

DISCUSSION PAPER SERIES

DP15934

Information and Wealth Heterogeneity in the Macroeconomy

Tobias Broer, Alexandre Kohlhas, Kurt Mitman and
Kathrin Schlafmann

MONETARY ECONOMICS AND FLUCTUATIONS

CEPR

Information and Wealth Heterogeneity in the Macroeconomy

Tobias Broer, Alexandre Kohlhas, Kurt Mitman and Kathrin Schlafmann

Discussion Paper DP15934

Published 17 March 2021

Submitted 15 March 2021

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Monetary Economics and Fluctuations

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Tobias Broer, Alexandre Kohlhas, Kurt Mitman and Kathrin Schlafmann

Information and Wealth Heterogeneity in the Macroeconomy

Abstract

We document systematic differences in macroeconomic expectations across U.S. households and rationalize our findings with a theory of information choice. We embed this theory into an incomplete-markets model with aggregate risk. Our model is quantitatively consistent with the pattern of expectation heterogeneity in the data. Relative to a full-information counterpart, our model implies substantially increased macroeconomic volatility and inequality. We show through the example of a wealth tax that neglecting the information channel leads to erroneous conclusions about the effects of policies. While in the model without information choice a wealth tax reduces wealth inequality, in our framework it reduces information acquired in the economy, leading to increased volatility and higher wealth inequality in equilibrium.

JEL Classification: D84, E21, E27, E62

Keywords: Heterogenous information, unemployment, incomplete markets, precautionary savings

Tobias Broer - tobias.broer@iies.su.se
IIES Stockholm University and CEPR

Alexandre Kohlhas - alexandre.kohlhas@iies.su.se
Institute for International Economic Studies

Kurt Mitman - kurt.mitman@iies.su.se
Stockholm University and CEPR

Kathrin Schlafmann - ksc.fi@cbs.dk
Copenhagen Business School, Institute for International Economic Studies and CEPR

Acknowledgements

For comments and useful suggestions, we thank Per Krusell, Kristoffer Nimark, Alessandra Peter, Mirko Wiederholt, and seminar and conference participants at the IIES, Bern, PSE, Oslo, SED, and the Econometric Society European Meeting 2020. This research received financial support from Ragnar Soderberg Stiftelsen. Mitman thanks the European Research Council (ERC Starting Grant 759482) for financial support. This paper was previously presented under the title "Heterogenous Information Choice in General Equilibrium."

Information and Wealth Heterogeneity in the Macroeconomy*

Tobias Broer[†]

Alexandre N. Kohlhas[‡]

Kurt E.W. Mitman[§]

Kathrin Schlafmann[¶]

March 2021

Abstract

We document systematic differences in macroeconomic expectations across U.S. households and rationalize our findings with a theory of information choice. We embed this theory into an incomplete-markets model with aggregate risk. Our model is quantitatively consistent with the pattern of expectation heterogeneity in the data. Relative to a full-information counterpart, our model implies substantially increased macroeconomic volatility and inequality. We show through the example of a wealth tax that neglecting the information channel leads to erroneous conclusions about the effects of policies. While in the model without information choice a wealth tax reduces wealth inequality, in our framework it reduces information acquired in the economy, leading to increased volatility and higher wealth inequality in equilibrium.

Keywords: Heterogenous information, unemployment, incomplete markets, precautionary savings

JEL Codes: D84, E21, E27, E62

*For comments and useful suggestions, we thank Per Krusell, Kristoffer Nimark, Alessandra Peter, Mirko Wiederholt, and seminar and conference participants at the IIES, Bern, PSE, Oslo, SED, FRB Minneapolis, CES-iFo Conference on Survey Data, and the Econometric Society European Meeting 2020. This research received financial support from Ragnar Soderberg Stiftelsen. Mitman thanks the European Research Council (ERC Starting Grant 759482) for financial support. This paper was previously presented under the title “Heterogenous Information Choice in General Equilibrium.”

[†]Paris School of Economics, Institute for International Economic Studies and CEPR. Email: tobias.broer@iies.su.se.

[‡]Institute for International Economic Studies. Email: alexandre.kohlhas@iies.su.se.

[§]Institute for International Economic Studies, CEPR and IZA. Email: kurt.mitman@iies.su.se.

[¶]Copenhagen Business School, Institute for International Economic Studies, Danish Finance Institute and CEPR. Email: ksc.fi@cbs.dk.

1 Introduction

The aftermath of the Great Recession brought inequality to the forefront of macroeconomic research. State-of-the-art models used by macroeconomists and policymakers now incorporate rich heterogeneity in income, wealth and employment status.¹ Recent empirical work has highlighted another central dimension along which households differ — in their expectations (e.g., Angeletos et al., 2020; Carroll, 2003; D’Acunto et al., 2019). The accuracy of those heterogeneous expectations correlates systematically with household wealth and labor market status (Vissing-Jorgensen, 2003). As a result, the interaction between household expectations and financial and employment status has potential implications for aggregate dynamics and the impacts of policies.

Workhorse macroeconomic models, however, maintain the strong assumption that households have full information about the state of the economy. Current models are thus incapable of answering questions such as: What are the aggregate effects of dispersed expectations? How do they depend on, and interact with, household heterogeneity and inequality in general equilibrium? How does the presence of heterogeneous, dispersed expectations change the efficacy of policy? In this paper, we develop a new framework to answer these questions.

To start, we document the extent of heterogeneity in household expectations in US micro data. We show that in a leading household survey (FRB NY Survey of Consumer Expectations) both the mean and reported uncertainty of stated forecasts of key macroeconomic variables differ substantially among households. Importantly, we document that the accuracy of household expectations is systematically related to wealth and employment status. In particular, all else equal, wealthier households have more accurate beliefs. But unlike the evidence in Vissing-Jorgensen (2003) and others this relationship is far from monotone in wealth — especially at the lower-end of the wealth distribution.

Motivated by this evidence, we embed dynamic optimal information choice into an otherwise standard business cycle model with incomplete markets. Households have to form expectations about future returns, wages and unemployment risk, and can acquire information about aggregate productivity subject to a fixed cost. The gains to acquiring this information depend

¹See Krueger et al. (2016) and Kaplan and Violante (2018) for recent surveys on macroeconomics and household heterogeneity.

on household wealth, employment status, and prior beliefs, leading to systematic differences in household expectations. Using this framework, we characterize the distribution of household expectations, and identify its consequences for wealth inequality, aggregate dynamics and the effects of policies.

Our framework builds on important work on the consequences of optimal information choice (e.g., Grossman and Stiglitz, 1980; Barlevy and Veronesi, 2000; Sims, 2003; Hellwig and Veldkamp, 2009; Mackowiak and Wiederholt, 2009; Veldkamp, 2011; Maćkowiak et al., 2018b). That literature has primarily restricted itself to studying the implications of once-and-for-all information choices that are identical across time and decision makers. The contribution of our paper is, in this context, to highlight the macroeconomic consequences of *dynamic, heterogeneous information choices*, and to quantify how they can profoundly shape macroeconomic outcomes.²

We show that differences in wealth and employment status naturally imply differences in information that make expectations consistent with the survey data. To understand the information acquisition decision, it is instructive to understand household savings decisions.

Consider first unemployed households, who dissave to smooth consumption. They have highly non-linear policy functions at low levels of wealth—leading down to the borrowing constraint (Carroll, 1997). At the borrowing constraint, households are hand-to-mouth, and thus do not value improved information, so never acquire it. As wealth rises, the marginal utility of unemployed households is still high, so savings mistakes are costly. Furthermore, because of the non-linearity of the policy function, uninformed savings can lead to large mistakes. Information, therefore, becomes very valuable. When wealth increases further, marginal utility falls (savings policy functions become approximately linear) since the household is no longer at risk of hitting the constraint due to a mistake. Mistakes are also smaller since the savings decision is less sensitive to the contemporaneous job-finding rate. The value of information thus falls. At the same time, as wealth rises further, two forces induce more information acquisition which counteracts this decrease: i) the cost of acquiring information relative to current

²Our study also contributes to the literature on dynamic macroeconomic models with heterogeneous expectations (see Branch and McGough (2018) for a survey). In particular, our work is similar in spirit to that which allows for differences in expectations due to random utility shocks (e.g., Brock and Hommes (1997), Brock and Hommes (1998)). In contrast to this line of work, the expectations in our framework arise from households' optimal information choices in an environment where decision rules do not aggregate.

wealth falls; and ii) the benefit of accurately predicting returns rises with the amount of savings. Eventually, these forces dominate, leading to upward-sloping probabilities of acquiring information at high levels of wealth.

Employed households always have (relatively) low marginal utility so the value of information starts off low and rises with wealth. Importantly, the ability of our framework to match the data crucially depends on precautionary effects tied to non-linear decision rules. This contrasts with previous analyses that, for tractability reasons, instead focus on linearized decision rules.³

We then show how that heterogeneous information choices substantially change the equilibrium properties of our economy relative to the full-information rational expectations benchmark. On the micro side, heterogeneous information choices feed back into wealth and income inequality in the economy, as differently informed households make disparate savings choices. As a result of this two-sided feedback, the introduction of household information choice exacerbates inequality. In particular, poor households with low information are unable to exploit periods of good labor market prospects and high returns to build up their financial wealth.

On the macro side, the presence of uninformed households leads to an increase in aggregate volatility. Under full information, household savings are pro-cyclical, but as the aggregate capital stock rises in booms, the return on savings falls, dampening the savings response. By contrast, uninformed households expectations about returns are sluggish, which makes household savings more pro-cyclical and the economy more volatile. That mechanism is dampened by increased information acquisition, because the benefits of information about the current state of the economy are higher when the economy is more volatile. These dynamics elucidate a more general feature of our framework: information acquisition decisions are *strategic substitutes*.⁴ In equilibrium, not all households acquire information in every period, leading to ten percent higher fluctuations in consumption and output relative to full information.

Next, we show that macroeconomic policies may have substantially different effects when

³In concurrent and related work, Auclert et al. (2020) and Carroll et al. (2020) analyze a heterogeneous-agent model with imperfect information, similar to that which we consider. However, unlike in our analysis, they assume an exogenous process for household information based on Mankiw and Reis (2002); Carroll (2003) (and hence do not allow households to make any information choices) and linearize the model.

⁴In a companion paper, Broer et al. (2021), we show how this feedback may lead to non-existence of a representative-agent equilibrium, and characterize the interaction between information choice, inequality and macro-dynamics more generally, using an analytical framework.

one accounts for heterogeneous household information choice. To do so, we consider the example of a wealth tax. In particular, we introduce a one percent per annum wealth tax on households. The direct impact of the tax is to reduce the average wealth of households. The indirect effect is to reduce information acquisitions by 30 percent in any quarter, as information acquisition on average rises in wealth. By reducing the information content in the economy, economic volatility increases by nine percent for the reasons discussed above. In contrast, in the full information case, the tax has virtually no impact on volatility, despite a similar fall in average wealth.

The effect of the wealth tax on household inequality in our benchmark economy is perhaps even more surprising: a one percent tax increases inequality by over two percent, whereas in the full information case the tax meaningfully decreases inequality. The reason inequality increases is because of the increased volatility, leading to larger overaccumulation of savings for uninformed, high wealth households. Our framework therefore also provides an explanation for why several countries did not see increases in wealth inequality following the abolition of previously instated wealth taxes (e.g., Jakobsen et al., 2020). Clearly, there may be alternative drivers of these lack of increases in inequality; however, our results do suggest that the effects of dynamic, heterogeneous information choice may substantially alter the relative costs and benefits of macroeconomic policies in unexpected directions. Thus, our findings imply a Lucas-type critique (Lucas Jr, 1976) for policy evaluation in full-information economies.

Finally, two wider implications of our theory are worth noting. First, in our analysis we, for simplicity, abstract from any behavioral drivers of information choice (e.g., Bordalo et al., 2016; Gabaix, 2019; Bordalo et al., 2017). Notwithstanding such alternative drivers, we show that households' rational incentives to acquire different information fundamentally alter the consequences of redistributive macroeconomic policies. We conjecture that behavioral heuristics, salience effects, and other behavioral drivers of agents' information choice would only increase the gap between the predictions of standard models and those relevant for macroeconomic policy.

Second, because of the complexity of computing rational expectations equilibria in neoclassical heterogeneous-agent economies, several authors have proposed dimensionality reduction methods. Most notably, Krusell and Smith (1998) propose constraining households to only form their expectations based on a limited set of moments. Through this lens, our approach

is to allow household themselves to decide which variables or moments to use to forecast the future. In this sense, our framework presents a natural evolution of the Krusell and Smith (1998) computational approach.

The rest of the paper proceeds as follows: Section 2 summarizes key patterns of expectations in US data. Section 3 presents a model of dynamic information choice in an environment with aggregate and idiosyncratic income risk and incomplete markets. Section 4 discusses the calibration of our model, while Section 5 presents our benchmark quantitative results and Section 6 studies the introduction of a wealth tax. We conclude in Section 7. An appendix contains a discussion of the related literature, as well as additional results and analysis.

2 Motivating Evidence

This section documents stylized facts about households’ expectations. Specifically, we analyze a household panel of economic expectations from the *Survey of Consumer Expectations* (SCE). The SCE is a monthly panel of 1,300 US households that began in 2013, in which households are asked to report their point and density forecasts for several macroeconomic and financial variables.⁵ Using the SCE data, we compute two types of forecast errors: 1) compared to realized macroeconomic outcomes; and 2) compared to professional forecasts from the *Survey of Professional Forecasters* (SPF), the main survey of macroeconomic expectations in the US.

We focus our analysis on variables most pertinent for households’ consumption and savings choices—unemployment and inflation—as those reflect both income and return risk that we later include in our structural framework. Further, those variables can be directly compared to forecasts from the SPF. We define a survey respondent’s forecast error as the difference between the realized outcome or the SPF value and the respondent’s forecast.⁶ A negative forecast error therefore corresponds to an over-estimate of the variable. Our analysis yields

⁵Armantier et al. (2017) provide details on the construction and scope of the SCE. Appendix 2 describes in further detail the data that we use in this section.

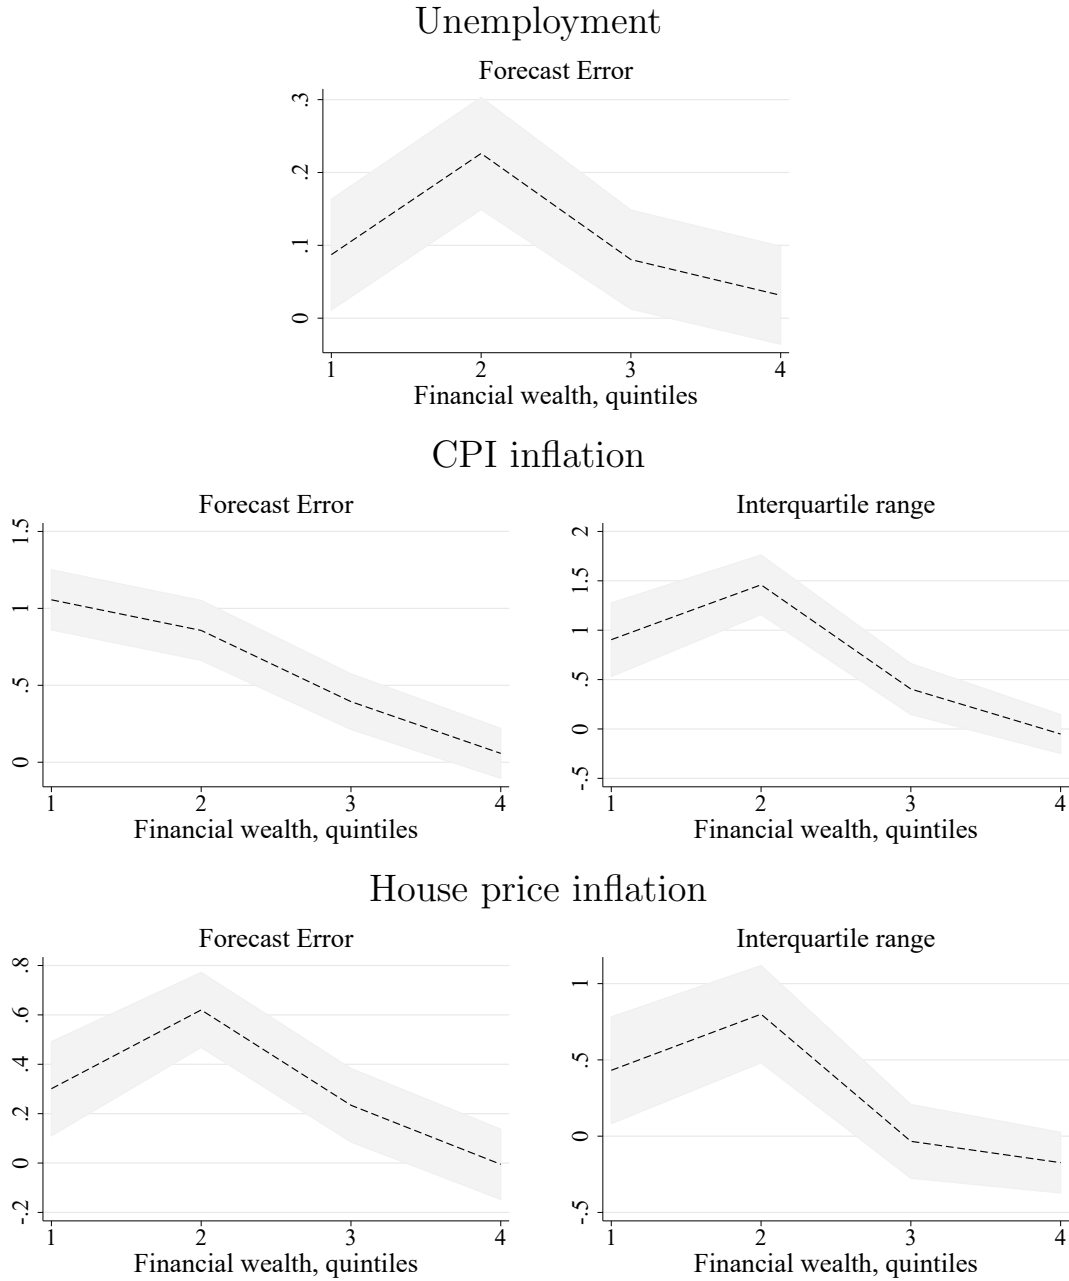
⁶Specifically, we use the year-over-year change in the US consumer price index and the S&P/Case-Shiller 20-City Composite Home Price Index, respectively. For CPI inflation, we compute forecast errors as the difference between the first-release outcome one-year ahead and respondents’ forecasts. For unemployment expectations, which are reported as the probability of rising unemployment over the next year, we define errors as the relative difference between a household’s forecasted probability and the average equivalent probability by professional forecasters. This reflects ample evidence that professional forecasters provide more accurate predictions than even those from modern statistical and economic models (e.g., Bhandari et al., 2019; Croushore and Stark, 2019). We therefore use professional forecasters stated probability as a proxy for the true, underlying probability. See Appendix A.1 for further details.

four stylized facts about economic expectations of US households:

1. *Household expectations are substantially less accurate than professional forecasts.* The median absolute inflation forecast error of households is more than twice as large as that of professional forecasters—equal to 1.8 and 0.7 percentage points (pp) for CPI inflation, respectively. The larger median forecast errors of households carry over to their unemployment forecasts.
2. *Household expectations are substantially more uncertain than professional forecasts.* When asked for their probability distribution over possible inflation realizations, households report substantially wider distributions. The median of the interquartile ranges of *individual* forecast distributions is more than triple that of professional forecasters—2.0pp vs. 0.6pp, respectively.
3. *Household expectations are more heterogeneous than professional forecasts.* The dispersion in household point forecasts for inflation and unemployment is substantially higher than professional forecasts. For example, the standard deviation of point forecasts of CPI inflation across households is about three times larger than across professionals.
4. *The accuracy and uncertainty of household expectations are inverse-U shaped in wealth.* We plot the average absolute forecast errors and the median household uncertainty by wealth quintile in Figure 1. Even controlling for demographics and labor force status, the wealthiest households make the smallest errors and are more certain of their forecasts. However, errors and uncertainty are non-monotone in wealth—in contrast to the results in Vissing-Jorgensen (2003). Apart from the forecasts of CPI inflation, all measures of accuracy and uncertainty are higher for households in the second wealth quintile than for the poorest households.

For the complete details of the construction of forecast errors and our empirical specification and controls, we refer the reader to Appendix 2. We now proceed by developing a theory of information choice that can rationalize the aforementioned stylized facts. We then embed that theory into an otherwise standard heterogeneous-agent incomplete-markets model to quantify the impact of household information acquisition for positive and normative questions.

Figure 1: Household expectations across the wealth distribution



Coefficients from the regression of forecast errors for outcome y , ν_{it}^y , on individual covariates (X_{it}) and their quintile in the wealth distribution \mathbb{Q}_i^W (where the top quintile is the omitted category), and time dummies η_t . The forecast error for outcome y is $\nu_{it}^y = \frac{P_i(y_{t+12} > y_t | t) - P_{SPF}(y_{t+12} > y_t | t)}{P_{SPF}(y_{t+12} > y_t)}$. The estimated regression is $\nu_{it}^y = \delta_W' \mathbb{Q}_i^W + \beta' X_{it} + \eta_t + \epsilon_{it}$. The figures plot the coefficients δ_W with 95% confidence intervals. Standard errors are clustered at the individual level. The full results of the regression are reported in Appendix Table V. The top panel shows estimates for the errors in individual unemployment forecasts (elicited as the “percent chance that 12 months from now the unemployment rate in the U.S. will be higher than it is now”) relative to the equivalent consensus forecast from professional forecasters (see Appendix 2 for detail). Row 2 presents estimates for absolute forecast errors (left panel) and the interquartile ranges (right panel) of individual forecast distributions for 12-month-ahead consumer price inflation. Row 3 presents equivalent results for house price (HP) inflation. Forecast errors for inflation equal the absolute difference between forecasts and outcomes, as detailed in Appendix 2.

3 Model Framework

In this section, we describe a basic incomplete-markets model with aggregate risk. The model closely follows the environment in Krusell and Smith (1998), but with a modified information structure. In particular, we assume that every period households have the option to acquire imperfect information about the state of the economy.

3.1 Households

The economy consists of a continuum of households of unit mass. Household preferences are described by the utility function:

$$\mathcal{U} = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{c_t^{1-\gamma} - 1}{1-\gamma} - \kappa(\mathcal{I}_t) \right], \quad (1)$$

where $\beta \in (0, 1)$ denotes the discount factor, c_t non-durable consumption at time t , κ_t denotes a utility cost of acquiring information \mathcal{I}_t , and $\gamma > 0$. The utility cost κ_t is distributed according to a type-I extreme value distribution with parameter α^κ , and i.i.d. across individuals and time.⁷ Each household is endowed with one unit of time, which it supplies inelastically to the labor market. Labor productivity ϵ_t is stochastic and can take on two values $\epsilon_t \in \{0, 1\}$, which we interpret as unemployment and employment, respectively. We assume that ϵ_t follows a two-state, first-order Markov process $\Pi_{z_{t+1}, \epsilon_{t+1} | z_t, \epsilon_t}$, which depends both ϵ_t and aggregate total factor productivity z_t (described below).⁸ A household earns wage w_t when employed and receives unemployment benefits μw_t when unemployed, where $\mu \in (0, 1)$. We assume that households cannot borrow but can only save in physical capital k_t , whose net return equals $r_t - \delta$, where $\delta \in (0, 1)$ denotes the rate of depreciation on capital. Household consumption choices are restricted by the per-period budget constraint:

$$c_t + k_{t+1} + \nu_t(\mathcal{I}_t) = r_t k_t + (1 - \tau_t) \epsilon_t w_t + \mu(1 - \epsilon_t) w_t + (1 - \delta) k_t, \quad (2)$$

where $\nu_t(\mathcal{I}_t)$ denotes a monetary cost of acquiring information \mathcal{I}_t , and τ the tax rate on labor income. We refer to the right-hand side of (2) as the household's *cash-at-hand*, and denote it

⁷We interpret the extreme value shocks as preference shocks to information acquisition to explain unobserved heterogeneity in the data.

⁸We make this assumption to allow for the share of households in a given idiosyncratic employment state to only depend on current total factor productivity z_t .

by y_t . Households maximize utility (1) subject to the budget constraint (2).

3.2 Firms and Markets

The production sector consists of a representative competitive firm, which maximizes profits. Output is produced using a Cobb-Douglas technology that aggregates labor services and capital:

$$Y_t = z_t K_t^\alpha (L_t)^{1-\alpha}, \quad (3)$$

where K_t and L_t denote economy-wide capital and labor in period t , respectively. Aggregate total factor productivity is stochastic, and follows a first-order Markov process that takes two values $z_t \in \{Z_l, Z_h\}$. The firm rents capital and labor in perfectly competitive factor models. Firm optimization therefore implies that:

$$w_t = z_t(1 - \alpha) \left(\frac{K_t}{L_t}\right)^\alpha, \quad r_t = z_t\alpha \left(\frac{K_t}{L_t}\right)^{\alpha-1}. \quad (4)$$

3.3 Government Policy

In our baseline analysis, the only role that the government has is to run a balanced-budget unemployment insurance scheme, such that $\tau_t = \frac{\mu u_t}{L_t}$, where L_t and $u_t = 1 - L_t$ are the employment and unemployment rates in the economy, respectively.

3.4 Timeline

At the start of each period, idiosyncratic shocks $(\epsilon_t, \kappa_t)_i$ and aggregate shocks z_t realize. We index these by $i \in [0, 1]$ when necessary for clarity. Households then choose which signals \mathcal{I}_t to acquire about the current state of the economy from a maximum signal set \mathcal{I}_t^{max} . We assume that \mathcal{I}_t^{max} does not contain sufficient information for households to perfectly learn the current state of the economy, but that it can, for example, include current market signals such as prices.⁹ A household's information set Ω_t accumulates according to $\Omega_t = \{\Omega_{t-1}, \mathcal{I}_t\}$. Next, firms rent capital and labor, production takes place and factor payments are made. Finally, conditional on their information choice, households make consumption and savings choices.

⁹An alternative approach is to instead allow households to flexibly design their optimal signal subject to a cost (e.g., Maćkowiak et al., 2018a). Although this approach has several advantages, it is computationally intractable for the model that we study (cf. Section 3.5).

3.5 Recursive Formulation of the Household Problem

Given the timeline and informational assumptions above, we can develop a recursive formulation of the household problem. Let $S = (\Gamma, z)$, where Γ denotes the cross-sectional distribution of capital and employment status. We denote an individual household's first-order belief about S by $\mathcal{P}_i(S)$.¹⁰ Household i 's second-order belief about household j 's belief is referred to as $\mathcal{P}_{ij}(S)$, and so on ad infinitum. Individual household beliefs are summarized by: $p = \left\{ \mathcal{P}_i, (\mathcal{P}_{ij})_{j \in [0,1]}, \dots, (\mathcal{P}_{ij\dots k})_{j, \dots, k \in [0,1]^{n-1}}, \dots \right\}$. Let \mathcal{P} denote the set of all such beliefs $\mathcal{P} = \left\{ (\mathcal{P}_i)_{i \in [0,1]}, (\mathcal{P}_{ij})_{i, j \in [0,1]^2}, \dots, (\mathcal{P}_{ij\dots k})_{i, j, \dots, k \in [0,1]^n}, \dots \right\}$. The aggregate state of the economy can then be described by $\Sigma = (S, \mathcal{P})$, while the individual state variables are characterized by $\sigma = (y, \epsilon, p)$, where y denotes household cash-at-hand. We denote next period's realization of variable x by x' , while we denote previous period's realization by x_{-1} .

At the beginning of the period, households choose what information to acquire $\mathcal{I} \in \mathcal{I}^{max}$:

$$V(y, \epsilon, p_{-1}, \Sigma_{-1}) = \max_{\mathcal{I}} \mathbb{E} [W(y - \nu(\mathcal{I}), \epsilon, p, \Sigma) + \kappa(\mathcal{I}) \mid \Omega_{-1}] \quad (5)$$

where V and W denote the household value functions before and after information choice, respectively. Information acquisition entails both a utility cost κ and a monetary cost ν as a function of information choice \mathcal{I} . We note that households' expectations in the first stage are computed using previous period's posterior beliefs p_{-1} , and hence information. We assume that households rationally use the equilibrium law of motion for the aggregate state, which we denote by H , i.e., $\Sigma = H(\Sigma_{-1}, z, \mathcal{I})$, and the exogenous transition matrix Π^z to form a prior about today's state variables from yesterday's posterior.

The assumption of type-1 extreme value shocks for the utility cost of information acquisition implies a parsimonious logistic choice function $\mathbb{P}_{\mathcal{I}}(\cdot)$ over information:

$$\mathbb{P}_{\mathcal{I}}(y, \epsilon, p_{-1}, \Sigma_{-1}) = \frac{e^{W(y - \nu(\mathcal{I}), \epsilon, p, \Sigma)}}{\sum_{\mathcal{I} \in \mathcal{I}^{max}} e^{W(y - \nu(\mathcal{I}), \epsilon, p, \Sigma)}}, \quad (6)$$

¹⁰Not to be confused with the powerset, \mathcal{P}_i here has a distribution with $\hat{\Gamma}_i$ and productivity \hat{z}_i as its typical elements, representing household i 's first order belief for the cross-sectional distribution and level of productivity, respectively.

yielding the standard value function:

$$V(y, \epsilon, p_{-1}, \Sigma_{-1}) = \frac{\gamma^E}{\alpha^\kappa} + \frac{1}{\alpha^\kappa} \log \left(\sum_{\mathcal{I} \in \mathcal{I}^{max}} e^{W(y-v(\mathcal{I}), \epsilon, p, \Sigma)} \right) \quad (7)$$

where γ^E is the Euler-Mascheroni constant.

After deciding on information choices, households choose consumption c and savings k' :

$$W(y, \epsilon, p, \Sigma) = \max_{c, k' \geq 0} u(c) + \beta \mathbb{E} [V(y', \epsilon', p, \Sigma) \mid \Omega] \quad (8)$$

subj. to

$$c + k' = y$$

$$y' = r(\Sigma')k' + w(\Sigma')L'\epsilon + (1 - \delta)k'$$

where the expectation in (8) is taken with respect to today's updated information set Ω . We let g denote the function that characterizes a household's savings choice $k' = g(\sigma, \Sigma)$, and ι the function that characterizes its information choice $\mathcal{I} = \iota(\sigma_i)$. Finally, today's posterior beliefs p are linked to yesterday's p_{-1} through Bayes' Rule and the information choice \mathcal{I} .

Recursive Imperfect Information Competitive Equilibrium

The definition of a Recursive Competitive Imperfect Information Equilibrium (RIICE) straightforwardly extends the standard definition of a Recursive Competitive Equilibrium: A RIICE is a law of motion H , a pair of individual value functions V, W , policy functions ι and g , as well as pricing functions $(r(\Sigma), w(\Sigma))$ such that: (i) (V, W, ι, g) solves the household's two-stage problem, (ii) r and w satisfy firm maximization, (iii) H is generated by policy functions ι and g , the markov processes Π and Π^z , and Bayes' Rule, using the information contained in \mathcal{I} and current beliefs in \mathcal{P} .

3.6 Computational Strategy

We now outline our procedure for computing RIICE equilibria. Our description here is non-technical, and we include it in the main text as it is intimately linked to the two-way feedback mechanisms mentioned in the introduction that is at the heart of this paper. The endogenous aggregate state variables of the economy, Γ and \mathcal{P} , are high-dimensional objects. Even a *full-*

information incomplete markets model with aggregate risk therefore presents a computational challenge, because of the infinite-dimensionality of Γ (the endogenous state variable in that model). Our *imperfect information* framework has a double infinity—the additional complexity arising from the entire set of (higher-order) beliefs, \mathcal{P} . The standard strategy for computing such models *without* information choice involves approximating the distribution Γ with a finite set of moments $\mathbf{m} \equiv (\mathbf{m}_1, \mathbf{m}_2, \dots, \mathbf{m}_n)$ (Krusell and Smith, 1998). Accurately forecasting those moments enables households to forecast future prices, which are necessary for solving the household problem. One interpretation of that solution method is one of “boundedly rational” expectations, as households only keep track of a set of moments of the distribution. Importantly, in those solution methods, the set of information available to households is *exogenously* set. By contrast, in our equilibrium concept, households *optimally choose* the information on which to form their “boundedly rational” expectations. Thus, one can interpret our model as the natural framework to study incomplete markets models with aggregate risk, as we provide a micro-foundation for the boundedly rational solution based on costly information acquisition by households. In particular, the Krusell-Smith economy can be seen as a special case of our model in which the cost of information is zero. The Krusell-Smith economy thus assumes common knowledge of household beliefs over both z and the set of moments \mathbf{m} , whereas in our framework we relax this common knowledge assumption. The RIICE framework then allows for the three-way interaction between inequality, information acquisition and aggregate dynamics on which we focus on in this paper.

Our computational strategy is as follows: Households form priors over the contemporaneous realization of productivity z and over a set of moments of Γ given by $\bar{\mathbf{m}}$. Given those priors, using Bayes’ Rule and the equilibrium law of motion H , households can form expectations about the future path of the wage and return on capital, necessary to solve their maximization problem. Households can then choose to acquire information about any combination of productivity and the moments in \mathbf{m} , which we include in \mathcal{I}^{max} . If *all* households acquire information about all possible $\{z_t, \mathbf{m}\}$ in *every period*, our equilibrium coincides with the equilibrium concept from Krusell and Smith (1998), as discussed above.

3.7 Numerical Solution Procedure

Here we provide a brief overview of the numerical procedure implementing the foregoing algorithm. To compute the equilibrium, we use an iterative procedure to solve for the equilibrium fixed point: First, we postulate a law of motion H for the aggregate state variables. Second, we solve the household’s two-stage problem conditional on H and the cross-sectional distribution of information, income and wealth. Third, using the resulting individual decision rules, we simulate a large number of households for a long number of periods. From this simulation, we then calculate time series for z_t and \mathbf{m} , and estimate a new law of motion H' . We iterate until convergence on the aggregate law of motion.

4 Calibration

The aim of our calibration exercise is to ensure that the model can accurately account for standard business cycle facts, as well as capture the rich heterogeneity in household expectations documented in Section 2. The model is set in discrete time with each period corresponding to one quarter.

Externally Calibrated Parameters

We choose standard parameters for the capital share α (0.36), and the depreciation rate δ (0.025). Following Krueger et al. (2016), we calibrate the structure of aggregate and idiosyncratic risk to capture key features of the unemployment and job-finding rates in the post-World War II US economy. We define ”booms” and ”busts” based on the observed unemployment dynamics in the data, as those more closely aligned to our model framework than traditional NBER dated recessions.¹¹ The productivity variable z_t is calibrated to match the difference in average US total factor productivity during booms and busts. We estimate the persistence of booms and busts to be 0.88 and 0.82, respectively, and the ratio of productivity values to $z_h/z_l = 1.027$. The individual transition probabilities in labor productivity ϵ_{it} are set to match US labor market transitions. We choose an unemployment rate in booms and busts equal to six and 10 percent, respectively. Monthly job-finding rates are set to match unemployment

¹¹We use an HP filter with smoothing parameter equal to 14.400 to construct the trend in the unemployment rate at monthly frequency. Boom periods are defined as periods with a below trend unemployment rate.

spells in the data, equal to 55 and 45 percent in boom and busts, respectively. The remaining transition probabilities are then pinned down by the requirement that the unemployment rate depends only on current productivity. Finally, we set the UI replacement rate μ equal to 0.40.

Internally Calibrated Parameters

For the set of moments that households use to forecast future wages and returns, we follow Krusell and Smith (1998) and use only the first moment of $\mathbf{m} = \int g(\sigma_i)\Gamma(d\sigma_i) = \bar{k}$. Even with this restricted set of moments, the model, in principle, suffers from the "infinite regress of expectations", described in e.g., Townsend (1983), which is induced by the public observation of the endogenous market outcome \bar{k}_t . To solve this problem, we exploit a feature of economies with full information: the sequence of shocks $\{z_s\}_{s=0}^{t-1}$ alone allows for very accurate predictions about the future capital stock \bar{k}_{t+j} (Den Haan, 2010). We therefore set $\mathcal{I}_t^{max} = z_t$, so that households simply decide each period whether or not to acquire information about the exogenous, current value of productivity z_t . Importantly, we check ex post that this maintained assumption allows households to form accurate posteriors about \bar{k}_t , and thereby make accurate forecasts of future wages and returns. We calibrate the discount factor $\beta = 0.99$ to generate a quarterly capital output ratio of 10 (Carroll et al., 2017).

We calibrate the degree of relative risk aversion and the information cost parameters κ_t and ν to qualitatively capture key features of the micro-data presented in Section 2 (cf. Section 5.3). We set $\gamma = 5$ and the monetary cost equal to $\nu = 0.0012$ (equivalent roughly to 0.05 percent of pre-tax wages) to match our empirical finding from Section 2 that forecast accuracy increases in wealth.¹² We set the scale parameter α^κ equal to $(1/3e^{-8})$ to capture the dispersion in expectations even for households with similar observable characteristics.

To see how household expectations compare to those observed in the SCE, we concentrate on expectations of future unemployment. As mentioned in Section 2, the SCE elicits expectations of future unemployment in the form of the "percent chance that 12 months from now the unemployment rate in the U.S. will be higher than it is now". We choose the same associated measure of "probabilistic errors" ν_{it}^u in our model economy: For all households, we compute the absolute difference between a household's perceived probability conditional on its current

¹²The benefit of additional information for wealthy households arises mainly from improved predictions about the future rate of return on capital. But when relative risk aversion is close to one, income and substitution effects largely cancel one another, and wealthy households do not value those improved predictions.

Table I: Expectational errors

	Mean Error	Standard deviation
Model	122	60
SCE Data	123	87

The table shows the mean and standard deviation of expectational errors ε_u in unemployment expectations. Individual errors are calculated as absolute difference between the individual perceived probability that the unemployment is higher in period $t + 4$ than in period t and the true probability, as a percentage fraction of the latter. For SCE data, we compare the survey responses to the consensus probability implied by individual density forecasts for the unemployment rate according to the Survey of Professional Forecasters.

information $P(u_{t+4} > u_t | p_{it})$ and the true probability $P(u_{t+4} > u_t | z_t)$, which is conditional on full information about current productivity. We then express this quantity as a percentage of the true probability, so that

$$v_{it}^u = \frac{P(u_{t+4} > u_t | p_{it}) - P(u_{t+4} > u_t | z_t)}{P(u_{t+4} > u_t | z_t)}$$

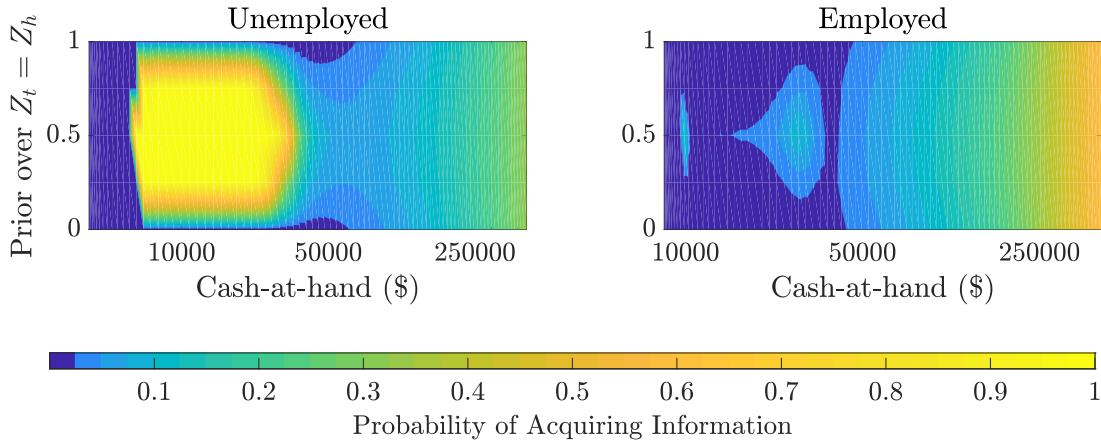
To construct an equivalent measure in the data requires a measure of the “true” probability of rising unemployment. Consistent with our analysis in Section 2, we proxy the true probability with the consensus estimate from the Survey of Professional Forecasters. This reflects ample evidence that professional forecasters provide more accurate predictions than even those from modern statistical and economic models (e.g., Bhandari et al., 2019; Croushore and Stark, 2019). In particular, our expectational error in the data equals the absolute difference between the probability of rising unemployment reported by SCE respondents and the consensus probability of a rising unemployment rate from the SPF.¹³ Table I compares the mean and standard deviation of the resulting relative errors in the model and the data. The dispersion in expectations is somewhat smaller than in the data, but overall the model replicates the first two moments of expectational errors in the data well.

5 Results

To understand the mechanisms driving our main quantitative findings, we proceed in three steps. First, we characterize the household information acquisition decision. Next, we show

¹³We compute the SPF probability as the average derived from individual predictive densities. To adjust for business cycle effects in our short data sample, we scale both errors by the average probability of rising unemployment over the sample / simulation period.

Figure 2: Information acquisition policies



For the unemployed (left panel) and the employed (right panel), and at mean prior about aggregate capital \bar{k} , the figure shows the probability of information acquisition for different values of the prior ϕ_1 and different values of individual cash on hand in the benchmark parameterization of the model.

how differential information choice affects savings decisions. Finally, we show how these two decisions interact to match the micro data and explain their impact on aggregate dynamics.

5.1 The Household Information Acquisition Decision

We start by analyzing when a household chooses to acquire information. This will provide important insight into the interaction between household heterogeneity and aggregate information dynamics. Decisions are most easily described by the probabilities of information acquisition before the realization of the extreme-value shocks. We plot the probabilities as a function of the household state variables, cash-at-hand and the prior over productivity, for the employed and unemployed, respectively, in Figure 2. Unsurprisingly, households with less informative prior expectations (closer to one-half) are more likely to acquire information.

Employed households—the precautionary savers in the economy—are less likely to acquire information, especially at low levels of wealth. These households are not at risk of hitting the borrowing constraint and have relatively low savings, so the costs from acquiring information outweigh the benefits. They also know that in the event of future job loss they have the option to acquire information—that option value further reduces the incentive to contemporaneously acquire information. As cash-at-hand (and hence wealth) rises, however, the costs of acquiring information relative to wealth fall and the benefits of predicting returns on increasing financial

wealth rise, increasing information acquisition probabilities.

Now, consider instead unemployed households. The unemployed are dissavers in the model, and when their cash-at-hand falls low enough they end up at the borrowing constraint. Hence, at very low values of cash-at-hand households almost never acquire information. Those households, pushed up against the borrowing constraint are not saving anyway, so have no benefit from acquiring information. However, as their cash-at-hand begins to rise they rapidly start to acquire information. The cost of making a savings mistake close to the borrowing constraint is very high (because marginal utility is high), so those households almost uniformly choose to acquire information. As wealth rises further, marginal utility falls, the savings policy function then becomes approximately linear, so the value of information initially falls (since the household is no longer at risk of hitting the constraint due to a mistake). The value of information then, however, starts to rise again with wealth for the same reasons as in the case of the employed.

After analyzing the household information acquisition decision, we now proceed to analyze how information affects the savings decisions of the households. This will provide an important second step in understanding how household information choice ultimately affects economic inequality and aggregate dynamics.

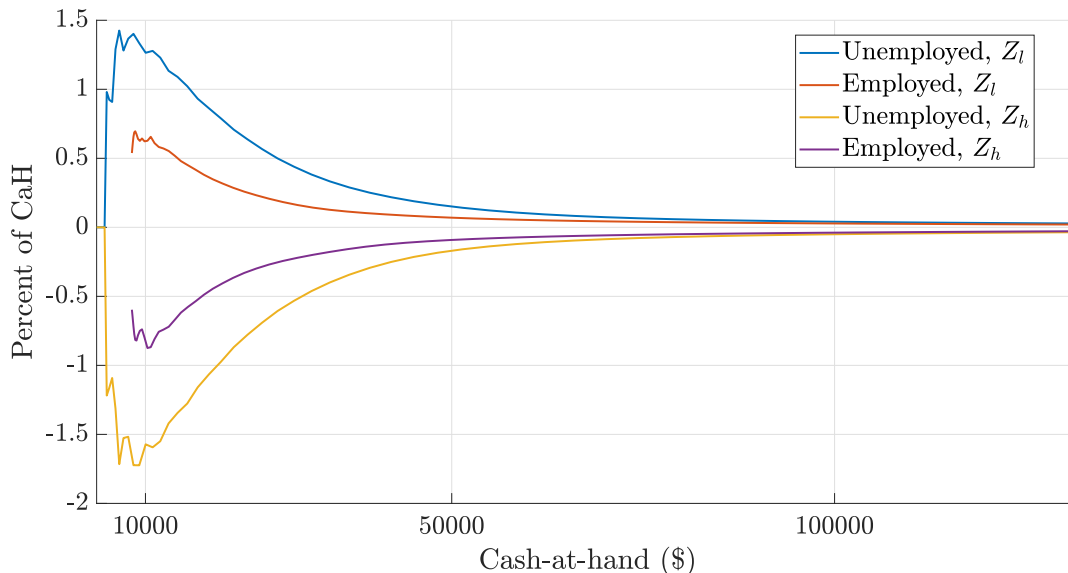
5.2 Saving Choices and Information

Consider the polar cases of a household who has just acquired information (“informed”) and a household with a 50/50 prior (“uninformed”). In Figure 3 we plot the difference in savings choices (informed minus uninformed) as a function of their cash-at-hand and their employment status. All else equal, informed households save more than uninformed in busts (Z_l), as they know that the risk of becoming unemployed or staying unemployed is higher, but less in booms (Z_h). The percentage difference in savings is, however, strongly heterogeneous across the distribution of cash-at-hand, and between unemployed and employed households.

Savings rates of informed households differ strongly from those of the uninformed at low but positive levels of wealth, where the precautionary-savings motive is stronger. This effect is larger for the unemployed, who benefit from information about current productivity to predict future job-finding rates. As cash-at-hand increases, the difference in savings rates between informed and uninformed households decreases because the precautionary motive is less

sensitive to the aggregate state, so being better able to predict the difference in idiosyncratic risk between booms and recessions is less valuable. It is important to note here that in this thought experiment informed and uninformed households have the same prior over the aggregate capital stock, so that the perceived difference in wages and returns from being informed is only due to the perceived difference in productivity and overall labor supply.

Figure 3: Difference in savings policies: informed - uninformed



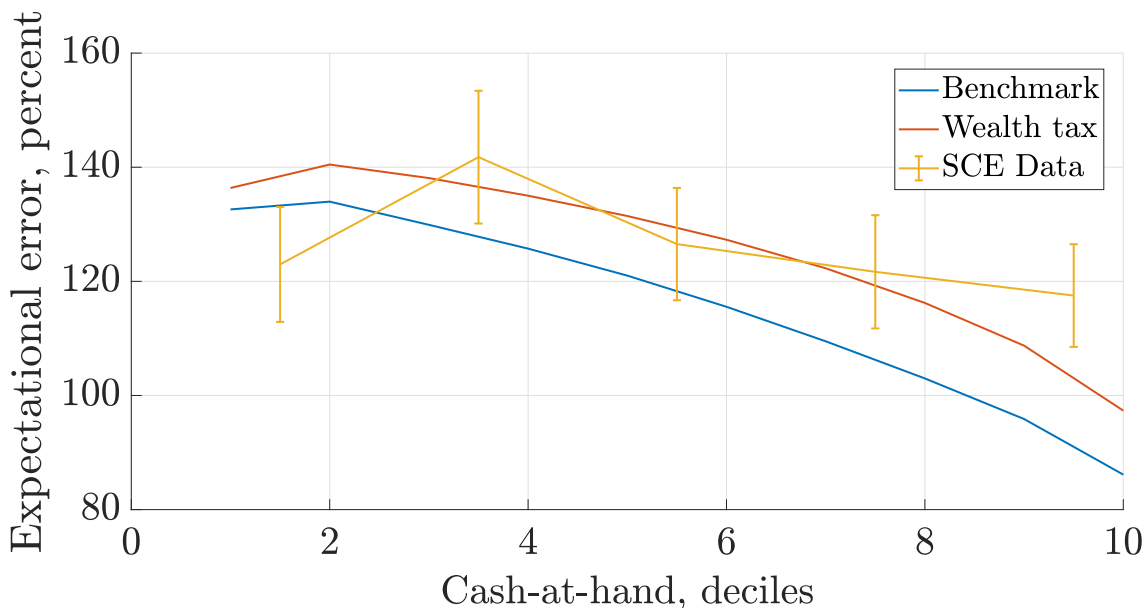
The figure shows the difference between savings by informed agents and uninformed agents ($\phi_1 = 0.5$) at a mean prior for aggregate capital \hat{k} and different values of individual cash on hand (along the bottom axis), in the benchmark parameterization of the model. The difference is defined as uninformed minus informed.

5.3 Accuracy of Expectations

We described how wealth and employment status affect a household’s decision to acquire information, and how a household’s savings are affected by the accuracy of its information. Before we turn to the economic consequences of the two-sided interaction that exists between household heterogeneity and information, we study how these forces shape the accuracy of household expectations across the wealth distribution. Specifically, Figure 4 shows how households’ information acquisition probabilities in equilibrium translate into a systematic relationship between the accuracy of households’ expectations and their wealth. The model generates an inverse-U shape, which closely matches that documented in the micro data in Section 2. The inverse-U shape is a result of two countervailing forces: First, the upward sloping part of the

curve is driven by the unemployed. The poorer households in the model are, on average, the unemployed, who at low levels of wealth acquire information with high probability. As those households find jobs, their wealth increases, but they also stop acquiring information, leading to the increase in average absolute forecast errors between the first and second decile of the wealth distribution. Second, beyond the second decile, the probability of acquiring information is monotonic in wealth. These are in effect the employed households that we discussed in Section 5.1. As wealth increases, the costs of making savings mistakes rise, and households' information acquisition probabilities increase. This then leads to the decline in average absolute forecast errors above the second decile visible in Figure 4.

Figure 4: Errors in unemployment expectations



The figure shows errors in unemployment expectations ε_u averaged within deciles of the distribution of individual capital holdings for the benchmark parameterisation (blue line), with a 1 percent p.a. wealth tax (red line) and the corresponding errors in the SCE (yellow line, with 95% confidence bars). Individual errors are calculated as the individual perceived probability that the unemployment rate is higher in period $t + 4$ than in period t minus the true conditional probability, as a percentage fraction of the latter probability, averaged across the sample / simulation.

5.4 Aggregate Implications

Our discussion thus far has focused on the dynamics of savings and information choice for an *individual* household. To understand how those decisions will impact wealth inequality and aggregate dynamics, it is important to understand how individual decisions interact to shape

the cross-sectional distribution of wealth and information. In order to build intuition for how those forces balance in general equilibrium, first consider the overall frequency of information acquisition. In any given period about 14% of households acquire information, so that on average households update their information about once every 7 quarters. As suggested by Figure 2, the probabilities with which households purchase information differ strongly across the equilibrium distribution of households. For example, on average, every period 10 percent of households update with a probability less than 5 percent, while another 10 percent update with a probability that exceeds one quarter.

Dynamically, households who choose not to acquire information will have priors (\hat{k}) about the level of the capital stock that are more tilted towards the long-run average level of aggregate capital. Hence, in booms, they will systematically underpredict the capital stock (and overpredict the return r), and vice-versa in recessions. We should note, however, that a hypothetical household that acquired information in every period would have an accurate estimate of the true capital stock and make negligible forecast errors as shown in Appendix A.3. The systematic interplay between information acquisition, priors and savings choices will impact both the aggregate time-series of capital and the distribution of capital holdings in the cross-section.

5.4.1 Aggregate Dynamics

In Table II, we contrast the aggregate dynamics of our benchmark economy with the full-information counterpart. Fluctuations in all aggregate variables are substantially more pronounced in the economy with heterogenous information. The standard deviation of the capital stock is about 40 percent higher. The stark difference can be understood in light of the foregoing discussion. Since uninformed households systematically underpredict the capital stock in booms, they *overpredict* the return on savings, and thus save more than if they had full information. The converse is true in recessions. Thus, in general equilibrium, the economy systematically overaccumulates capital in booms and underaccumulates in recessions, leading to the much larger variance of capital, output and consumption.

The presence of imperfect information thus serves as an amplifying force—it induces weaker mean-reversion of the capital stock relative to the full-information economy. These dynamics elucidate a more general feature of our framework: information acquisition decisions are *strate-*

gic substitutes. The individual benefits of information rise with the dispersion of the capital stock. Thus when the average share of information in the economy increases, the dispersion of the capital stock falls, and so does the incentive to acquire information. In a companion paper, (Broer et al., 2021), we show how this may imply non-existence of homogeneous-information (representative-agent) equilibria in neoclassical economies.

Table II: Moments of aggregate quantities, average

	St dev K	St dev Y	St dev I	St dev C	Info unemp	Info emp
Benchmark	5.13	3.45	7.77	3.18	0.16	0.14
Full information	3.61	3.16	7.11	2.92	1.00	1.00
Difference, percent	41.92	9.23	9.33	9.02	-84.13	-86.31

The table shows the standard deviations of the time series for natural logarithms of capital (column one), output (column two), investment (column three) and consumption (column four), as well as the average percentage fraction of unemployed (column five) and employed agents (column six) that acquire information in a given period, averaged across a long simulation of the benchmark economy (top row), and its full-information counterpart (middle row), all in percent, plus their percentage difference (bottom row).

5.4.2 Wealth Inequality

Compared to the full-information economy, inequality is substantially higher in our heterogenous-information model—the Gini coefficient of wealth increases from 0.35 to 0.40.¹⁴ In particular, the right tail of the wealth distribution holds substantially more wealth with heterogenous information: the wealth share of the top 1% and 10% is about 5% higher in our economy compared to full-information. The reason for this increased wealth inequality is that limited information dampens the correlation between returns and savings rates. Indeed, the proceeding discussion about aggregate savings dynamics plays out similarly at the individual level. The combination of the extreme-value shocks and the systematic tilt towards the mean in uninformed households’ capital beliefs implies that when households are far away from the borrowing constraint and have roughly linear policy functions, the slope of the savings policy function behaves as if it was stochastic. The slope of the policy function is pinned down by the discount factor and the expected rate of return. Since households only infrequently update

¹⁴It is well known (see, e.g., Krueger et al., 2016) that unemployment risk is not sufficient to generate the concentration in the wealth distribution observed in the data. We are confident that our qualitative results would be unchanged if we were to add the necessary ingredients to generate higher wealth inequality. In fact, the inverse-U shape in information choice is likely to be more pronounced, as currently the wealth poor in the model hold too much wealth relative to the data.

information, their information about the expected rate of return is effectively stochastic. The behavior is therefore analogous to the mechanism described in Piketty and Saez (2003) of how exogenous random savings rates can generate Pareto tails in the wealth distribution. Here, however, heterogenous information choice provides a microfoundation for that type of seemingly random savings behavior, as opposed to other models that either assume exogenously stochastic savings rates, discount factors or returns on savings.

6 A Wealth Tax

The foregoing discussion highlighted the interplay between precautionary savings, information choice and the aggregate economy. In this section we illustrate how policy reforms that affect precautionary-savings motives will in turn alter information acquisition decisions, potentially changing the information content of the economy and shaping aggregate dynamics in unanticipated ways.

For our counterfactual, we consider a wealth tax. Such a tax has been hotly debated by policy makers in recent years and recently introduced as a policy proposal in the U.S. Congress. One of the main arguments of the proponents of the tax is that it will reduce inequality and be an efficient way to finance increased government spending. As such, we consider the counterfactual policy experiment in which the government imposes a linear wealth tax $\tau_k > 0$ on beginning-of-period capital holdings to finance government spending.¹⁵ Household cash-at-hand y is therefore given by the expression:

$$y = rk(1 - \tau)\bar{l}\epsilon w + \mu(1 - \epsilon_t)w + (1 - \delta - \tau_w)k \quad (9)$$

Because the probability of information acquisition is strongly heterogeneous in individual wealth (Figure 2), the wealth tax changes the average level of information in the economy. Although information acquisition policies for *a given wealth level* are approximately unaffected by the tax, a wealth tax reduces the mass of households that acquire information every period by reducing average savings, and hence wealth. In effect, the introduction of a wealth tax moves the mean of the distribution of individual capital to the left, resulting in a shift in expectation

¹⁵For continuity with the previous sections, we assume that the spending is unvalued by households. For the positive statements of this section, this is isomorphic to assuming an additively separable utility function over G .

errors as shown by the “wealth tax” line in Figure 4. Table III shows that the effects of this movement on household information choices is powerful. It amounts to a reduction in the number of households who acquire information every period by 15 and 30 percent for a one-half and one percent p.a. wealth tax, respectively. The fall for the employed is somewhat more pronounced, as their information acquisition probability does not show the pronounced hump shape at low asset levels found for the unemployed. They are therefore more sensitive to the reduction in wealth that is implied by the wealth tax.

Table III: Impact of a wealth tax

	Mean \bar{k}_t	St dev \bar{k}_t	St dev Y_t	90/10	99/1	Info acqu.
Limited information, 0.5 %	-5.2	3.6	1.3	-0.5	-1.4	-15.5
Full information, 0.5 %	-5.3	0.7	0.3	-2.3	-2.2	0.0
Limited information, 1 %	-10.0	8.0	2.8	2.0	2.0	-30.4
Full information, 1 %	-10.2	1.3	0.5	-4.3	-7.3	0.0

The table shows percentage changes of equilibrium moments implied by the introduction of respectively, a 0.5 and 1 percent p.a. wealth tax ($\tau_k = 0.00125, 0.0025$ for the quarterly tax) in the limited-information (rows 1 and 3) and full-information (rows 2 and 4) economies. The moments are the mean capital stock (column one), the standard deviations over time of natural logarithms of capital (column two) and output (column three), and 90/10 and 99/1 percentile ratios of the cross-sectional standard deviation of wealth holdings (columns four and five), and the average share of individuals acquiring information every period (column six, for the limited-information economies only).

Table III further shows quantitatively that the change in household information materially alters the consequences of the wealth tax on macroeconomic volatility and inequality relative to the full-information counterpart. Specifically, capital and output volatility are roughly unchanged under full-information. By contrast, by further dampening the mean-reversion of capital, the reduction in information acquisition substantially *increases* the volatility of both capital and output in our environment. The difference in the effect of the tax on inequality between the two environments is even more pronounced. There are two opposing forces that explain our results. First, the direct impact of the wealth tax disproportionately affects high wealth households, lowering their share of wealth. Second, poorer households on average are less informed, and hence make worse savings choices. With a 0.5 percent tax, the two effects approximately offset each other. With a one percent tax, the information effect dominates, such that the wealth tax actually *increases* wealth inequality: both percentile ratios in Table III increase by two percent.

While we abstain from making welfare statements about the desirability of such a policy, our positive findings indicate that policymakers should proceed with caution when evaluating the consequences of tax policy. More generally, the experiment illustrates how macroeconomic policies may have important additional effects in environments with heterogeneous, endogenous information. By changing the distribution of agents' information and expectations, macroeconomic policies fundamentally alter an economy's responsiveness to shocks, as well as individual agents' decision rules. Moreover, these additional effects may be quantitatively important — both from a positive and a normative point of view.

7 Conclusion

The frontier of macroeconomic research continues to recognize and incorporate salient dimensions of heterogeneity to provide a more complete and accurate understanding of the economy. In this paper, we illustrated how the interaction between heterogeneity in expectations and precautionary savings gave rise to new qualitative and quantitative insights about inequality and macroeconomic dynamics. Our findings have important implications for both the heterogeneous-agent macro literature and the heterogeneous-expectations literature. For the former, our policy experiment provides a Lucas-type critique (Lucas Jr, 1976) to policy analysis in incomplete markets models: any policy that has a first order impact on the wealth distribution will in turn affect the information content of the economy, with implications for dynamics and the cross-section. For the latter, studying information choice with linear policy rules misses the two-way feedback between heterogeneity in expectations and the distribution of agents. Our framework provides a laboratory to push both strands of the literature forward to explore new questions in macroeconomics.

Our analysis is positive in nature, but raises interesting normative questions. Particularly, information acquisition has obvious externalities in our environment through the implied change in the dynamic properties of prices and aggregate quantities. Does this mean policymakers should subsidize information? Should such subsidies target a particular part of the population? We leave these exciting and important questions for future research.

References

- ANGELETOS, G.-M., Z. HUO, AND K. A. SASTRY (2020): “Imperfect Macroeconomic Expectations: Evidence and Theory,” Tech. rep., National Bureau of Economic Research.
- ARMANTIER, O., G. TOPA, H. W. VAN DER KLAUW, AND B. ZAFAR (2017): “An Overview of the Survey of Consumer Expectations,” *Economic Policy Review*, 23, 51–72.
- AUCLERT, A., M. ROGNLIE, AND L. STRAUB (2020): “Micro jumps, macro humps: Monetary policy and business cycles in an estimated HANK model,” Tech. rep., National Bureau of Economic Research.
- BARLEVY, G. AND P. VERONESI (2000): “Information acquisition in financial markets,” *The Review of Economic Studies*, 67, 79–90.
- BHANDARI, A., BOROVIČKA, AND P. HO (2019): “Survey Data and Subjective Beliefs in Business Cycle Models,” Tech. rep.
- BORDALO, P., K. COFFMAN, N. GENNAIOLI, AND A. SHLEIFER (2016): “Stereotypes,” *The Quarterly Journal of Economics*, 131, 1753–1794.
- BORDALO, P., N. GENNAIOLI, AND A. SHLEIFER (2017): “Memory, attention, and choice,” *The Quarterly journal of economics*.
- BRANCH, W. A. AND B. MCGOUGH (2018): “Heterogeneous expectations and micro-foundations in macroeconomics,” in *Handbook of computational economics*, Elsevier, vol. 4, 3–62.
- BROCK, W. A. AND C. H. HOMMES (1997): “A Rational Route to Randomness,” *Econometrica*, 65, 1059–1095.
- (1998): “Heterogeneous beliefs and routes to chaos in a simple asset pricing model,” *Journal of Economic dynamics and Control*, 22, 1235–1274.
- BROER, T., A. KOHLHAS, K. MITMAN, AND K. SCHALFMANN (2021): “The Impossibility of Krusell-Smith Equilibria,” Tech. rep.
- CARROLL, C., J. SLACALEK, K. TOKUOKA, AND M. N. WHITE (2017): “The distribution of wealth and the marginal propensity to consume,” *Quantitative Economics*, 8, 977–1020.

- CARROLL, C. D. (1997): “Buffer-stock saving and the life cycle/permanent income hypothesis,” *The Quarterly journal of economics*, 112, 1–55.
- (2003): “Macroeconomic expectations of households and professional forecasters,” *the Quarterly Journal of economics*, 118, 269–298.
- CARROLL, C. D., E. CRAWLEY, J. SLACALEK, K. TOKUOKA, AND M. N. WHITE (2020): “Sticky expectations and consumption dynamics,” *American economic journal: macroeconomics*, 12, 40–76.
- CROUSHORE, D. AND T. STARK (2019): “Fifty Years of the Survey of Professional Forecasters,” *Economic Insights*.
- DEN HAAN, W. J. (2010): “Assessing the accuracy of the aggregate law of motion in models with heterogeneous agents,” *Journal of Economic Dynamics and Control*, 34, 79–99.
- D’ACUNTO, F., D. HOANG, M. PALOVITA, AND M. WEBER (2019): “IQ, expectations, and choice,” Tech. rep., National Bureau of Economic Research.
- GABAIX, X. (2019): “Behavioral inattention,” in *Handbook of Behavioral Economics: Applications and Foundations 1*, Elsevier, vol. 2, 261–343.
- GROSSMAN, S. J. AND J. E. STIGLITZ (1980): “On the impossibility of informationally efficient markets,” *The American economic review*, 70, 393–408.
- HELLWIG, C. AND L. VELDKAMP (2009): “Knowing what others know: Coordination motives in information acquisition,” *The Review of Economic Studies*, 76, 223–251.
- JAKOBSEN, K., K. JAKOBSEN, H. KLEVEN, AND G. ZUCMAN (2020): “Wealth Taxation and Wealth Accumulation: Theory and Evidence from Denmark,” *The Quarterly Journal of Economics*, 135, 329–388.
- KAPLAN, G. AND G. L. VIOLANTE (2018): “Microeconomic heterogeneity and macroeconomic shocks,” *Journal of Economic Perspectives*, 32, 167–94.
- KRUEGER, D., K. MITMAN, AND F. PERRI (2016): “Macroeconomics and household heterogeneity,” in *Handbook of Macroeconomics*, Elsevier, vol. 2, 843–921.
- KRUSELL, P. AND A. SMITH (1998): “Income and Wealth Heterogeneity in the Macroecon-

- omy.” *Journal of Political Economy*, 106.
- LUCAS JR, R. E. (1976): “Econometric policy evaluation: A critique,” in *Carnegie-Rochester conference series on public policy*, North-Holland, vol. 1, 19–46.
- MAĆKOWIAK, B., F. MATĚJKA, AND M. WIEDERHOLT (2018a): “Dynamic Rational Inattention: Analytical Results,” *Journal of Economic Theory*, 176, 650–692.
- (2018b): “Survey: Rational inattention, a disciplined behavioral model,” .
- MACKOWIAK, B. AND M. WIEDERHOLT (2009): “Optimal sticky prices under rational inattention,” *American Economic Review*, 99, 769–803.
- MANKIW, N. G. AND R. REIS (2002): “Sticky information versus sticky prices: a proposal to replace the New Keynesian Phillips curve,” *The Quarterly Journal of Economics*, 117, 1295–1328.
- PIKETTY, T. AND E. SAEZ (2003): “Income inequality in the United States, 1913–1998,” *The Quarterly journal of economics*, 118, 1–41.
- SIMS, C. A. (2003): “Implications of rational inattention,” *Journal of monetary Economics*, 50, 665–690.
- VELDKAMP, L. L. (2011): *Information choice in macroeconomics and finance*, Princeton University Press.
- VISSING-JORGENSEN, A. (2003): “Perspectives on behavioral finance: Does” irrationality” disappear with wealth? Evidence from expectations and actions,” *NBER macroeconomics annual*, 18, 139–194.

A Appendix

A.1 SCE data, sample selection, and motivating evidence

The SCE is a monthly internet survey of about 1300 “household heads”, defined as the person in a household who owns, is buying or rents the home. Subjects are chosen from the respondents to the Consumer Confidence Survey (CCS), itself based on the universe of US postal addresses, to match demographic targets from the American Community Survey, and remain in the survey for up to 12 months. The SCE core module contains monthly information about households’ expectations about key macroeconomic and individual variables, including consumer price and house price inflation, unemployment, interest rates, etc. Importantly, a yearly module also asks the survey respondents for key financial variables, including their net wealth.

Variable definitions

We use three expectational variables from the SCE: CPI inflation, house price inflation, and unemployment expectations. The former two ask respondents for their best guess for a variable’s outcome, in addition to the probability of it falling into a number of bins. The exact questions are:

- CPI inflation:

“What do you expect the rate of inflation to be over the next 12 months? Please give your best guess.”, followed by *“In your view, what would you say is the percent chance that, over the next 12 months the rate of inflation will be...”*

- house price inflation:

“By about what percent do you expect the average home price to [increase/decrease]? Please give your best guess.”, followed by *“And in your view, what would you say is the percent chance that, over the next 12 months, the average home price nationwide will...”*

We calculate forecast errors as the absolute difference between individual best guesses and the actual (12-month-ahead) outcomes of US consumer price index inflation and inflation of the S&P/Case-Shiller 20-City Composite Home Price Index, respectively. We use the measures of interquartile ranges of individual forecasts provided by the SCE.

For unemployment expectations, the survey does not ask for point forecasts but elicits beliefs about the probability that the national unemployment will rise:

- unemployment:

“What do you think is the percent chance that 12 months from now the unemployment rate in the U.S. will be higher than it is now?”

To construct errors ν_{it} of individual unemployment forecasts $P(u_{t+12} > u_t|t)$, we would ideally compare household i 's response to the true probability $P(u_{t+12} > u_t|t)$. However, the latter is unobserved. Consistent with ample evidence that professional forecasters provide more accurate predictions than those from modern statistical and economic models (Stark, 2010; Bandhari et al, 2020), we proxy the true probability by the consensus forecast from the SPF, which we denote $P_{SPF}(u_{t+12} > u_t|t)$. In particular, we calculate each forecaster's belief about the probability of rising unemployment (using the probabilistic answers in the variable PRUNEMP), and then average over forecasters. Finally, since the data was collected during a time of steadily falling unemployment, we scale the difference between a household's expectations and the consensus forecast of professional forecasters by the average consensus forecast to make the measure comparable to the model-implied probabilities that are calibrated to a different time period. We thus compute the error in unemployment forecasts as

$$\nu_{it} = \frac{P_i(u_{t+12} > u_t|t) - P_{SPF}(u_{t+12} > u_t|t)}{P_{SPF}(u_{t+12} > u_t)} \quad (10)$$

In addition to the expectational variables we use the following household characteristics: gender, age, dummies that take values 1 if the household head reports to have a college degree, or to participate in the labor market (in the sense that she / he is either employed or unemployed), respectively. We also use a measure of household net financial wealth, which we construct as the difference between a household's total financial assets and non-mortgage debt.¹⁶ We deflate the resulting quantities by the level of the US consumer price index.

¹⁶The question about financial assets is “Approximately what is the total current value of your [and your spouse's/partner's] savings and investments (such as checking and savings accounts, CDs, stocks, bonds, mutual funds, Treasury bonds), excluding those in retirement accounts?”, that for mortgage debt “Approximately, what is the total amount of outstanding loans against your home(s), including all mortgages and home equity loans?”, that for total debt “Approximately, what is the total amount of your [and your spouses/partners] current outstanding debt?”.

Table IV: Macroeconomic expectations in the SCE and the SPF

Prob rising unemployment	Median	SD		
SCE	40	22.7		
SPF	30	17.8		
CPI inflation	Median error	Median IRQ	SD error	SD IQR
SCE	1.85	2.06	2.47	4.15
SPF	0.72	.56	0.65	0.25

The table shows moments of individual perceived probability distributions according to data from the Survey of Consumer Expectations (SCE) and of individual forecast distributions from the Survey of Professional Forecasters (SPF). Rows two and three show the median and standard deviation of individual SCE unemployment forecasts (elicited as the “percent chance that 12 months from now the unemployment rate in the U.S. will be higher than it is now”) and the equivalent moments inferred from individual SPF forecast distributions, which enter equation (10). Rows five and six show, again for SCE and SPF, respectively, the median error of individual CPI inflation expectations / forecasts (column 2), the standard deviation of these errors (column 2), the median interquartile ranges calculated from individual distributions (column 3), and their standard deviations (column 4).

Sample restriction

We do not perform any sample selection other than dropping households whose median inflation expectations lie in the extreme bins (higher than 12 or lower than -12 percent) respectively.

Details about motivating evidence

Table IV illustrates that households’ 12-month unemployment and inflation expectations from the SCE are on average less accurate than professional forecasts (stylized fact 1). Households attach on average a higher probability to rising unemployment than professional forecasters, implying larger forecast errors during a sample period where unemployment actually declined steadily. We find a similar picture for CPI inflation: The median of household point forecast errors are substantially larger for households than for professional forecasters—equal to 1.8 and 0.7 percentage points (pp), respectively.

Furthermore, table IV demonstrates that household expectations are substantially more uncertain than professional forecasts (stylized fact 2). When elicited for their probability distribution over possible inflation realizations, households report substantially wider distributions. The median of the interquartile ranges of *individual* forecast distributions is more than triple that of professional forecasters—2.0pp vs. 0.6pp.

Table IV also shows that household expectations are substantially more heterogeneous than

Table V: Estimates of equation (11)

	(1)	(2)	(3)	(4)	(5)
	UE, Abs error	Inflation, Abs Error	Inflation, IQR	HP Inflation, Abs Error	HP Inflation, IQR
Male	0.0466 (1.01)	-0.354*** (-2.79)	-0.776*** (-3.68)	-0.292*** (-2.78)	-0.182 (-0.92)
College Degree	-0.132** (-2.53)	-0.506*** (-3.36)	-0.741*** (-3.04)	-0.386*** (-3.08)	-0.252 (-1.10)
Participation	0.128*** (2.67)	0.243 (1.56)	-0.435 (-1.57)	0.153 (1.31)	-0.390 (-1.45)
Fin Wealth, 1st quintile	0.0872 (1.15)	1.102*** (5.55)	0.915** (2.47)	0.252 (1.37)	0.402 (1.15)
Fin Wealth, 2nd quintile	0.226*** (2.94)	0.914*** (4.60)	1.471*** (4.90)	0.639*** (4.13)	0.783** (2.53)
Fin Wealth, 3rd quintile	0.0804 (1.18)	0.400** (2.17)	0.412 (1.60)	0.218 (1.47)	-0.0630 (-0.26)
Fin Wealth, 4th quintile	0.0314 (0.47)	0.0823 (0.50)	-0.0459 (-0.23)	0.0163 (0.11)	-0.160 (-0.80)
Constant	0.973 (1.60)	1.766 (1.05)	8.081** (2.10)	1.887 (1.21)	10.74*** (3.13)
r2	0.0419	0.0954	0.0765	0.0280	0.0358
N	9139	8618	8618	7537	7537

t statistics in parentheses

* $p < .1$, ** $p < .05$, *** $p < .01$

The table shows estimates of equation (11) for individual forecast errors and for interquartile ranges (IQR) of the individual forecast distributions in SCE data. Column 1 shows estimates for the errors in individual unemployment forecasts (elicited as the “percent chance that 12 months from now the unemployment rate in the U.S. will be higher than it is now”) relative to the equivalent consensus forecast from professional forecasters (equation (10)). Columns 2 and 3 present results for 12-month-ahead CPI inflation, columns 4 and 5 for house price inflation, where forecast errors equal the absolute difference between forecasts and outcomes. All regressions also include a cubic function in age, and year-month-dummies. T-statistics are shown in parentheses. Standard errors are clustered at the individual level; stars denote conventional significance levels: * ($p < .1$), ** ($p < .05$), *** ($p < .01$).

SPF forecasts (stylised fact 3). Specifically, household unemployment expectations and point forecasts for CPI inflation have a substantially higher standard deviation than the forecasts of professionals. For example, the standard deviation of forecast errors for CPI inflation across households is about three times larger than across professionals.

Stylized fact 4 states that the accuracy and uncertainty of household expectations are inverse-U shaped in wealth. To identify the relationship between individual characteristics and expectations of macroeconomic outcomes, we estimate the following regression

$$y_{it} = \beta \mathbb{X}_{it} + \varepsilon_{it} \quad (11)$$

where y_{it} denotes a moment of individual i 's perceived distribution of an outcome in period $t+12$ as elicited in period t , β is a vector of coefficients to be estimated. \mathbb{X}_{it} is a set of regressors and ε_{it} is an error term. As regressors we include dummies that equal 1 when the household's

Table VI: Parameterization

Parameter	Value
<i>Externally calibrated parameters</i>	
Capital share (α)	0.36
Depreciation rate (δ)	0.025
Persistence of booms	0.88
Persistence of busts	0.82
Ratio of productivity between booms and bust (z_h/z_l)	1.027
Unemployment rate in booms	0.06
Unemployment rate in busts	0.06
Monthly job-finding rate in booms	0.55
Monthly job-finding rate in busts	0.45
Unemployment insurance replacement rate(μ)	0.40
<i>Internally calibrated parameters</i>	
Discount factor (β)	0.99
Relative risk aversion (γ)	5
Monetary cost of information (ν)	0.0012
Scale parameter of utility cost of information (α^κ)	$1/3e^{-8}$

head is a man, has a college degree, or participates in the labor market, plus a cubic function in age, year-month-dummies, and indicator variables for quintiles of the distribution of real net financial wealth

Table V reports the estimates for equation (11) when y_{it} equals the forecast error in unemployment as constructed above (column 1), absolute errors and individual interquartile ranges for CPI inflation expectations (columns 2 and 3, respectively) and for house price inflation expectations (columns 4 and 5). We cluster standard errors at the individual level. The coefficients associated with the wealth-quintile dummies are depicted in Figure 1 in the main text.

A.2 Calibration Parameters

Table VI summarizes the parameters of the calibrated model.

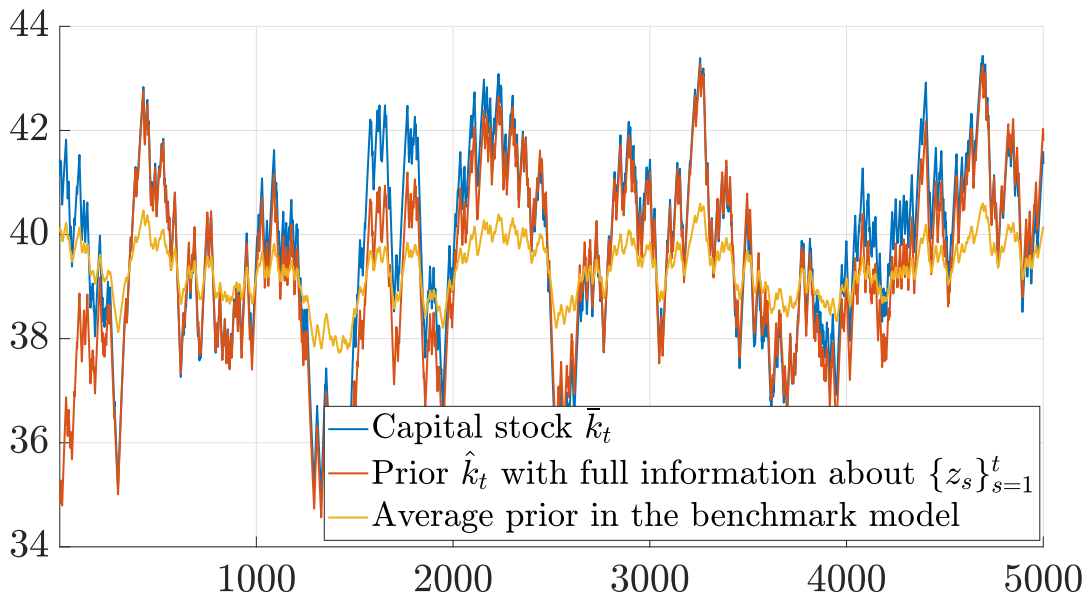
A.3 Time Series for Capital and Priors

Limited information makes individual prior expectations about the current capital stock move much more slowly than the actual capital stock. In particular, households who choose not to

acquire information will have priors (\hat{k}) about the level of the capital stock that are more tilted towards the long-run average level of aggregate capital. Hence, in booms, they will systematically underpredict the capital stock (and overpredict the return r), and vice-versa in recessions.

Importantly, however, this sluggishness is not a consequence of our maintained assumption that households estimate the current capital stock only from the information they acquire about productivity. In fact, for economies with full information, Den Haan (2010) shows that the sequence of shocks $\{z_s\}_{s=0}^{t-1}$ alone allows for very accurate predictions about the future capital stock \bar{k}_{t+j} . We verify that this holds also in our setup. Figure 5 depicts the time series of the actual capital stock (blue line), the prior belief of an individual that acquires information in every period (red line) and for comparison the average prior belief in our benchmark economy (yellow line). An individual that always acquires information would have prior beliefs that closely track the realized value (with a correlation of 0.95).¹⁷

Figure 5: Mean capital \bar{k}_t - realization and priors



Based on a simulation of the benchmark model, the figure shows time series of the capital stock \bar{k}_t (blue line), the prior about current aggregate capital \hat{k}_t of households who acquire information about the current productivity state every period (red line), and the average prior in the benchmark economy (yellow line).

¹⁷In the figure, we start \hat{k} at an arbitrary value of 35, and discard the initial 200 periods to calculate the correlation, to demonstrate that the strong correlation does not depend on an accurate initial point prior.