive skills and perceptions thereof to an otherwise standard heterogeneous agent New Keynesian (HANK) model with

¹See, e.g., Werning (2015), Kaplan et al. (2018), Auclert (2019), Bayer et al. (2023a), Luetticke (2021), Hagedorn et al. (2019), Patterson (2023), Almgren et al. (2022), Holm et al. (2021) on shock transmission and policy efficacy, and Dávila and Schaab (2023a), McKay and Wolf (2023), Bhandari et al. (2021), Bilbiie (forthcoming), Smirnov (2022), Acharya et al. (2023), Yang (2022) on optimal policy design.

²Important exceptions include models allowing for heterogeneity in preferences (e.g., Auclert et al. (2020), Aguiar et al. (forthcoming), Kaplan and Violante (2022)), as we discuss below.

incomplete markets, idiosyncratic productivity risk, borrowing constraints, and a nominal rigidity in the form of sticky wages. This framework allows us to unpack potential mechanisms underlying our empirical findings, and to derive macroeconomic implications of heterogeneity in cognitive skills and overconfidence. We model cognitive skills heterogeneity as differences in average labor market productivity and overconfidence as overweighting the probability of reaching a better productivity state and underweighting the probability of reaching a worse state. Motivated by micro data on the prevalence of persistent overconfidence and the strong correlation between cognitive skills and overconfidence, and in the interest of parsimony, we calibrate our baseline model such that 62% of households are high-skilled with well-calibrated beliefs about future productivity while the remaining 38% are low-skilled and overconfident. We also calibrate the parameter governing the degree of overconfidence, by matching our finding that overconfident households are about 1.2 times as likely to be overly-optimistic about their future financial situations than rational households.

Accounting for heterogeneity in cognitive skills and overconfidence substantially improves the model's empirical fit. In contrast to standard one-asset HANK models and to a HANK model with heterogeneity in skills but not in beliefs about them, our model jointly matches total wealth in the economy, high HtM prevalence, and an average quarterly marginal propensity to consume (MPC) in the consensus range of 15-25% (e.g., [Jappelli and Pistaferri \(2010\)](#), [Havranek and Sokolova \(2020\)](#)). This holds even when all wealth is liquid and held in a single asset.³

Existing one-asset HA(NK) models struggle to match these data moments jointly because if the supply of assets is large enough to match the average wealth in the economy, the price of the asset is so low that almost all households accumulate a sufficient buffer stock to make the borrowing constraint nonbinding ([Auclert et al. \(forthcoming\)](#), [Kaplan and Violante \(2022\)](#)). This makes HtM status counterfactually rare and implies that most households have low MPCs. Consequently, standard models produce an average MPC that is too low.

Our model achieves reconciliation because overconfident households underestimate their insurance needs and consequently perceive the price of the asset as too high to merit accumulating a sufficient buffer stock. Even when the supply of assets is high, many overconfident households choose to do little if any precautionary saving and often end up being HtM, consistent with our empirical findings. HtM status in our model thus is often due at least partly to "bad decisions" and not only to the "bad luck" that drives standard models. Our results are driven by differences in overconfidence rather than by differences in skills: removing heterogeneity in overconfidence from the model by imposing rational beliefs for all households, while retaining heterogeneous average

³Our model also accounts well for other untargeted wealth inequality statistics. It produces more and more empirically realistic inequality than its rational counterpart, better matching empirical wealth shares—e.g., of the top 10% or the bottom 50%. Moreover, our model does not suffer from the "missing middle" problem ([Kaplan and Violante, 2022](#)) of an implied wealth distribution that is too polarized compared to the data. E.g. a standard one-asset HANK model predicts median wealth that is about an order of magnitude smaller than the data's. Our model matches this (untargeted) moment well.

skill levels, fails to match the average MPC and delivers very few HtM households.

A standard practice for better reconciling HANK models with the data is to introduce a second, illiquid asset that can be adjusted only infrequently (Kaplan and Violante (2014), Kaplan et al. (2018), Bayer et al. (2019), Auclert et al. (forthcoming)). This approach produces a liquidity premium that is arguably too high, as discussed in Kaplan and Violante (2022). We show that a two-asset version of our model can fit the data with a substantially lower liquidity premium, because overconfident households underestimating their individual income risk implies that they also underestimate the shadow value of future liquidity and thereby put downward pressure on the equilibrium liquidity premium.

In contrast to standard models, our model also generates empirically realistic shares of HtM households throughout the income distribution, even though we do not explicitly target this. Because overconfidence is a key predictor of HtM status, our model produces significant shares of higher-income HtM households, in line with the data. Standard models produce either far too few HtM households throughout the income distribution (when calibrated to match average wealth) or far too much HtM polarization by income (when directly targeting the average MPC). The reason is that standard models match the average MPC by making practically all low-income households HtM—a side effect of households being HtM due only to "bad luck".

Our model thus requires only one additional parameter—and no free parameters, as we discipline overconfidence using our new survey evidence—to substantially improve the performance of existing HANK models. The mechanism that allows us to better match key features of micro and macro data—lower-skilled households' undersaving due to overconfidence about their future financial situations—generates important and distinct implications for macroeconomic policies as well.

We start our policy analysis by considering unexpected transfer payments intended to stimulate private consumption by targeting households with high MPCs. Given the difficulty of empirically identifying the general-equilibrium effects of transfer policies, they are usually evaluated using models. For untargeted transfers, these models sensibly match the average MPC (e.g., Kaplan and Violante (2014), Wolf (forthcoming)). But with targeted transfers, average MPCs are no longer sufficient to assess the aggregate effects: it is the *distribution* of MPCs, across targeted vs. non-targeted groups, that matters. We illustrate this by modeling a stimulus payment targeted to the bottom income quartile and estimate a transfer multiplier of 0.9 in general equilibrium. This contrasts with standard HANK models, which either under- or overestimate the share of HtM households in low-income groups and thus under- or overestimate the average MPC of transfer recipients: the low-MPC standard HANK model predicts a multiplier of 0.5 and its low-wealth, high-MPC counterpart a multiplier of 2.4.

Next we show that heterogeneity in overconfidence also has implications for fiscal policies that more directly impact household self-insurance decisions in steady-state. The key mechanism is that our financially constrained households are mostly overconfident and hence value additional

insurance less than the constrained and rational households in standard models.

First, we show that providing public insurance through minimum income benefits does not crowd out private precautionary savings as strongly as predicted by rational models. Overconfident households undervalue this insurance because they underestimate their probability of reaching bad income states and therefore reduce any existing buffer stock only mildly. The introduction of 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 standard models.

Second, we consider indirect insurance provision through government debt issuance (e.g., [Woodford \(1990\)](#), [Aiyagari and McGrattan \(1998\)](#)). Higher government debt levels reduce households' self-insurance cost by reducing the cost of liquid assets. But the induced increase in precautionary saving is muted in our model because overconfident households undervalue the insurance function of cheaper assets. Thus even at high public debt levels, many overconfident 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. In a standard model, low-wealth households are eager to save themselves away from the borrowing constraint and increase their saving strongly in response to cheaper liquidity. This drives down the HtM share strongly and increases the wealth share held by the bottom half of the distribution. These contrasting effects have normative implications as well: the optimal government debt level is substantially lower with heterogeneous overconfidence, irrespective of whether we consider a model in which households can only save in government bonds or also in productive capital.

Related literature. We contribute to five strands of literature. One considers how cognitive skills heterogeneity affects the macroeconomy. So far, this literature is largely empirical and focused on growth ([Hanushek and Woessmann \(2008\)](#)). [D'Acunto et al. \(2019, 2023a,b\)](#) bring cognitive skills heterogeneity to the empirical study of economic fluctuations, showing it plays key roles in how households form their inflation expectations and respond (or not) to information and incentives provided by policy interventions. We empirically link heterogeneity in cognitive skills to heterogeneity in forecasted and realized financial situations, including HtM status. We then build a model capturing key features of that micro heterogeneity and use it to quantitatively study macro dynamics, the wealth distribution, and policy design and effectiveness. Altogether, we show that accounting for heterogeneity in beliefs about cognitive skills can greatly improve a model's ability to fit the micro and macro data.

A second strand considers potential psychological sources of liquidity or poverty traps and their macro implications. Work on aspirations as reference points ([Dalton et al., 2016](#); [Genicot and Ray, 2017, 2020](#)) has focused on how excessive pessimism can dampen growth, while we focus on how excessive optimism affects stabilization and macroeconomic policies. [Sergeyev et al. \(2024\)](#) consider how financial stress and naivete about financial stress can create persistent financial constraints and impact wealth inequality and fiscal multipliers. We consider a different decision

making mechanism than work on aspirations or stress, focusing on biased beliefs rather than behavioral preferences or the neglect of one or more key parameters. Our mechanism is relatively easy to validate empirically and incorporate into an otherwise standard quantitative model.

A third strand focuses on differences between perceived vs. actual idiosyncratic labor market risk. So far, this literature has focused on various beliefs about the labor market and a subset of important macro applications. [Balleer et al. \(2022\)](#) show that working-age individuals in the U.S. are "vastly over-optimistic about their own labor market prospects" (p. 1). [Mueller et al. \(2021\)](#) find optimistic bias about job-finding rates, especially for the long-term unemployed, and little evidence for downward revision of these beliefs when remaining unemployed. [Wang \(2023\)](#) shows how calibrating a standard incomplete-markets model to consumers' perceived rather than actual income risk is better able to account for observed wealth inequality.⁴ Our contributions are uncovering the role of cognitive skills heterogeneity in shaping biased perceptions about risk and future financial situations, and building a general equilibrium model that can jointly fit key features of micro and macro data and quantitatively evaluate and guide policy.

Fourth, we contribute to the development of macro models seeking to use insights from behavioral economics to improve predictive and prescriptive power. Most work in this vein focuses on a representative behavioral agent.⁵ Behavioral HANK models tend to allow for heterogeneity only in the budget constraint, with a homogeneous behavioral or information friction about an aggregate variable only.⁶ [Pfäuti and Seyrich \(2023\)](#) study a case of heterogeneous behavioral biases, but focus on expectations about aggregate variables in that case. [Guerreiro \(2023\)](#) allows for heterogeneous attention, but focuses on a case where households hold rational expectations about their idiosyncratic shocks. [Ilut and Valchev \(2023\)](#) develop a model of imperfect reasoning and introduce this into an [Aiyagari \(1994\)](#) economy. In contrast to our framework, their households are ex-ante identical and so HtM status is driven by adverse idiosyncratic productivity shocks—by bad luck.⁷

⁴The evidence on income forecast errors is more mixed. [Souleles \(2004\)](#) finds evidence of over-optimism in the 1986-1995 Michigan Survey of Consumers (SOC) (see especially his Figure 4 Panel A), using its short panel component to pair one 12-month forecast with a 6-month realization for some respondents. [Rozsypal and Schlafmann \(2023\)](#), using six additional years of SOC data, find that the direction and magnitude of forecast errors vary with income level. [d'Haultfoeuille et al. \(2021\)](#) find that "individuals tend to be right on average about their future earnings", using four-month forecasts and subsequent realizations in the first three waves of the New York Fed's Survey of Consumer Expectations (from 2015). They nevertheless strongly reject rational expectations after accounting for measurement error and aggregate shocks. [Caplin et al. \(2024\)](#) find close alignment between survey income forecasts and administrative data realizations in Denmark.

⁵See, e.g., [Woodford \(2013\)](#), [Gabaix \(2014\)](#), [Woodford \(2019\)](#), [Gabaix \(2020\)](#), [Bordalo et al. \(2020\)](#), [Lian \(2023\)](#), and [Boutros \(2023\)](#).

⁶See, e.g., [Farhi and Werning \(2019\)](#), [Auclert et al. \(2020\)](#), [Angeletos and Huo \(2021\)](#), [Laibson et al. \(forthcoming\)](#), and [Pfäuti and Seyrich \(2023\)](#).

⁷[Ilut and Valchev \(2023\)](#)'s households do not know their optimal policy function and estimate it based on costly (and noisy) deliberation signals. Once households become HtM, they are likely to remain so because they hold excessively high beliefs about their optimal consumption that induce them to dissave and remain at the borrowing constraint. In contrast, HtM households in our setup tend to differ systematically from households away from the borrowing constraint, consistent with what we find in the micro data. Additionally, our model features nominal rigidities and allows for two assets. We also take a step beyond the crucial one of matching key empirical moments

A fifth and parallel strand considers (persistent) heterogeneity in reduced-form or presumed-classical preferences. [Aguiar et al. \(forthcoming\)](#) find that allowing for heterogeneity in patience and the elasticity of intertemporal substitution helps match several empirical facts about the behavior of HtM households. They suggest that behavioral factors might provide a potential micro-foundation for their modeling choices. [Krueger et al. \(2016\)](#) and [Auclert et al. \(2020\)](#) introduce permanent heterogeneity in patience and—in the case of [Auclert et al. \(2020\)](#)—in average skills to better match wealth inequality data. [Kekre and Lenel \(2022\)](#) show that heterogeneity in risk aversion can help account for observed heterogeneity in portfolio choice. [Kaplan and Violante \(2022\)](#) show that heterogeneity in risk aversion can produce similar results to heterogeneity in discount factors in terms of HtM shares and MPCs. They also show, however, that allowing for heterogeneity in risk aversion or in discount factors does not solve the standard HANK’s “missing middle problem” of producing a wealth distribution that is too polarized. We show that allowing for heterogeneity in overconfidence, in contrast, fills in the missing middle. Furthermore, our micro data does not favor patience or risk aversion alone as an empirically likely key margin of heterogeneity; e.g., their correlations with HtM status are relatively weak compared to cognitive skills and overconfidence, both qualitatively and quantitatively.

Overall, we show that accounting for observed *fundamental* differences between financially constrained and unconstrained consumers is crucial for understanding macroeconomic fluctuations and general equilibrium. This 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 productivity shocks, and HtM tends to be a relatively transitory state. In our model, households are financially constrained in part because they overestimate their own abilities, leading to a systematic relationship between cognitive skills, overconfidence, and persistent HtM status. Accounting for this relationship turns out to matter greatly for policy as well.

Outline. We detail our data and empirical findings in [Section 2](#). [Section 3](#) shows how we introduce cognitive skills and overconfidence 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

In this section, we document several new facts regarding consumers’ cognitive skills, beliefs about these skills and future financial situations, and how they relate to other forms of persistent heterogeneity and to six measures of hand-to-mouth status. We later use these facts to help discipline and test our model. We show both unweighted and sampling probability-weighted estimates, by demonstrating use cases for our model: analyzing positive and normative implications for fiscal policy.

following [Solon et al. \(2015\)](#).

2.1 Data

Our micro data source is the American Life Panel, a long-running online panel that goes to great lengths to obtain a nationally representative sample of U.S. adults.

We measure cognitive skills and overconfidence about cognitive skills using data from the modules in [Stango and Zinman \(2023, 2024\)](#), henceforth SZ, which elicited behavioral biases and cognitive abilities, together with questions about household financial condition (that we use here to construct some of our measures of HtM status), from the same 845 panelists in two survey rounds administered in 2014 and 2017. The SZ modules sample only working-age adults (aged 18-60 in 2014), which maps well into our model's focus on labor-market productivity. We bring in additional variables—regarding standard measures of HtM status not covered in the SZ modules, and standard measures of consumer sentiment that we use to measure subjective financial condition and expectations thereof—using various other ALP surveys administered from 2010 through 2022. We start by detailing our key variable definitions and prevalences, including comparisons to other work where applicable. We then describe the key micro empirical regularities that shape and discipline our model.

Cognitive skills. We measure cognitive skills for SZ panelists with standard tests for fluid intelligence ([McArdle et al. \(2007\)](#)), numeracy ([Banks and Oldfield \(2007\)](#)), cognitive control/executive function ([MacLeod \(1991\)](#), [Miyake and Friedman \(2012\)](#)), and crystallized intelligence in the form of financial literacy ([Lusardi and Mitchell \(2014\)](#)).⁸ We then extract a single common factor (a.k.a. "g" or generalized intelligence) to use as a summary statistic for cognitive skills, as is customary given that various cognitive skills measures are strongly related, both conceptually and empirically ([Jensen, 1998](#); [Stango and Zinman, 2023](#)).⁹

Overconfidence. We measure overconfidence for SZ panelists 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. Later in that survey they take a standard 15-question "number series" test of fluid intelligence ([McArdle et al. \(2007\)](#)).¹⁰ Respondents are overconfident on average, with 70 percent providing a better-than-average percentile.¹¹

⁸For details on test questions, please see the Data Appendix to [Stango and Zinman \(2023\)](#).

⁹Results are very similar, qualitatively and quantitatively, if we use the first principal component of cognitive skills instead of the first common factor.

¹⁰Number series scores correlate strongly with those from other fluid intelligence tests like IQ and Raven's.

¹¹The 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 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).

We are most interested in *heterogeneity* in overconfidence and measure it in two ways. One is the degree of overconfidence, defined as the self-assessed rank minus the actual rank so that a higher value of this "oc percentile rank" indicates more overconfidence. The second maps into a key model input: the population share of households exhibiting persistent overconfidence. To estimate this input we flag the 38 percent of respondents who are above-median rank in both 2014 and 2017 as "oc in both rounds" (the standard error on this prevalence estimate is 4pp).¹²

We are not aware of any other quantitative estimate of the share of consumers who are persistently overconfident about their ability, or some closely related object, in a plausibly representative national sample of the working-age population. Huffman et al. (2022) estimate that 45 to 48 percent of managers are over-confident about their performance in a repeated high-stakes workplace tournament held by a single employer. Moschini et al. (2023) find widespread over-optimism about college completion among 18 year-olds in the 1997 NLSY. Various theories explain how overconfidence can persist even in the presence of feedback (e.g., Heidhues et al. (2018) or Zimmermann (2020)).

Subjective financial condition forecasts and realizations. We link overconfidence about cognitive skills to consumers' forecasts of their future financial situation. The ALP elicits such forecasts, and subsequent realizations, in many of its survey modules, allowing us to build a panel of 21,586 forecast-realization pairs, provided by 3,467 ALP panelists (including many SZ panelists, as detailed below), across fourteen surveys administered in January and July from July 2010 to January 2016.

The ALP elicits forecasts with a question that 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)): "... do you think that a year from now you will be better off financially, or worse off, or just about the same as now?". These forecasts are highly correlated with expected income growth in the relatively small number of ALP surveys that also elicit an income forecast (Appendix Table A1). We measure realizations a year later with "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?". Both forecasts and forecast errors tilt strongly optimistic in the aggregate, regardless of the time period in our sample (Appendix

¹²Data limitations preclude us from estimating prevalence more precisely, by directly comparing each respondent's forecasted to actual percentile, because the forecast's integer percentile support is much more granular than the 15-question test realization's support.

Table A2).¹³ Forecast errors are persistent,¹⁴ and there is only modest evidence of learning over relatively long periods of time.¹⁵ Nor is there evidence of substantial overcorrection.¹⁶

Being especially interested in persistent heterogeneity across consumers, we construct three household-level measures of persistent optimism about financial situations. The first two are indicators equaling one if the proportion of potentially optimistic forecast errors (weakly) exceeds 0.5. The third is the proportion itself. We estimate that 27 to 40 percent of the sample are persistently optimistic in the SZ overlap sample. The SZ sample is key for our subsequent analysis because we have the requisite measures of overconfidence about cognitive skills only for those panelists. We obtain similar estimates of persistent optimism prevalence in the broader ALP sample.

Hand-to-Mouth status. To assess whether someone is (persistently) HtM, we use six different measures of financial constraints. Some of them have been used in previous work, others are new. Two of the six measures are from the two SZ modules. The other four we pull in from other survey modules completed by SZ respondents, so that we can link those additional HtM measures to cognitive skills and overconfidence thereon.

We start by detailing the two HtM measures from the SZ modules. For each of these, we create indicators for whether someone exhibits the symptom of HtM status in both 2014 and 2017. The first measure indicates *severe financial distress*, defined as reporting that any of four events happened in the previous 12 months: forced move, late payments, hunger, or foregone medical care. An estimated 28 or 31 percent of our sample exhibits this indicator in both 2014 and 2017 (for standard errors on these and other estimates of HtM prevalence see Table 1 Columns (7) and (8)). Our second measure classifies a household as HtM if its liquid net worth is less than half of total monthly household income. About 40 or 47 percent of our sample exhibits this indicator in both 2014 and 2017. Kaplan and Violante (2022) obtain a similar estimate, of 41 percent, in a snapshot from the 2019 Survey of Consumer Finances.

¹³Appendix Table A2 shows that forecasts are more than twice as likely to predict improvement (27 to 30 percent of observations) as deterioration (10 to 14 percent of observations). Forecast errors are roughly three times more likely to be in an optimistic than pessimistic direction; to see this, focus on the "same" realization column to allow for the possibility of forecast errors in either direction, and note that an estimated 13 to 18 percent of the sample forecasted better and ended up the same, while only 4 to 7 percent forecasted worse and ended up the same. Our findings are consistent with the evidence of persistent and strong excessive optimism, from several decades of consumer sentiment data across many wealthy countries, in Claus and Nguyen (2023), following Souleles (2004)'s similar findings from 1978-1996 U.S. data. The one counterexample we know of is Hyttinen and Putkuri (2018)'s evidence of aggregate mean-zero forecast errors from Finland.

¹⁴Appendix Table A3 shows that 74 percent of consecutive forecast errors are the same (both optimistic, both realistic, or both pessimistic), and that 53 percent of panelists who make an optimistic forecast error in the previous period make the same error in the next period.

¹⁵Comparing the first to last forecast-realization pair we observe for panelists with multiple pairs, Appendix Table A4 shows that the accuracy rate increases from 55 to 62 percent and the optimistic slant decreases from $16/21 = 77$ percent to $13/18 = 72$ percent.

¹⁶Appendix Table A3 shows that optimists are about 9 times more likely to get better-calibrated than to overcorrect with a pessimistic forecast error.

The third measure of HtM status is indicating strong agreement with the statement "I live from paycheck to paycheck" in a 2012 survey. An estimated 56 or 59 percent of our sample does so. Our fourth measure is closely related and draws on two questions asked in nine COVID-era modules administered May 2020-July 2022. The mean proportion of these modules in which a panelist exhibits paycheck-to-paycheck behavior is about 40 or 44 percent.¹⁷ Our fifth measure indicates whether someone lacks precautionary savings, defined as reporting not having emergency or rainy day funds set aside to cover 3-months of expenses. An estimated 63 or 72 percent of respondents, who completed both surveys where this question was asked, indicate this in either 2012 or 2018.¹⁸ Our sixth measure is based on whether the panelist indicates having difficulty dealing with expense shocks, measured as the proportion of 3 surveys from 2011, 2012, and 2018 where they do not express the highest confidence or certainty that they could cover an unexpected \$2,000 need arising in the next month. The mean proportion across panelists is about 51 or 59 percent, as compared to [Sergeyev et al. \(2024\)](#)'s estimate that 54 percent of U.S. households would have difficulty covering an unexpected \$2,000 emergency expense in 2022.

Overall, our estimates of HtM prevalence square well with those from prior work. They also suggest that we have measures of financial constraints of varying severity, which will be useful for exploring the robustness of our results below.

2.2 Key Correlations

We now use the above variables to estimate the key micro empirical relationships that shape and discipline our model.

2.2.1 Empirical strategy

In estimating empirical relationships between variables, we focus on pairwise correlations, for two reasons. One is empirical: pairwise correlations are easier to interpret when all of the variables of interest are correlated with each other; conversely, multivariate estimates are likely subject to confounds from over-controlling and multicollinearity. The other is conceptual: for modeling purposes, we are interested in identifying a proxy for persistent and relatively fundamental consumer heterogeneity (like overconfidence about cognitive skills) that can reproduce key empirical pat-

¹⁷For each panelist-survey we define an indicator that =1 if panelists respond "Very difficult" or "Somewhat difficult" to "In the past month, how difficult has it been for you to cover your expenses and pay all your bills?" or, on the followup question "Suppose now you have an emergency expense that costs \$400. Based on your current financial situation, how would you pay this expense?" they report one or more expensive options: credit card revolving, small-dollar credit, or that they wouldn't be able to pay for it. For each panelist we then take the ratio of the count of indicators to the count of completed surveys, across the nine modules.

¹⁸The indicator for lacking precautionary savings is strongly serially correlated within-person across the two surveys, with a tetrachoric correlation of 0.82 (s.e.=0.05) in the sample with nonmissing overconfidence. Unsurprisingly then, correlations are statistically indistinguishable if we define the measure as lacking precautionary saving in both 2012 and 2018. We report results for the either 2012 or 2018 version in the main table, in the interest of showing a measure that indicates relatively high HtM prevalence.

Table 1: Pairwise correlations between persistent HtM measures and cognitive skills or persistent overconfidence about skills

	CS rank: cf		1=Oc both rounds		Oc pctlile rank		Row variable, unw.	Row variable, w.
	Unw.	Weighted	Unw.	Weighted	Unw.	Weighted	Pop. share	Pop. share
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
1=(Severe financial distress)	-0.34	-0.29	0.18	0.27	0.19	0.18	0.28	0.31
s.e.	0.04	0.07	0.06	0.12	0.04	0.08	0.02	0.04
N	841	841	813	813	813	813	813	813
1=(Low net worth)	-0.40	-0.37	0.25	0.20	0.23	0.09	0.40	0.47
s.e.	0.04	0.06	0.06	0.10	0.04	0.07	0.02	0.03
N	788	788	760	760	760	760	760	760
1=(paycheck-to-paycheck c. 2012)	-0.29	-0.50	0.15	0.01	0.15	0.17	0.59	0.56
s.e.	0.07	0.08	0.10	0.24	0.07	0.12	0.03	0.08
N	263	263	255	255	255	255	255	255
paycheck-to-paycheck, COVID	-0.38	-0.28	0.22	0.20	0.30	0.29	0.40	0.44
s.e.	0.02	0.02	0.05	0.09	0.05	0.08	0.04	0.03
N	527	527	516	516	516	516	516	516
1=(Lacks precautionary savings)	-0.30	-0.30	0.11	0.09	0.18	0.19	0.63	0.72
s.e.	0.07	0.12	0.10	0.16	0.07	0.11	0.03	0.04
N	272	272	262	262	262	262	262	262
Difficult covering \$2k expense	-0.40	-0.43	0.23	0.31	0.22	0.25	0.51	0.59
s.e.	0.04	0.06	0.07	0.09	0.05	0.07	0.04	0.03
N	499	499	485	485	485	485	485	485

Note: CS = cognitive skills, measured as the common factor of four standard tests; OC= overconfidence re: relative performance in a cognitive skills test (see Section 2.1 for details). Weighted estimates use the sampling probability for the last SZ module. In Columns 5 and 6, we use Obviously Related Instrumental Variables to account for measurement error by having the two measurements of o/c rank (taken in 2014 and 2017) instrument for each other (Gillen et al. (2019); Stango and Zinman (2023)). We do not take the same approach to the o/c indicator in Columns 3 and 4, because measurement error-IV does not work well on misclassification error. Fully non-IV correlations estimated using tetrachoric or Pearson. See Section 2.1 for details on HtM measure definitions. The two non-indicator HtM variables are each defined as the proportion of indicators across multiple surveys, so for population share estimates we take the mean of the estimated population shares for each component indicator used in creating that variable.

terns in the aggregate (like patterns of forecast errors and financial constraints). The proxy can be useful, for modeling purposes, whether or not it has a causal relationship with the other variables of interest. We address measurement error in cognitive skills, overconfidence, and other potential sources of fundamental and persistent heterogeneity in decision making by using SZ’s repeated measurements as instruments for each other where advisable, following [Gillen et al. \(2019\)](#) and [Stango and Zinman \(2023\)](#).¹⁹

2.2.2 Cognitive skills and HtM status

As noted at the outset, cognitive skills heterogeneity has been linked to some variables of macroeconomic interest in prior work but not explicitly to HtM status and its persistence within-household over time.²⁰ Columns (1) and (2) in [Table 1](#) take steps towards filling that gap. We estimate unweighted and sampling-probability-weighted correlations between our cognitive skills summary measure and each of our six HtM measures, finding a negative sign on all 12 point estimates. All of them are larger than $|0.27|$ and most have t-stats of $|4|$ or more.

2.2.3 Overconfidence, forecasting, and HtM status

Given that cognitive skills heterogeneity alone is unlikely to help fit the macro data (as we show formally in [Section 4](#)), we now consider whether overconfidence about cognitive skills is a potential underpinning or proxy for the strong relationship between cognitive skills and persistent HtM documented in [Table 1](#) Columns 1 and 2. Indeed, overconfidence in relative performance is the behavioral bias most strongly correlated with cognitive skills out of the 17 biases measured in the SZ data ([Stango and Zinman \(2023\)](#)). Overconfidence could be a key link between cognitive skills and consumer behavior that has been overlooked so far.

Table 2: Optimistic forecast errors are more prevalent among the overconfident

(Optimist share overconfident)	Optimism measure	
(Optimist share not oc)	$1 = (\text{Prop. Opt. FEs} > 0.5)$	$1 = (\text{Prop. Opt. FEs} \geq 0.5)$
Unweighted	1.25	1.20
Weighted	1.08	1.06

Note: Sample is the 462 Stango-Zinman panelists who also provide the requisite data, in other ALP modules, to measure at least two potentially optimistic forecast errors. Overconfidence re: relative performance in a cognitive skills test (see [Section 2.1](#) for details). Weighted estimates use the sample probability weights from the last Stango-Zinman module.

[Table 2](#) links overconfidence about cognitive skills to over-optimism about one’s own future

¹⁹Measurement error IV is advisable for smooth measures but not for discrete ones—the latter are subject to misclassification error that is non-classical.

²⁰Recall that 5 of our 6 HtM measures explicitly capture persistence. Because HtM status is so persistent, results on HtM snapshots are similar and we do not report them below.

financial situation. We see that persistent optimism about one’s own future financial condition—as measured by our two indicators—is about 1.06 to 1.25 times more prevalent among persistently overconfident households than in the rest of the population. In our model calibration, we will use this ratio of relative over-optimism to discipline overconfidence.

Table 2 suggests that the strong negative relationship between cognitive skills and HtM status in columns 1 and 2 in Table 1 may be due at least in part to overconfidence. Columns 3-6 in Table 1 provide empirical support for that conjecture. Here we estimate 24 correlations: (6 HtM measures \times 2 overconfidence measures \times weighted or unweighted). All 24 point estimates are positively signed, and 17 have t -stats strictly greater than two.²¹ Relatedly, Grohmann et al. (2023) find that overconfident participants save less in a lab experiment.

2.3 Other sources of fundamental heterogeneity?

Other papers have put forth more classical sources of relatively fundamental heterogeneity as candidates for macro modeling; see e.g., Krueger et al. (2016), Auclert et al. (2020), Aguiar et al. (forthcoming), Kaplan and Violante (2022), and Andreou et al. (2023) on patience, and Kaplan and Violante (2022) and Kekre and Lenel (2022) on risk aversion. But we find that the micro data favors focusing on cognitive skills and overconfidence over patience or risk aversion. Stango and Zinman (2023)’s findings point to cognitive skills heterogeneity as the most likely source or summary statistic for heterogeneity in various behavioral biases, and moreover show that overconfidence in relative performance is the bias that has the strongest correlation with cognitive skills. Here we look directly at relationships between our other key micro variables for macro modeling on the one hand, and patience or risk aversion on the other. We do not find evidence of a robust relationship between those classical decision inputs and persistent over-optimism about financial condition, subject to the caveat that any nulls are imprecisely estimated (Appendix Table A6). Turning to HtM status, although we do find some evidence of potentially meaningful correlations with patience or risk aversion, overall the relationships are less robustly strong across our six HtM measures than they are with cognitive skills or overconfidence, both statistically and quantitatively, and patience has a surprising positive correlation with our pre-COVID measure of living paycheck-to-paycheck (Appendix Table A7). Nor is patience a good proxy for overconfidence (Appendix Table A8 Columns 1 and 2). Risk aversion might be, but the two different measures of presumed-classical risk aversion in the SZ data have opposite-signed correlations with overconfidence (Appendix Table A8 Columns 3-6), despite being positively correlated >0.2 with each other.

²¹Consistent with Tables 1 and 2, Table A5 shows strong correlations between over-optimism about financial condition and HtM status. All 30 point estimates are positive, most have t -stats > 3 , and 28 have t -stats larger than 2.

2.4 Summary of results from micro data

To summarize, we find that persistent HtM status decreases strongly with cognitive skills and increases with overconfidence thereon, and that overconfident consumers tend to be persistently too optimistic about their future financial situation. Together with prevalent overconfidence, and the strong negative correlation between cognitive skills and overconfidence found in prior work, these findings suggest that accounting for consumer heterogeneity in cognitive skills and/or overconfidence could be important for understanding macroeconomic fluctuations. We next develop a model to explore this possibility formally and quantitatively.

3 Model

We now develop an augmented HANK model, using our new results in Section 2, together with consensus estimates of key macro variable moments, to shape and discipline the model. Aside from adding heterogeneity in cognitive skills and overconfidence about these skills, 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 $t = 1, 2, \dots$. We first focus on the case in which households can only save in one asset—a liquid bond issued by the government. Later on, we introduce a second asset in the form of illiquid productive capital.

Households. There is a unit mass of households subject to idiosyncratic risk, incomplete markets, and borrowing constraints. We allow for permanent heterogeneity in households’ cognitive skills (modelled as productivity) and overconfidence about these cognitive skills (specifically about idiosyncratic productivity).²² An individual household’s productivity of permanent type g in period t are denoted by $\bar{e}_g e_t$, where \bar{e}_g captures permanent differences across groups in average productivity levels, and e_t captures idiosyncratic productivity. The stochastic component e_t follows a Markov process with time-invariant transition matrix \mathcal{P} . The process for e_t is the same for all households and the mass of households in state e is always equal to the probability of being in state e in the stationary equilibrium, $p(e)$.

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

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

²²We assume that heterogeneity in cognitive skills and overconfidence is permanent given the results in [Stango and Zinman \(2024\)](#). Consistent with that, [Hoffman and Burks \(2020\)](#) also find, among truckers, that workers’ over-optimistic beliefs about their productivity are very persistent.

subject to

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

$$b_t \geq -\underline{b}, \quad (2)$$

where c_t denotes consumption, n_t hours worked, r_t the net real interest rate, w_t the real wage, τ_t the income tax rate, and V the value function. We assume a standard CRRA utility function where the parameters γ , φ , and β denote relative risk aversion, the inverse Frisch elasticity of labor supply, and the time discount factor, respectively. These parameters as well as the exogenous borrowing limit \underline{b} are the same for all households and time-invariant.

The expectations operator $\tilde{\mathbb{E}}_{g,t}$ is our key innovation, and we discuss it next.

Cognitive skills and overconfidence. We model heterogeneity in cognitive skill levels as different average productivities \bar{e}_g , given the strong (negative) correlation between cognitive skills and income in the data (Stango and Zinman, 2023).

All households observe their current productivity $\bar{e}_g e_t$ but overconfident households have biased beliefs about the transition probabilities $p(e_{t+1}|e_t)$. Specifically, overconfident 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 makes overconfident households too optimistic about their expected future productivity, relative to rational households with the same productivity and idiosyncratic risk.

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 productivities are ordered such that $e_1 < e_2 < \dots < e_J$. To capture overconfidence with only one additional parameter independent of the number of states, we assume that an overconfident household's *perceived* transition probabilities \tilde{p}_{ij} are given by

$$\tilde{p}_{ij} \equiv \begin{cases} \alpha p_{ij}, & \text{if } i < j \\ \frac{1}{\alpha} p_{ij}, & \text{if } i > j \\ 1 - \sum_{j \neq i} \tilde{p}_{ij}, & \text{if } i = j, \end{cases} \quad (3)$$

where the parameter $\alpha \geq 1$ captures overconfidence. If $\alpha > 1$, the household assigns too much weight to reaching a better state (this is the case $i < j$) and too little weight to reaching a worse state ($i > j$). The perceived probability of staying in the same state ($i = j$) ensures that the probabilities sum to 1.²³ We discuss an alternative modelling approach in Section 4.3, where the degree of overconfidence depends on the distance between the states. Note that the rational

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

expectations case is captured by setting $\alpha = 1$ and thus nested in our setup.²⁴

An immediate implication is that overconfident households will more often be overly optimistic about their financial situation (specifically income, in the model) compared to rational households, consistent with the empirical findings reported in Section 2.2.3. We will use our empirical estimate of the relative share of optimists among overconfident and rational households from Table 2 to calibrate α below (in 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. (forthcoming), which is based on Erceg et al. (2000)).²⁵ Each worker provides $n_{k,t}$ hours of work to a continuum of unions indexed by $k \in [0, 1]$. Each union aggregates efficient units of work into a union-specific task

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

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

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

$$N_t = \left(\int_k N_{k,t}^{\frac{\epsilon-1}{\epsilon}} \right)^{\frac{\epsilon}{\epsilon-1}} \quad (4)$$

and sells these services to firms at price w_t .

We model wage stickiness by imposing a quadratic utility cost $\frac{\psi}{2} \int_k \left(\frac{W_{k,t}}{W_{k,t-1}} - 1 \right)^2 dk$ that shows up in the household's utility function. A union sets a common nominal wage $W_{k,t}$ per efficient unit for each of its members.

In doing so, the union trades-off the marginal disutility of working given average hours against the marginal utility of consumption given average consumption. The union then calls upon its members to supply hours. We assume the union ensures that each household supplies the same amount of hours.

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

$$Y_t = N_t, \quad (5)$$

²⁴Modelling overconfidence as in (3) 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.

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

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

$$w_t = 1. \quad (6)$$

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

$$1 + \pi_t = 1 + \pi_t^w, \quad (7)$$

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

$$T_t - \bar{T} = \vartheta \frac{B_t - \bar{B}}{\bar{Y}}, \quad (8)$$

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

$$B_t + T_t = (1 + r_t)B_{t-1}. \quad (9)$$

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 overconfident consumers change the effectiveness of temporarily increasing fiscal transfers in Section 5.1, .

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

$$\sum_{g,e} \mu_g p(e) \int c_t \Psi_{g,t}(b_{t-1}, e_t) = Y_t \quad (10)$$

$$\sum_{g,e} \mu_g p(e) \int b_t \Psi_{g,t}(b_{t-1}, e_t) = B_t, \quad (11)$$

where μ_g denotes the mass of agents of type g .

3.1 Calibration

One period in the model corresponds to a quarter. We calibrate the standard parameters to values often used in the literature. For the idiosyncratic productivity process, we follow [McKay et al. \(2016\)](#) and set the autocorrelation of e_t to $\rho_e = 0.966$ and the variance to $\sigma_e^2 = 0.033$. We then discretize this process into an eleven-states Markov chain using the [Rouwenhorst \(1995\)](#) method. 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 \(2022\)](#)).

Table 3: Persistent overconfidence: prevalence and relationship to income

	Overconfident in both survey rounds?			
	Yes	No	Yes	No
	Unweighted	Unweighted	Weighted	Weighted
Population share	0.34 (0.02)		0.38 (0.04)	
Mean Income	\$51,182	\$79,765	\$42,035	\$77,145
N	817	817	817	817

Note: Standard errors in parentheses. Weighted estimates use the sampling probability for the last SZ module. Income is the sample mean of each panelist’s mean income across the two SZ modules.

We set the share of overconfident households to 0.38 (Table 3), using the higher estimate from our data in light of [Huffman et al. \(2022\)](#)’s finding of even higher prevalence in a high-stakes workplace tournament.²⁶ Based on prior work showing strong negative correlations between cognitive skills and overconfidence about those skills (see [Ehrlinger et al. \(2008\)](#); [Stango and Zinman \(2023\)](#) with additional results here in Table A9), and in the interest of parsimony, we collapse permanent heterogeneity in skills and confidence to two types: overconfident with low skills, and rational with high skills. We normalize the average productivity of the high-skilled and rational households to $\bar{e}_2 = 1$ and set the average skill level of the low-skilled and overconfident households to $\bar{e}_1 = 0.55$, based on our weighted estimates of average income for overconfident vs. rational households in Table 3: $0.55 = \frac{42,000}{77,000}$.

Following equation (3), we capture the degree of overconfidence in the overconfident and low-skilled group with one parameter, α . To calibrate α , we target our estimates from Table 2 that overconfident households are more likely to have optimistic one-year forecast errors about their financial situation,²⁷ using a medium value of 1.18 as our target. This results in $\alpha = 2$. Below

²⁶Using our lower estimate of 34% changes our quantitative results only slightly. For example, it changes the share of HtM from 29% to 27% and the average MPC from 0.16 to 0.15.

²⁷Note that in the stationary equilibrium of our model a household that is overly optimistic about its future idiosyncratic productivity is also overly optimistic about its future financial situation (defined as labor income plus asset income). The reason is that wages, hours worked, and asset returns are constant and therefore the only

Table 4: Stationary equilibrium calibration

Parameter	Description	Value
R	Steady state real rate (annualized)	4%
γ	Risk aversion	2
φ	Inverse of Frisch elasticity	2
\underline{b}	Borrowing limit	0
$\frac{\bar{B}}{4Y}$	Average wealth to average income	4.1
<u>Idiosyncratic risk</u>		
ρ_e	Persistence of idiosyncratic risk	0.966
σ_e^2	Variance of idiosyncratic risk	0.033
<u>Permanent heterogeneity</u>		
μ_g	Mass of households	{0.38, 0.62}
\bar{e}_g	Cognitive skills	{0.55, 1}
α	Degree of overconfidence	2

Note: Calibration summary for our one-asset model using two groups to capture permanent heterogeneity: households in group one have relatively low average skill levels $\bar{e}_1 < \bar{e}_2$ and are overconfident ($\alpha > 1$), group two is relatively high-skilled and has rational expectations ($\alpha = 1$).

we consider several alternative parameterizations of heterogeneity in cognitive skills and overconfidence and find similar results. Table 4 summarizes our baseline calibration.

4 Stationary Equilibrium Predictions

We now consider our model’s ability to fit various key moments from macro and micro data, as compared to HANK models that abstract from cognitive skills or belief heterogeneity or both.

4.1 Hand-to-Mouth Shares and Average MPCs

We start by considering the effects of permanent heterogeneity in cognitive skills and overconfidence on the share of Hand-to-Mouth (HtM) households and the implied average marginal propensity to consume (MPC) of households.²⁸

Table 5 compares predictions across four different models: our baseline model with heterogeneity in cognitive skills and overconfidence ("*HANK: CS + OC*", in Column 1), a standard HANK model (Column 2) with no heterogeneity in permanent productivity levels ($\bar{e}_g = \{1, 1\}$) and full rationality ($\alpha = 1$), a HANK model with permanent heterogeneity in skill levels but full rationality ("*HANK: CS*", Column 3), and a HANK model with a group of permanently overconfident households but no skill heterogeneity ("*HANK: OC*", Column 4).²⁹ We start by comparing our

possible variation in a household’s financial situation comes from changes in idiosyncratic productivity.

²⁸Here we define HtM as holding less liquid wealth than half of average monthly income.

²⁹When comparing these four different models, we take the standard approach of recalibrating the discount factor such that all 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 all models.

model to the standard HANK, and then use the other two models to help unpack the differences.

Table 5: MPCs and shares of HtM households across the models.

	HANK: CS + OC	Standard HANK	HANK: CS	HANK: OC
	(1)	(2)	(3)	(4)
HtM Share	0.29	0.03	0.04	0.28
Avg. MPC	0.16	0.04	0.04	0.18
HtM rational HHs	0.02	0.03	0.03	0.02
Avg. MPC rat. HHs	0.03	0.04	0.04	0.03
HtM OC HHs	-	-	-	0.70
Avg. MPC OC HHs	-	-	-	0.43
HtM rat. HHs Low-Skilled	-	-	0.04	-
Avg. MPC rat. HHs LS	-	-	0.03	-
HtM OC HHs LS	0.74	-	-	-
Avg. MPC OC HHs LS	0.38	-	-	-

Note: MPCs refer to MPCs out of a \$500 dollar stimulus check. "HANK: CS + OC" is our baseline model (one-asset, with heterogeneity in cognitive skills and overconfidence); "Standard HANK" denotes a standard one-asset model that abstracts from heterogeneity in skills and overconfidence; "HANK: CS" adds heterogeneity in skills only to Standard HANK; "HANK: OC" adds heterogeneity in overconfidence only to Standard HANK.

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

In contrast, our model with skill and belief heterogeneity (Column 1) produces an average MPC and a HtM share that are both multiple times larger than in the standard one-asset HANK 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 16%. Our predicted share of HtM households, 0.29, is in the range of our estimated empirical share based on our most conservative definition of HtM status: those with severe financial distress (Table 1).

Column 3 shows that skill heterogeneity alone does not drive our model's ability to fit the data better. If we introduce skill heterogeneity but keep all households rational (i.e., well-calibrated about their productivity), the average MPC and the HtM share are very similar to those produced by the standard HANK model. The reason is that a rational household still has a strong incentive to self-insure regardless of its average productivity.

Column 4 shows that our model's allowance for belief heterogeneity drives its improved performance. Specifically, keeping average productivity homogeneous but allowing some households to be overconfident about their future idiosyncratic productivity generates average MPCs and HtM

shares that are consistent with the data. The mechanism is that overconfident households overestimate their expected income; i.e., they perceive their income risk to be lower than it actually is. Overconfident households thus accumulate less precautionary savings than rational households facing the same actual income risk. Consequently, and in line with our empirical findings in Section 2, overconfident households are much more likely to end up being HtM than rational households in our model (74% of overconfident households are HtM, while only 2% of rational households are). This also results in a high average MPC for the group of low-skilled, overconfident households (0.375, vs. 0.027 for the rational households), driving up the aggregate average MPC. These results are consistent with [Bernard \(2023\)](#)'s empirical finding that a lack of cognitive sophistication is positively correlated with MPCs.

4.2 "Missing Middle Problem" and the Top 10% Wealth Share

Standard one-asset HANK models can generate a high average MPC by restricting wealth to be many multiples lower than consensus estimates ([Wolf \(forthcoming\)](#), [Kaplan and Violante \(2022\)](#), [Seidl and Seyrich \(2023\)](#)). This restriction also produces an excessively polarized wealth distribution ([Kaplan and Violante, 2022](#)). One way to see this "Missing Middle" problem is that median wealth to mean annual earnings is about an order of magnitude smaller in standard HANK models than in the data. We offer further confirmation of this finding by recalibrating the standard HANK model used in Table 5 Column 2 to match the average MPC produced by our one-asset model with skill and belief heterogeneity. Matching the average MPC requires setting total wealth to income to 0.7 instead of 4.1, and delivers a median wealth-to-average annual income ratio of 0.2 vs. about 1.5 in the data ([Kaplan and Violante, 2022](#)).

Our one-asset model with heterogeneity in cognitive skills and overconfidence fills in the missing middle: it predicts a median wealth-to-average annual income ratio of 1.4 that is close 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. Overconfident households tend to be HtM and thus account for most of the bottom, 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 45% of wealth, as compared to the empirical estimate of 49% ([Kaplan and Violante, 2022](#)). Overall, our model produces a wealth distribution that matches the data well.

Discount factor heterogeneity. As illustrated by [Krueger et al. \(2016\)](#), [Aguiar et al. \(forthcoming\)](#), or [Kaplan and Violante \(2022\)](#), ex-ante heterogeneity in discount factors β can help the rational model account for some of the MPC patterns observed in the data. [Aguiar et al. \(forthcoming\)](#) suggest that behavioral frictions could provide a microfoundation for the low β of some households. Yet our empirical evidence in Section 2 points towards overconfidence and not impatience as having a strong connection to HtM status. Similarly, [D'Acunto et al. \(2023a\)](#) find links

between other key macro variables and cognitive skills that cannot be explained via heterogeneity in patience.

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

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

Proof. See Appendix B. ■

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

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. But this comes at the cost of also producing two unattractive features. First, it requires using the discount factor of the impatient households as a free parameter to match the average MPC. Second, it tends to produce wealth distributions with a missing middle (as shown by [Kaplan and Violante \(2022\)](#)).

We further show that the two models can also differ vastly in their normative implications (Section 5.2.2), highlighting that it matters for the optimal debt level *why* households differ in their savings behavior and HtM status.

4.3 Extensions

We now show that our results are robust to: (i) Accounting for the empirical finding that 11% of households are persistently *underconfident*, and (ii) An alternative specification of overconfidence that is state dependent. We then extend the model to incorporate a second, productive asset.

4.3.1 Underconfident households.

Our survey data suggests that 11% of households are persistently underconfident, defined as underestimating their cognitive skills in both survey rounds. We extend our model to account for this by setting $\alpha_{uc} < 1$ for 11% of households and adding a symmetric target to its calibration: we now not only target overconfident households being 1.18 times as likely to be optimistic about their future situations than their rational counterparts, but also underconfident households being 1.18 times as likely to be too pessimistic. We again allow the discount factor to adjust keep the real interest rate at 4% annually.³⁰

³⁰This requires a discount factor of 0.981 instead of 0.982.

Incorporating underconfident households actually increases the overall HtM share slightly from 29.2% to 30.0% and the average MPC from 16.3% to 16.7%. In partial equilibrium, one would see effects in the opposite direction, because underconfident households overestimate their precautionary savings motive compared to rational households. This pushes underconfident households to save more than the rational households they are replacing in the model. The underconfident are then slightly less likely to end up HtM (recall that HtM probability is quite low for rational households in any model). In general equilibrium, the added savings demand from underconfident households pushes up the asset price, crowding out savings from the larger mass of households close to the borrowing constraint.

Overall, extending the model by accounting for underconfident households further illustrates how adding heterogeneity in beliefs about skills can help improve model performance in general equilibrium. But given the small share of underconfident households in the data, adding them to our model has only small quantitative effects.

4.3.2 Alternative way of modelling overconfidence.

In our baseline specification of overconfidence (equation (3)), the degree of overconfidence is the same for all overconfident households, independent of their current state or skill level. We now allow for dependence of the following form:

$$\tilde{p}_{ij} \equiv \begin{cases} \alpha^{(e_j - e_i)} p_{ij}, & \text{if } i \neq j \\ 1 - \sum_{j \neq i} \tilde{p}_{ij}, & \text{if } i = j. \end{cases} \quad (12)$$

As in our baseline specification, when $\alpha > 1$, 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.³¹

We again calibrate α to match the empirical finding that overconfident households are about 1.18 times as likely to be overly-optimistic about their future financial situation than rational agents. This implies $\alpha = 2.65$. The predicted average MPC is 0.175 and thus largely unchanged from our baseline estimate of 0.163. The predicted HtM share is now about 10 percentage points higher, at 39.2%, and thus closer to the empirical shares of more expansive definitions of HtM (see Table 1).

4.3.3 Overconfidence 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

³¹This 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.

(2022), Auclert et al. (forthcoming)). 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 a standard two-asset HANK model.

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, we add capital to the production function.

The household’s budget constraint now reads:

$$c_t + \frac{b_t}{1 + r_t} + k_t = b_{t-1} + (1 + r_t^k)k_{t-1} + (1 - \tau_t)w_t \bar{e}_g e_t n_t, \quad (13)$$

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

$$b_t, k_t \geq 0.$$

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

$$Y_t = K_{t-1}^\chi N_t^{1-\chi}, \quad (14)$$

where χ denotes the capital share in production.

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

$$\sum_{g,e} \mu_g p(e) \int k_t \Psi_{g,t}(k_{t-1}, e_t) = K_t. \quad (15)$$

Calibration. We maintain the same values for each of the parameters that also appear in our baseline model (except for the discount factor). Table 6 shows our calibration of the additional parameters and the discount factor. We set the capital share to $\chi = 0.318$ and the quarterly depreciation rate to $\delta = 0.0175$ as in Bayer et al. (forthcoming). We then use the per-period

capital market participation probability λ and the discount factor β to jointly target the average wealth-to-annual income ratio of 4.1 and the liquid asset-to-annual income of 0.2 as in [Kaplan and Violante \(2022\)](#).

Table 6: Calibration two-asset model

Parameter	Description	Value
χ	Capital share	0.318
δ	Depreciation rate	0.0175
λ	Capital market participation rate	0.37
β	Discount factor	0.992

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

Stationary Equilibrium Results. Table 7 shows the influence of overconfident households on the stationary equilibrium (Column 1). We start by explaining the mechanisms underlying our results, and then compare the empirical fit of our model to standard models.

The share of HtM households is now 0.38, as compared to 0.29 in our baseline model, because the illiquid asset's higher return induces some savers to substitute from the liquid asset. This is mostly driven by "wealthy HtM" households who would not be HtM in a one-asset model and now choose to save only in the illiquid asset ([Kaplan et al. \(2018\)](#)). In contrast, the average MPC increases by only about one percentage point here relative to our baseline model, to 0.171, indicating that in the two-asset model, on average, constrained households are not as far off their Euler equation and thus spend less out of a \$500 windfall. There is again a stark difference between the behavior of rational and overconfident households in our model. Rational households accumulate liquid assets to self-insure before saving in the illiquid asset. Overconfident households remain much more likely to be HtM (77% vs. 14%) because they foresee little value in accumulating a liquid buffer stock and hence prioritize the illiquid asset's higher return if they do save.

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

	Two-asset HANK with overconfidence	Rational two-asset HANK	
	(1)	(2)	(3)
		Calibrated as (1)	Recalibrated
HtM	0.38	0.23	0.27
Avg. MPC	0.17	0.06	0.15
return gap (annualized)	2.3%	4.4%	9.3%
HtM rat. HHs	0.14	0.23	0.27
Avg. MPC rat. HHs	0.04	0.06	0.15
HtM OC HHs	0.77	-	-
Avg. MPC OC HHs	0.39	-	-

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

Column 2 presents a standard two-asset HANK model for comparison, keeping all the parameters the same as in our model except for recalibrating β to target the mean wealth-to-annual income ratio of 4.1.³² Unlike our model, this produces a HtM share of 0.23 and average quarterly MPC of 0.06 that are substantially below the lower end of the consensus ranges of empirical estimates.³³ Targeting an average MPC at the lower end, e.g. 0.15, requires an annualized return gap of 9.3% (Column 3).³⁴ Our model produces a much lower return gap of 2.3% because overconfident households underestimate precautionary savings needs and thus require a much smaller premium on illiquid assets, thereby driving demand for the illiquid asset up and its return down.

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

³²This requires quarterly $\beta = 0.989$, as compared to 0.992 in our model.

³³Compared to our model, the average MPC in the rational model is also lower conditional on the share of HtM households, because overconfident households have higher MPCs conditional on their current state given their lower perceived income risk.

³⁴In targeting the quarterly average MPC of 0.15 we set $\beta = 0.9805$, $\lambda = 0.15$, and $\delta = 0.00875$.

5 Policy Implications of the Systematic Relationship between Cognitive Skills, Overconfidence, and HtM Status

We now show that heterogeneous overconfidence matters for the design and effectiveness of fiscal policy tools seeking to stimulate and/or insure consumption. There are two key mechanisms. First, our model produces more higher-income HtM households, as overconfidence is a key predictor of HtM status even conditional on income. This has implications for the effectiveness of income-targeted transfer payments on stimulating private consumption. Second, overconfident households undervalue self-insurance and hence are less responsive to changes in precautionary savings incentives. This dampens crowdout when the government provides insurance (we consider a minimum income benefit as an example), but makes it more difficult to induce households in the neighborhood of the borrowing constraint to self-insure when the government provides liquidity (through higher public debt levels).

5.1 The distribution of HtM households and targeted transfers


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

As such, matching the HtM-income distribution is important for accurately evaluating the stimulative efficacy of targeted transfer policies (Figure 1). The figure’s gray hashes show the relationship between HtM status and income found in our microdata (with plus/minus two standard errors also in grey).³⁵ The standard one-asset HANK model with high wealth, depicted by the black diamonds, unsurprisingly underestimates the HtM shares at all income levels: almost everyone saves their way out of low wealth for precautionary reasons, even at low incomes. The standard HANK model recalibrated to generate the same average MPC as our baseline model, depicted by the red squares, produces more HtM households but far too many of them are low-income. Low income predicts HtM status counterfactually strongly in standard models because households become HtM solely due to bad luck. Our model (depicted by the blue dots) also has overconfident households choosing to be HtM throughout the bottom three income quartiles, thereby better matching empirical estimates.

Our model’s more realistic depiction of the HtM-income gradient implies that transfers targeted to low-income households are less effective at stimulating consumption than the standard HANK

³⁵We use our "severe financial stress" empirical HtM measure for this comparison because it yields the same aggregate HtM share as our baseline model, thereby giving the model an opportunity to match the HtM shares along the income distribution. Although the levels of our different HtM measures differ quite substantially, their relative steepness along the income distribution are similar. Appendix Figure A1 shows this for our HtM measure based on liquid net worth-to-income.

Figure 1: HtM shares along the income distribution



Note: "HANK: CS + OC" is our baseline model. "Standard HANK, low wealth" is the standard HANK model recalibrated to match the average MPC of our baseline model. "Data" shows our prevalence estimates for the severe financial distress HtM measure.

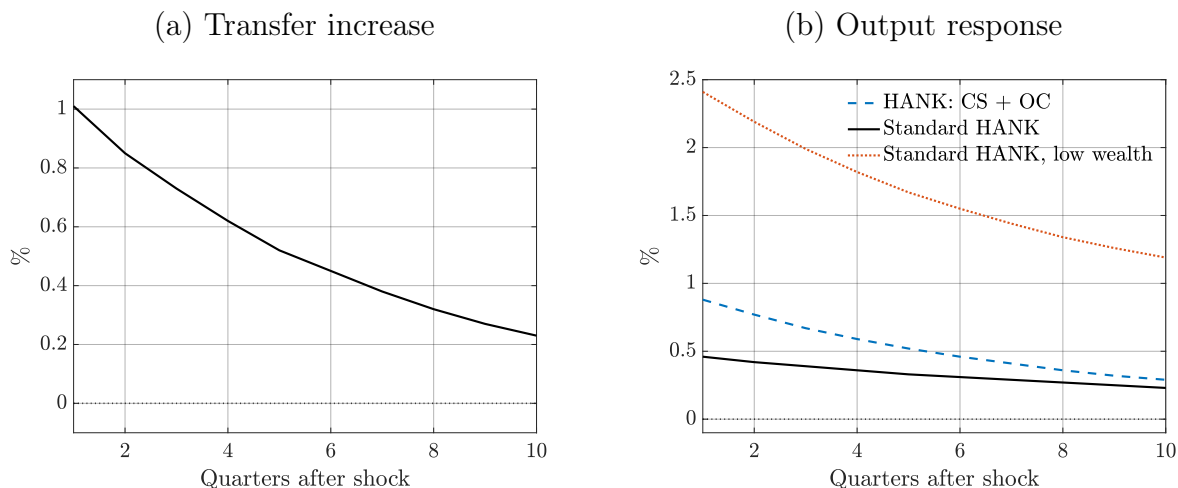
model with the same average MPC would imply, because income is a much weaker predictor of MPCs in our model. Consider a surprise lump-sum transfer to each household in the bottom income quartile, in an aggregate amount of 1 percent of steady-state output on impact, following an AR(1) process with a persistence parameter of 0.8, and financed in the short-run by higher debt which is then slowly repaid with higher taxes. Figure 2(a) shows the exogenous path of the transfers, and 2(b) shows output, in terms of percentage deviations from steady-state output. Our model (dashed blue line) predicts an output response that is less than half as strong than in the standard HANK model with the same average MPC (dotted orange line): the standard HANK model implies a transfer multiplier of about 2.4 on impact whereas our model implies 0.9.³⁶ These comparisons highlight that the average MPC is not a sufficient statistic for analyzing targeted stimulus: the distribution of MPCs is important as well.

A second channel further weakens the effectiveness of targeted transfers in our model: muted relaxation of the precautionary saving motive. The persistence in the transfer payments provides some insurance, temporarily decreasing households' precautionary savings motive. This motive is prevalent and strong in the standard model with classically rational households, further increasing spending and total output.³⁷ But overconfident households undervalue the insurance because they underestimate the likelihood of being income-eligible to exercise the insurance option in the future. They thus *perceive* their precautionary savings motive to be less relaxed than rational households

³⁶The low-MPC standard HANK model (solid black line) produces a multiplier of about 0.5, due to its low MPC across all income groups.

³⁷See e.g. Bayer et al. (2023b)'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) and Dengler and Gehrke (2023) find similar results for temporary increases in unemployment benefits and "short-term work", both of which can be understood as targeted transfers although they are not lump-sum and thus additionally have distortionary effects. Beraja and Zorzi (2024) analyze potential size-dependency for stimulus transfers.

Figure 2: Targeted Transfer Shocks



Note: This figure shows the effects of a positive transfer shock (left panel) on total output (right panel). Both are expressed in percentage deviations from steady state output. "HANK: CS + OC" is our baseline model. "Standard HANK, low wealth" is the standard HANK model recalibrated to match the average MPC of our baseline model..

would, and as such do not increase their spending as much.³⁸ This second channel mirrors recent empirical evidence that cognitive constraints can limit the effectiveness of macro policies designed to induce behavior change through incentive changes (D’Acunto et al., 2023a).

5.2 Precautionary Savings Behavior and Fiscal Insurance Policies

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

5.2.1 Minimum income benefits as public insurance

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

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

where \bar{y} is the median income in the stationary equilibrium and $0 \leq a_1, a_2 \leq 1$. Transfers thus decrease in individual income at the withdrawal rate a_2 and no transfers are paid to households

³⁸The relaxation of the precautionary savings motive is also an important driver in the standard HANK model with low average MPCs (black-solid line in Figure 2). But the MPCs are so low in that model, across all income quartiles, that it still predicts a smaller effect on aggregate output than our model.

whose labor income satisfies $w_t n_{i,t} e_{i,t} \geq \frac{a_1}{a_2} \bar{y}$. Following Bayer et al. (2023c), we set $a_1 = 0.5$ and $a_2 = 0.8$. and assume for simplicity that these transfers do not distort labor supply.

Total government transfer payments are then:

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

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

Table 8: Effects of introducing public insurance

	HANK: CS + OC	Standard HANK	Standard HANK, low wealth
	(1)	(2)	(3)
HtM Share	0.29	0.03	0.30
Avg. MPC	0.16	0.04	0.16
Bottom50W	2.7%	12.8%	3.0%
Real rate	4%	4%	4%
HtM Share with MIB	0.32	0.09	0.40
Avg. MPC with MIB	0.15	0.06	0.26
Bottom50W with MIB	1.6%	9.2%	1.3%
Real rate with MIB	5.0%	5.5%	6.9%

Note: MPCs refer to MPCs out of a \$500 dollar stimulus check. "HANK: CS + OC" is our baseline model (one-asset, with heterogeneity in cognitive skills and overconfidence), "Standard HANK" denotes a standard one-asset model, in which we abstract from heterogeneity in skills and overconfidence, "Standard HANK low wealth" is the same HANK model but with restricted liquidity to match the average MPC of "HANK: CS + OC". "... with MIB" refers to the stationary equilibrium in the models with public insurance via minimum income benefits (MIB).

Table 8 compares the stationary equilibrium effects of MIB on the average MPC and HtM share in our baseline model (Column 1) to a standard rational one-asset HANK model (Column 2). We also again consider a standard HANK model in which we reduce the amount of wealth such that it produces the same average MPC in the absence of transfers as our model does (Column 3).

In the two standard models, targeted transfers crowd-out self-insurance precautionary savings in the stationary equilibrium quite strongly. Households correctly forecast the probability of a bad productivity shock and thus internalize the insurance value of receiving a transfer in that state, reducing their precautionary savings accordingly. This increases the average MPC by more than 50% in either standard model, and the HtM share also increases substantially (by 6pp from the low base in Column 2, and by 10pp on the base of 30 in Column 3). Crowd-out is also reflected by the large increase in the equilibrium real interest rate from 4% to 6.9%. This higher rate is required to induce non-HtM households to hold the liquidity foregone by those moving to the borrowing constraint in response to the policy. Overall then, under standard HANK, introducing minimum income benefits as social insurance produces an economy with substantially higher interest rates,

less precautionary savings, more HtM households, and a higher average MPC.

In our model, crowd-out and its concomitant effects are dampened because overconfident households underpredict their probability of reaching a low-productivity state in which they receive a transfer. The average MPC even slightly decreases from 0.163 to 0.151,³⁹ while the share of HtM households only mildly increases from 29.2% to 32.1%. The real interest rate increase is also substantially smaller, rising only to 5.0%.

5.2.2 Liquidity Provision and the Optimal Public Debt Level

Fiscal policy can also facilitate private insurance, by issuing more government debt (e.g., [Woodford \(1990\)](#)). More debt increases the supply of liquid assets and thus of self-insurance possibilities for households. But this increase in liquidity supply has muted effects in our model compared to the rational HANK model. Figure 3a shows the share of HtM households, and 3b the share of wealth held by the poorest 50% of households, as a function of the government debt level in steady state.⁴⁰

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

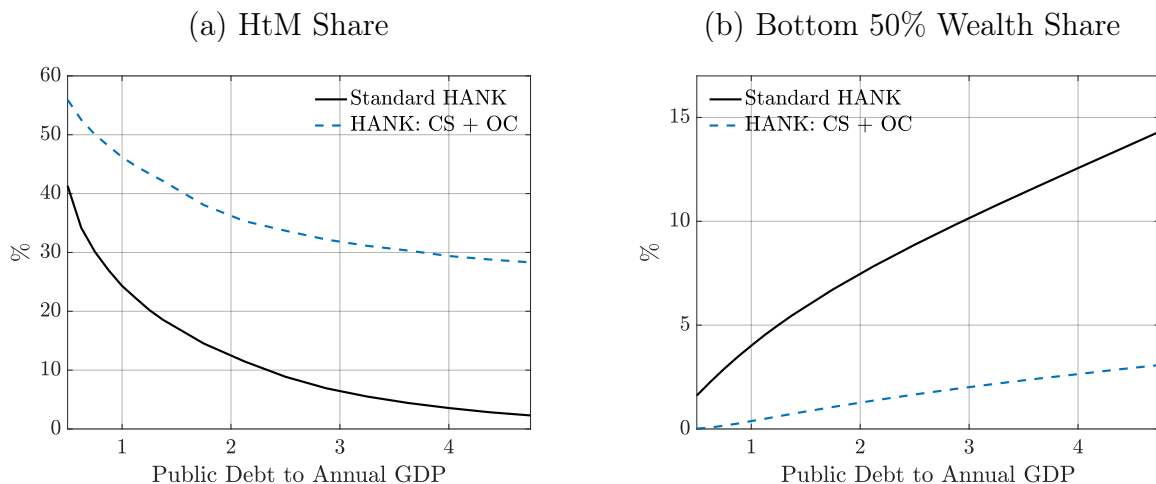
The dashed blue lines in Figure 3 illustrate the much weaker household response to liquidity provision in our model. The share of HtM households has a relatively flat slope with respect to debt supply, and it plateaus well above zero; e.g., it is about 0.29 at a debt-to-GDP ratio of 4, compared to nearly zero in the standard model. The bottom 50% wealth slope is remarkably flat, reaching only about a 3% share at a debt-to-GDP ratio of 4 compared to about 13% in the standard model. Even when liquidity is abundant, overconfident households do not tend to save themselves out of being liquidity constrained because they still perceive the liquid asset price as too high compared to their underestimated income risk.

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

³⁹There are two opposing effects of the introduction of MIB on the average MPC: on the one hand, the effective lower income risk reduces households' MPC conditional on their individual state. On the other hand, there are more households in individual states with higher MPCs as MIB crowd out precautionary savings. In the rational models, the latter dominates whereas in our baseline model, the first effect dominates because minimum income benefits only mildly crowd out households' precautionary savings.

⁴⁰When varying the supply of government debt, we fix the discount factor β as calibrated in Table 4 and let the interest rate adjust to clear the bond market.


Figure 3: The Implications of Higher Government Debt



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

maximize the average expected discounted lifetime utility of households.⁴¹

Figure 4: Public Debt and Social Welfare



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

Figure 4 shows that average welfare peaks at a much lower debt level in our model compared to the standard one-asset HANK model: optimal debt is about 135% of annual GDP, compared to about 400% in the standard HANK model. Since overconfident households underestimate their

⁴¹Such an objective function takes into account aggregate efficiency, risk-sharing, and intertemporal-sharing (Dávila and Schaab, 2023b). 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".

income risk and therefore have a dampened response to the liquidity supply increase even when they are at or close to the borrowing constraint, the very households that the social planner would like to save more are the least responsive ones. This diminishes the social benefit of higher government debt compared to the standard model. Even though we abstract from many important channels here—and therefore, our quantitative estimates should be interpreted with caution—the mechanism through which heterogeneity in overconfidence reduces the optimal debt level likely holds in richer models as well.⁴²

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. In the model with discount factor heterogeneity, the optimal debt level is 2.5 times as high as in our baseline model because the households who benefit more from government liquidity provision (those with higher discount factors, because they value precautionary savings more) also get de facto higher social welfare weights in a utilitarian welfare function (because their future utility is discounted less). As such, accounting for the strong empirical relationships between overconfidence, savings behavior and HtM status in Table 1, rather than relying on heterogeneity in patience (and its weaker empirical links to HtM status in Table A7), can matter greatly for optimal policy.

6 Conclusion

We analyze implications of heterogeneity in cognitive skills and self-perceptions thereof for households' savings behavior and financial situations, macroeconomic fluctuations, and fiscal policy. We start with U.S. micro data and find that lower-skilled households systematically overestimate their skills and are persistently overly optimistic about their future financial situations. They are also substantially more likely to be persistently HtM.

Guided by these findings, we then introduce persistent heterogeneity in skills and overconfidence into a HANK model and uncover a systematic reason why many households are persistently HtM: "bad decisions", not just "bad luck". Accounting for this reason, in the form of overconfidence about future productivity, 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 model also matches the income-HtM distribution whereas the rational model does not. Our two-asset HANK model matches the data

⁴²In a robustness exercise, we analyze the optimal debt level in our two-asset model and its rational counterpart. Overconfidence again reduces the optimal debt level significantly, although for both models the level of optimal debt is lower than in the respective one-asset models due to crowdout of productive capital. See e.g., [Aiyagari and McGrattan \(1998\)](#), [Davila et al. \(2012\)](#), [Angeletos et al. \(2023\)](#), or [Woodford \(1990\)](#) for analyses of optimal public liquidity provision.

with a lower, and perhaps more empirically realistic, liquidity premium than required in other models. It turns out that our key innovation is the overconfidence of low-skilled households rather than their lower productivity level. Thus, our model requires only one additional parameter—the degree of overconfidence of low-skilled households, as disciplined by our empirical findings—to substantially improve the empirical fit of existing HANK models.

We also show that accounting for the underlying reason why some households are persistently financially constrained matters greatly for fiscal policies. This is particularly pronounced for policies that affect the precautionary savings incentives of households, because overconfident households undervalue insurance and thus have muted responses to changes in such incentives. It also matters for income-targeted transfers, because in our model—as in the micro data—income is much less strongly correlated with HtM status and hence much less of a summary statistic for the MPC than in standard models.

One consideration for future work on normative questions—we mostly consider positive ones in this paper—is whether overconfidence may not be all bad, from a welfare perspective (as in, e.g., [Brunnermeier and Parker \(2005\)](#)). If it is not all bad, quantitative welfare modeling might seek to account for the benefits. Regardless, our finding that overconfidence correlates strongly with persistent and severe financial distress suggests important costs—costs that might be amplified by financial stress ([Sergeyev et al. \(2024\)](#)).

We also stop short of examining different combinations of macroeconomic policies in the presence of permanent heterogeneity across households—but our model provides a framework for doing so going forward. Consideration of monetary policy, and fuller consideration of fiscal policy, likely will require accounting for an additional source of heterogeneity: beliefs about aggregate variables. Some recent papers find empirical links between heterogeneity in expectations about such variables and cognitive skills ([D’Acunto et al., 2023a,b](#)). Modeling such links should be a fruitful new line of inquiry.

References

- ACHARYA, S., E. CHALLE, AND K. DOGRA (2023): “Optimal monetary policy according to HANK,” *American Economic Review*, 113, 1741–1782.
- AGARWAL, S. AND B. MAZUMDER (2013): “Cognitive abilities and household financial decision making,” *American Economic Journal: Applied Economics*, 5, 193–207.
- AGUIAR, M. A., M. BILS, AND C. BOAR (forthcoming): “Who are the Hand-to-Mouth?” *Review of Economic Studies*.
- AIYAGARI, S. R. (1994): “Uninsured idiosyncratic risk and aggregate saving,” *The Quarterly Journal of Economics*, 109, 659–684.
- AIYAGARI, S. R. AND E. R. MCGRATTAN (1998): “The optimum quantity of debt,” *Journal of Monetary Economics*, 42, 447–469.
- ALMGREN, M., J.-E. GALLEGOS, J. KRAMER, AND R. LIMA (2022): “Monetary policy and liquidity constraints: Evidence from the euro area,” *American Economic Journal: Macroeconomics*, 14, 309–340.
- ANDREONI, J. AND C. SPRENGER (2012): “Estimating time preferences from convex budgets,” *American Economic Review*, 102, 3333–3356.
- ANDREOU, E., M. DEMETRIADOU, AND A. TRYPHONIDES (2023): “The Cross Section of Household Preferences and the Marginal Propensity to Consume: Evidence from high-frequency data,” *Working Paper*.
- 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. (2019): “Monetary policy and the redistribution channel,” *American Economic Review*, 109, 2333–67.
- 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. ROGNLIE, AND L. STRAUB (2020): “Micro jumps, macro humps: Monetary policy and business cycles in an estimated HANK model.” *Working Paper*.
- (forthcoming): “The intertemporal Keynesian cross,” *Journal of Political Economy*.
- BALLEER, A., G. DUERNECKER, S. FORSTNER, AND J. GOENSCH (2022): “The Effects of Biased Labor Market Expectations on Consumption, Wealth Inequality, and Welfare,” *Working Paper*.
- BANKS, J. AND Z. OLDFIELD (2007): “Understanding pensions: Cognitive function, numerical ability and retirement saving,” *Fiscal Studies*, 28, 143–170.
- 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 (2023a): “The liquidity channel of fiscal policy,” *Journal of Monetary Economics*, 134, 86–117.
- (forthcoming): “Shocks, frictions, and inequality in US business cycles,” *American Economic Review*.
- BAYER, C., B. BORN, R. LUETTICKE, AND G. J. MÜLLER (2023b): “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 (2023c): “A HANK2 model of monetary unions,” *Working Paper*.
- BAYER, C., R. LÜTTICKE, L. PHAM-DAO, AND V. TJADEN (2019): “Precautionary savings, illiquid assets, and the aggregate consequences of shocks to household income risk,” *Econometrica*, 87, 255–290.
- BENIGNO, P. AND L. PACIELLO (2014): “Monetary policy, doubts and asset prices,” *Journal of Monetary Economics*, 64, 85–98.
- BERAJA, M. AND N. ZORZI (2024): “Durables and Size-Dependence in the Marginal Propensity to Spend,” *Working Paper*.
- BERNARD, R. (2023): “Mental accounting and the marginal propensity to consume,” *Working Paper*.
- 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., D. EVANS, M. GOLOSOV, AND T. J. SARGENT (2021): “Inequality, Business Cycles, and Monetary-Fiscal Policy,” *Econometrica*, 89, 2559–2599.
- BILBIIE, F. O. (forthcoming): “Monetary policy and heterogeneity: An analytical framework,” *Review of Economic Studies*.
- 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., 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.
- 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*.
- CAPLIN, A., V. GREGORY, E. LEE, S. LETH-PETERSEN, AND J. SÆVERUD (2024): “Subjective Earnings Risk,” *Working Paper*.
- 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.

- DALTON, P. S., S. GHOSAL, AND A. MANI (2016): “Poverty and aspirations failure,” *The Economic Journal*, 126, 165–188.
- DÁVILA, E. AND A. SCHAAB (2023a): “Optimal Monetary Policy with Heterogeneous Agents: Discretion, Commitment, and Timeless Policy,” *Working Paper*.
- (2023b): “Welfare Assessments with Heterogeneous Individuals,” *Working Paper*.
- DAVILA, J., J. H. HONG, P. KRUSELL, AND J.-V. RÍOS-RULL (2012): “Constrained efficiency in the neoclassical growth model with uninsurable idiosyncratic shocks,” *Econometrica*, 80, 2431–2467.
- DENGLER, T. AND B. GEHRKE (2023): “Short-time work and precautionary savings,” *Working Paper*.
- D’HAULTFOEUILLE, X., C. GAILLAC, AND A. MAUREL (2021): “Rationalizing rational expectations: Characterizations and tests,” *Quantitative Economics*, 12, 817–842.
- DOHMEN, T., A. FALK, D. HUFFMAN, AND U. SUNDE (2010): “Are risk aversion and impatience related to cognitive ability?” *American Economic Review*, 100, 1238–1260.
- D’ACUNTO, F., D. HOANG, M. PALOVIITA, AND M. WEBER (2019): “Cognitive abilities and inflation expectations,” *AEA Papers and Proceedings*, 109, 562–66.
- (2023a): “Cognitive Constraints and Economic Incentives,” *Working Paper*.
- (2023b): “IQ, expectations, and choice,” *The Review of Economic Studies*, 90, 2292–2325.
- EHRLINGER, J., K. JOHNSON, M. BANNER, D. DUNNING, AND J. KRUGER (2008): “Why the unskilled are unaware: Further explorations of (absent) self-insight among the incompetent,” *Organizational Behavior and Human Decision Processes*, 105, 98–121.
- 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.
- FARHI, E. AND I. WERNING (2019): “Monetary policy, bounded rationality, and incomplete markets,” *American Economic Review*, 109, 3887–3928.
- 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.
- GENICOT, G. AND D. RAY (2017): “Aspirations and inequality,” *Econometrica*, 85, 489–519.
- (2020): “Aspirations and economic behavior,” *Annual Review of Economics*, 12, 715–746.
- 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.
- GROHMANN, A., L. MENKHOFF, C. MERKLE, AND R. SCHMACKER (2023): “Earn more tomorrow: Overconfidence, income expectations and consumer indebtedness,” *Working Paper*.
- GUERREIRO, J. (2023): “Belief Disagreement and Business Cycles,” *Working Paper*.

- HAGEDORN, M., J. LUO, I. MANOVSKII, AND K. MITMAN (2019): “Forward guidance,” *Journal of Monetary Economics*, 102, 1–23.
- HANUSHEK, E. A. AND L. WOESSMANN (2008): “The role of cognitive skills in economic development,” *Journal of Economic Literature*, 46, 607–668.
- 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.
- HEIDHUES, P., B. KŐSZEGI, AND P. STRACK (2018): “Unrealistic expectations and misguided learning,” *Econometrica*, 86, 1159–1214.
- HOFFMAN, M. AND S. V. BURKS (2020): “Worker overconfidence: Field evidence and implications for employee turnover and firm profits,” *Quantitative Economics*, 11, 315–348.
- HOLM, M. B., P. PAUL, AND A. TISCHBIREK (2021): “The transmission of monetary policy under the microscope,” *Journal of Political Economy*, 129, 2861–2904.
- 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.
- JAPPELLI, T. AND L. PISTAFERRI (2010): “The consumption response to income changes,” *Annual Review of Economics*, 2, 479–506.
- JENSEN, A. R. (1998): “The g Factor: The Science of Mental Ability,” *Westport, CT: Praeger*.
- JORDÀ, Ò., K. KNOLL, D. KUVSHINOV, M. SCHULARICK, AND A. M. TAYLOR (2019): “The rate of return on everything, 1870–2015,” *The Quarterly Journal of Economics*, 134, 1225–1298.
- 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.
- KEKRE, R. (2023): “Unemployment insurance in macroeconomic stabilization,” *Review of Economic Studies*, 90, 2439–2480.
- KEKRE, R. AND M. LENEL (2022): “Monetary policy, redistribution, and risk premia,” *Econometrica*, 90, 2249–2282.
- KRUEGER, D., K. MITMAN, AND F. PERRI (2016): “Macroeconomics and household heterogeneity,” in John Taylor and Harald Uhlig (eds) *Handbook of Macroeconomics*, Elsevier, vol. 2,

843–921.

- LAIBSON, D., P. MAXTED, AND B. MOLL (forthcoming): “Present bias amplifies the household balance-sheet channels of macroeconomic policy,” *Quarterly Journal of Economics*.
- LIAN, C. (2023): “Mistakes in future consumption, high MPCs now,” *American Economic Review: Insights*, 5, 563–581.
- LUETTICKE, R. (2021): “Transmission of monetary policy with heterogeneity in household portfolios,” *American Economic Journal: Macroeconomics*, 13, 1–25.
- LUSARDI, A. AND O. S. MITCHELL (2014): “The economic importance of financial literacy: Theory and evidence,” *American Economic Journal: Journal of Economic Literature*, 52, 5–44.
- MACLEOD, C. M. (1991): “Half a century of research on the Stroop effect: An integrative review.” *Psychological Bulletin*, 109, 163.
- 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. WOLF (2023): “Optimal Policy Rules in HANK,” *Working Paper*.
- MIYAKE, A. AND N. P. FRIEDMAN (2012): “The nature and organization of individual differences in executive functions: Four general conclusions,” *Current Directions in Psychological Science*, 21, 8–14.
- MOSCHINI, E., G. RAVEENDRANATHAN, AND M. XU (2023): “Optimism About Graduation and College Financial Aid,” *Working Paper*.
- 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.
- PATTERSON, C. (2023): “The matching multiplier and the amplification of recessions,” *American Economic Review*, 113, 982–1012.
- PFÄUTI, O. AND F. SEYRICH (2023): “A behavioral heterogeneous agent new keynesian model,” *Working Paper*.
- ROUWENHORST, K. G. (1995): “Asset Pricing Implications of Equilibrium Business Cycle Models,” in Thomas F. Cooley (ed) *Frontiers of Business Cycle Research*, Princeton University Press, 294–330.
- ROZSYPAL, F. AND K. SCHLAFMANN (2023): “Overpersistence bias in individual income expectations and its aggregate implications,” *American Economic Journal: Macroeconomics*, 15, 331–371.
- SEIDL, H. AND F. SEYRICH (2023): “Unconventional Fiscal Policy in a Heterogeneous-Agent New

- Keynesian Model,” *Journal of Political Economy Macroeconomics*, 1, 633–664.
- SERGEYEV, D., C. LIAN, AND Y. GORODNICHENKO (2024): “The Economics of Financial Stress,” *Working Paper*.
- SMIRNOV, D. (2022): “Optimal Monetary Policy in HANK,” *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).
- STANGO, V. AND J. ZINMAN (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*.
- WANG, T. (2023): “Perceived versus calibrated income risks in heterogeneous-agent consumption models,” *Working Paper*.
- WERNING, I. (2015): “Incomplete markets and aggregate demand,” *Working Paper*.
- WOLF, C. (forthcoming): “Interest Rate Cuts vs. Stimulus Payments: An Equivalence Result,” *Journal of Political Economy*.
- WOODFORD, M. (1990): “Public debt as private liquidity,” *The American Economic Review*, 80, 382–388.
- (2013): “Macroeconomic analysis without the rational expectations hypothesis,” *Annual Review of Economics*, 5, 303–346.
- (2019): “Monetary policy analysis when planning horizons are finite,” *NBER Macroeconomics Annual*, 33, 1–50.
- YANG, Y. (2022): “Redistributive inflation and optimal monetary policy,” *Working paper*.
- ZIMMERMANN, F. (2020): “The dynamics of motivated beliefs,” *American Economic Review*, 110, 337–363.

A Additional Results

Table A1: Subjective financial condition forecasts are strongly positively correlated with income forecasts

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

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

Table A2: Household financial condition forecasts and forecast errors tilt optimistic

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

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

Table A3: Household financial condition forecast errors are persistent

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

Note: Sample is 10,546 forecast error pairs from 2,469 panelists. Sample is smaller here than in Appendix Table A2 because here we require ≥ 2 forecast-realization pairs per panelist and only include realizations of "about the same", to allow for the sharpest feasible test of persistence, by holding realizations constant and allowing for forecast errors in either direction (thereby minimizing measurement error from censoring).

Table A4: Household financial condition forecast learning?

Panel A. First forecast - realization pair	Realization this year			
<u>Forecast last year</u>	Better	Same	Worse	Total
Better	0.09	0.16	0.06	0.31
Same	0.05	0.40	0.12	0.57
Worse	0.01	0.05	0.06	0.12
Total	0.15	0.61	0.23	1
Panel B. Last forecast - realization pair	Realization this year			
<u>Forecast last year</u>	Better	Same	Worse	Total
Better	0.10	0.13	0.04	0.28
Same	0.06	0.46	0.09	0.61
Worse	0.01	0.05	0.06	0.11
Total	0.17	0.65	0.18	1

Note: Sample includes only the 3073 panelists with multiple forecast-realization pairs.

Table A5: Pairwise correlations between persistent optimism about financial condition and HtM measures, using all data for non-SZ modules

	Proportion optimistic forecast errors						Row variable pop. share	
	1=(≥ 0.5)		1=(> 0.5)		Unw.	Weighted	Unw.	Weighted
	Unw.	Weighted	Unw.	Weighted				
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	
1=(lives paycheck-to-paycheck c. 2012)	0.14	0.21	0.14	0.13	0.14	0.17	0.48	0.50
s.e.	0.05	0.07	0.05	0.07	0.04	0.05	0.02	0.02
N	1068	1068	1068	1068	1068	1068	1068	1068
Lives paycheck-to-paycheck, COVID era	0.19	0.16	0.17	0.11	0.15	0.10	0.38	0.39
s.e.	0.04	0.05	0.04	0.05	0.03	0.03	0.02	0.02
N	1086	1086	1086	1086	1086	1086	1086	1086
1=(Lacks precautionary savings in 2012 and 2018)	0.34	0.32	0.34	0.31	0.30	0.27	0.36	0.39
s.e.	0.05	0.07	0.05	0.07	0.04	0.05	0.02	0.02
N	864	864	864	864	864	864	864	864
1=(Lacks precautionary savings in 2012 or 2018)	0.36	0.34	0.39	0.33	0.36	0.35	0.58	0.62
s.e.	0.05	0.06	0.05	0.07	0.04	0.05	0.02	0.02
N	864	864	864	864	864	864	864	864
Difficulty covering \$2k emergency expense	0.17	0.14	0.19	0.15	0.16	0.12	0.48	0.52
s.e.	0.03	0.04	0.03	0.04	0.02	0.03	0.01	0.02
N	2480	2480	2480	2480	2480	2480	2480	2480

Note: Here we combine all the data we have on potentially optimistic financial condition forecast errors and HtM measures. Weighted estimates use the mean sampling weight across all financial condition realizations per panelist.

Table A6: Pairwise correlations between persistent optimism about financial condition and patience or risk aversion

	Patience		RA: lotteries		RA: scale	
	Unw. (1)	Weighted (2)	Unw. (3)	Weighted (4)	Unw. (5)	Weighted (6)
1=(Prop. optimistic FCEs>0.5)	-0.05	-0.11	-0.05	-0.12	-0.07	-0.20
s.e.	0.07	0.13	0.06	0.10	0.05	0.09
N	447	447	468	468	465	465
1=(Prop. optimistic FCEs \geq 0.5)	-0.01	-0.01	-0.06	-0.12	-0.06	-0.15
s.e.	0.07	0.14	0.06	0.10	0.05	0.09
N	447	447	468	468	465	465
Prop. optimistic forecast errors	-0.12	-0.13	-0.09	-0.15	-0.05	-0.16
s.e.	0.07	0.14	0.06	0.11	0.05	0.08
N	447	447	468	468	465	465

Notes: Persistent optimism measures based on panelists with multiple potentially optimistic forecast errors (see Section 2.1 for details). Patience is the average savings rate across 24 convex time budget choices ([Andreoni and Sprenger, 2012](#)). Risk aversion (RA) is based on the [Barsky et al. \(1997\)](#) lifetime income gamble elicitation (Columns 3 and 4) or the [Dohmen et al. \(2010\)](#) financial risk-taking scale (Columns 5 and 6). Weighted estimates use sampling probability from the last SZ module. We use Obviously Related Instrumental Variables to account for measurement error in the column variables by using the two measures of each (taken in 2014 and 2017) to instrument for each other ([Gillen et al., 2019](#); [Stango and Zinman, 2023](#)).

Table A7: Pairwise correlations between persistent HtM measures and patience or risk aversion

	Patience		RA: lotteries		RA: scale	
	Unw. (1)	Wtd. (2)	Unw. (3)	Wtd. (4)	Unw. (5)	Wtd. (6)
1=(Persistent severe financial distress)	-0.01	-0.08	0.04	0.08	0.11	0.03
s.e.	(0.06)	(0.14)	(0.05)	(0.12)	(0.04)	(0.09)
N	780	780	832	832	818	818
1=(Persistent low net worth)	-0.03	-0.07	0.14	0.03	0.06	0.08
s.e.	(0.06)	(0.10)	(0.05)	(0.09)	(0.04)	(0.07)
N	734	734	778	778	765	765
1=(paycheck-to-paycheck c. 2012)	0.06	0.38	0.05	-0.16	0.01	0.07
s.e.	(0.10)	(0.17)	(0.09)	(0.31)	(0.07)	(0.16)
N	233	233	260	260	256	256
paycheck-to-paycheck, COVID era	-0.13	-0.01	0.13	0.01	0.08	0.05
s.e.	(0.07)	(0.12)	(0.06)	(0.10)	(0.05)	(0.08)
N	493	493	519	519	516	516
1=(Lacks prec. saving in 2012 or 2018)	-0.22	-0.19	0.07	-0.08	0.11	0.05
s.e.	(0.08)	(0.13)	(0.08)	(0.14)	(0.07)	(0.11)
N	254	254	269	269	264	264
Difficult covering \$2k emerg. expenses	-0.15	-0.04	0.11	0.13	0.14	0.15
s.e.	(0.07)	(0.12)	(0.06)	(0.11)	(0.05)	(0.08)
N	462	462	491	491	487	487

Note: Patience is the average savings rate across 24 convex time budget choices ([Andreoni and Sprenger, 2012](#)). Risk aversion is based on the the [Barsky et al. \(1997\)](#) lifetime income gamble elicitation (Columns 3 and 4) or the [Dohmen et al. \(2010\)](#) financial risk-taking scale (Columns 5 and 6). Weighted estimates use sampling probability from the last SZ module. We use Obviously Related Instrumental Variables to account for measurement error in the column variables by using the two measures of each (taken in 2014 and 2017) to instrument for each other ([Gillen et al., 2019](#); [Stango and Zinman, 2023](#)). HtM measures are detailed in Section 2. Weighted estimates use the sampling probability for the last SZ module.

Table A8: Pairwise correlations between overconfidence and patience or risk aversion

	Patience		RA: lotteries		RA: scale	
	Unwtd.	Weighted	Unwtd.	Weighted	Unwtd.	Weighted
	(1)	(2)	(3)	(4)	(5)	(6)
1=Oc both rounds	0.04	-0.01	0.16	0.24	-0.08	-0.20
s.e.	(0.06)	(0.14)	(0.05)	(0.12)	(0.04)	(0.07)
N	758	758	807	807	813	813
Oc percentile rank	0.00	-0.01	0.24	0.31	-0.15	-0.32
s.e.	(0.07)	(0.12)	(0.06)	(0.12)	(0.05)	(0.08)
N	758	758	807	807	813	813


Notes: See Section 2 for details on overconfidence measures. Patience is the average savings rate across 24 convex time budget choices [Andreoni and Sprenger \(2012\)](#). Risk aversion is based on the [Barsky et al. \(1997\)](#) lifetime income gamble elicitation (Columns 3 and 4) or the [Dohmen et al. \(2010\)](#) financial risk-taking scale (Columns 5 and 6). Weighted estimates use sampling probability from the last SZ module. We use Obviously Related Instrumental Variables to account for measurement error in the column variables, and in overconfidence percentile rank, by using the two measures of each (taken in 2014 and 2017) to instrument for each other ([Gillen et al., 2019](#); [Stango and Zinman, 2023](#)).

Table A9: Persistent overconfidence: Correlations with cognitive skills

	1 = oc both rounds		oc percentile rank	
	Unweighted	Weighted	Unweighted	Weighted
	(1)	(2)	(3)	(4)
<u>Cognitive skill measures</u>				
<u>Summary: 1st common factor</u>	-0.64	-0.63	-0.77	-0.74
s.e.	0.03	0.05	0.04	0.06
N	817	817	817	817
<u>Summary: 1st principal component</u>	-0.55	-0.54	-0.82	-0.83
s.e.	0.03	0.05	0.03	0.05
N	733	733	733	733
<u>Component: Fluid intelligence</u>	-0.72	-0.73	-1.05	-1.07
s.e.	0.03	0.05	0.03	0.06
N	817	817	817	817
<u>Component: Numeracy</u>	-0.36	-0.45	-0.57	-0.66
s.e.	0.04	0.07	0.05	0.08
N	798	798	798	798
<u>Component: Financial literacy</u>	-0.32	-0.24	-0.47	-0.36
s.e.	0.04	0.09	0.04	0.09
N	813	813	813	813
<u>Component: Executive function</u>	-0.32	-0.41	-0.44	-0.60
s.e.	0.05	0.07	0.05	0.09
N	749	749	749	749

Note: Overconfidence re: relative performance in a cognitive skills test (see Section 2 for details). All cognitive skills measures are percentile ranks. of each of the component measures shown in the table (see [Stango and Zinman \(2023\)](#) for details on component measures). Weighted estimates use the sampling probability for the last SZ module. All cognitive skills measures, and overconfidence percentile rank, use Obviously Related Instrumental Variables to account for measurement error by having the two rank measures (taken in 2014 and 2017) instrument for each other ([Gillen et al. \(2019\)](#), [Stango and Zinman \(2023\)](#)). We do not take the same approach to the overconfidence indicator in Columns (1) and (2), because measurement error-IV does not work well on misclassification error. An IV point estimate of a correlation can exceed $|1|$, as it does in two instances here.

Figure A1: Distribution of HtM along the income distribution



Note: This figure shows the share of hand-to-mouth households along the income distribution for our "low net worth" HtM measure in the data (black dashes). It also shows the share of HtM households in our baseline model with overconfidence (blue dots) and in the standard HANK model recalibrated to match the average MPC of our baseline model (red squares). We redefine the HtM measure in both models such that the aggregate HtM share matches the data.

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:

$$\beta \tilde{E}_t [u'(c_{t+1}^1)] = \hat{\beta} E_t [u'(\hat{c}_{t+1}^1)], \quad (16)$$

where a hat " $\hat{\cdot}$ " denotes the economy with heterogeneity in patience. Similarly, for household 2:

$$\beta \tilde{E}_t [u'(c_{t+1}^2)] = \hat{\beta} E_t [u'(\hat{c}_{t+1}^2)], \quad (17)$$

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 (16) implies

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

(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

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

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.⁴³ 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. ■

⁴³A 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 (16) implies a discount factor ratio of 0.86 whereas equation (16) implies a ratio of 0.75.