Expectation and Wealth Heterogeneity in the Macroeconomy*

Tobias Broer   Alexandre Kohlhas   Kurt Mitman   Kathrin Schlafmann

December, 2022

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 macroeconomic 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 top-end wealth inequality in equilibrium.

JEL codes: D84, E21, E27, E63

Keywords: Expectations, heterogeneity, information

1 Introduction

Expectations have been at the bedrock of modern macroeconomics since the "rational expectations revolution" pioneered by Robert E. Lucas, Jr., in the 1970s. The prevailing paradigm—the full-information and rational expectations framework—assumes that all households, at all moments in time, have the same expectations about the macroeconomy. Building

*First draft: April 2020. Previous title: *Heterogeneous Information Choice in General Equilibrium.* Broer: PSE (tobias.broer@psemail.eu). Kohlhas: Oxford University (alexandre.kohlhas@economics.ox.ac.uk). Mitman: IIES, Stockholm University (kurt.mitman@iies.su.se). Schlafmann: CBS (ksc.fi@cbs.dk). For comments and useful suggestions, we thank Adrien Auclert, Zhen Huo, Per Krusell, Tommaso Monacelli, Kristoffer Nimark, Alessandra Peter, Morten Ravn, Matthew Rognlie, Kjetil Storesletten, Laura Veldkamp, Gianluca Violante, Mirko Wiederholt, and seminar and conference participants at Bern, Bocconi, CBS, CEPR Fall Meeting '21, CES-iFO Conference on Survey Data, Copenhagen, ESEM '20, FRB Minneapolis, IIES, LMU, NBER Summer Institute '22, Northwestern, Oslo, Oxford, PSE, SED, Virginia, Zheijang. This research received financial support from *Ragnar Soderberg Stiftelsen.* Broer and Mitman thank the Banque de France Chair and the European Research Council (ERC Starting Grant 759482), respectively, for financial support.

on the work of Muth (1961) and several others, Mankiw *et al.* (2003) contrast this prediction with survey data on expectations, showing instead the profound dispersion of expectations that exists among households. Recent empirical work has stressed that household expectations are not only heterogeneous but also correlate systematically with household economic characteristics. This creates systematic heterogeneity in both the level and accuracy of expectations across the distribution of households (e.g., Carroll, 2003; Coibion *et al.*, 2018; Weber *et al.*, 2022). Given the importance of expectations to macroeconomics, it is central to have a theory of expectation formation that is consistent with the data.

In this paper, we develop a theory of information choice that we embed into a standard incomplete-markets model with aggregate risk. Our framework can capture both the rich differences in expectations, observed in the data, as well as those that exist in wealth, income, and employment status. Heterogeneity in income, wealth, and employment status on its own significantly impacts the response of the economy to shocks and affects the efficacy and transmission of economic policy (see e.g., Krueger *et al.*, 2016; Kaplan and Violante, 2018). We use our quantitative-theoretical framework—disciplined by survey data—to quantify the importance of the nexus between expectations and household heterogeneity in understanding aggregate fluctuations and the distribution of wealth. We then explore how the presence of heterogeneous expectations changes the efficacy of macroeconomic policies.

To start, we provide new evidence on heterogeneity in household expectations using US micro-level data. We show that in a leading household survey (FRB NY Survey of Consumer Expectations) both the mean and self-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 household wealth and other economic characteristics: All else equal, wealthier households have more accurate expectations; however, unlike the evidence in Carroll (2003) and Vissing-Jorgensen (2003), this relationship is far from monotone—especially at the lower-end of the wealth distribution, where the accuracy of expectations is declining with wealth.[1]

Next, we embed dynamic information choice into an otherwise standard business cycle model with incomplete markets to explore if households' heterogeneous incentives to acquire information rationalize our empirical results. In the model, households form expectations about future returns, wages, and unemployment risk, to determine their optimal consumption-savings choices, and acquire costly information about the state of the economy to do so. The

[1] A burgeoning literature has begun to empirically document the various ways in which household and firms expectations differ from one another and within-groups (e.g., Coibion *et al.*, 2018; Coibion *et al.*, 2020; Reis *et al.*, 2020; Andrade *et al.*, 2022; and Macaulay and Moberly, 2022). We make two contributions to this line of research. First, we provide new evidence on the systematic (non-monotone) relationship between the accuracy of expectations and household wealth. Second, we provide the first, to our knowledge, general equilibrium model that allows for the analysis of the macroeconomic implications of systematic differences in expectations.

information that households can acquire approximates the optimal signal that households would choose to design. The gains to acquiring this information depend on household wealth, employment status, and prior beliefs, leading to systematic heterogeneity in expectations. Using this framework, we characterize the distribution of household expectations.

Our framework builds on prior work on the financial and macroeconomic consequences of costly information choices (Grossman and Stiglitz, 1980; Sims, 2003; Hellwig and Veldkamp, 2009; Veldkamp, 2011; and Maćkowiak *et al.*, 2021). This 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, heterogenous information choices*, and to quantify how they can profoundly shape macroeconomic outcomes.

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

Consider first *unemployed households*, who dissave to smooth consumption. At low levels of wealth, unemployed households have highly non-linear policy functions, leading down to the borrowing constraint (Carroll, 1997). At the borrowing constraint, households are hand-to-mouth, and hence do not value improved information about the future, so never acquire it. As wealth rises, the marginal utility of consumption for unemployed households remains high, making savings mistakes extremely costly. Additionally, because of the non-linearity of the savings policy function, uninformed savings can lead to large errors. Additional information, as a consequence, becomes highly valuable and households purchase information frequently.

When wealth increases further, marginal utility eventually falls (policy functions become approximately linear) as 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 falls. At the same time, as wealth rises, two forces induce more information acquisition that counteracts this decrease: (i) the cost of acquiring information relative to current wealth falls; and (ii) the benefit of accurately predicting returns rises with the amount of savings. Eventually, these forces dominate, leading to increasing information acquisition at high levels of wealth.

The comparative statics for *employed households* are similar to those of wealthy unemployed households: Employed households always have (relatively) low marginal utility of consumption compared to unemployed. The value of additional information about the state of the economy, as a consequences, starts off low and then rises with wealth.

Crucially, we show that the ability of our framework to match the survey evidence on household expectations depends on pre-cautionary effects tied to non-linear decision rules. This contrasts with previous analysis (summarized in e.g., Veldkamp, 2011) that, for tractability purposes, instead has focused on linear decision rules. Solving heterogenous-agent models with aggregate risk and non-linear decision rules is challenging. Our problem adds a further layer of complexity by allowing for heterogeneity also in expectations and higher-order beliefs. We develop a tractable method to tackle these challenges.[2]

We show that heterogeneity in information choices substantially changes the equilibrium properties of the economy relative to the full-information benchmark, in which all households have full information and hence common expectations about the state of the economy.

On the micro side, heterogeneous information choices feed back into wealth and income inequality, 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 little information are unable to exploit periods of good labor market prospects and high returns to build up financial wealth. The introduction of heterogeneous, incomplete information mitigates the lack of wealth inequality that exists in standard frameworks with incomplete markets relative to the data.

On the macro side, the presence of uninformed households leads to an increase in aggregate volatility, due to a stronger endogenous propagation of shocks. 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 to adjust, which makes household savings more procyclical and the economy more volatile. This mechanism is itself somewhat dampened by increased information acquisition, due the benefits of information about the economy being higher when the economy is more volatile. These dynamics elucidate a more general feature of our framework: Information acquisition decisions are *strategic substitutes*. In equilibrium, not all households acquire information in every period, leading to 11-16 percent higher fluctuations in consumption and output relative to the full-information case.

The consequences of macroeconomic policies may further be substantially different once one accounts for households' heterogeneous information choices. To demonstrate this, we consider the example of a wealth tax, modeled on the French tax system and the recent proposal in the US Congress.[3] 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

[2] Auclert *et al.* (2020) and Carroll *et al.* (2020) make the extreme opposite assumption by analyzing a model with non-linear policy rules, but exogenously heterogeneous information. They assume an exogenous process for household information based on Mankiw and Reis (2002) and Carroll (2003), and linearize the model in their solution. We consider the exogenous information case in Section 5 and 6, and show how the endogeneity of household information choices profoundly alters the macroeconomic consequences of incomplete information.

[3] For a description of the French wealth tax that used to operate, see, for example, `https://www.service-public.fr/particuliers/vosdroits/N20074`. The "Warren 2021 proposal" can be found here: `https://www.congress.gov/bill/senate-bill/510`.

is to reduce information acquisitions by 27 percent per quarter, as information acquisition on average rises in wealth. By reducing the information content in the economy, economic volatility increases by 4 percent. In contrast, in the full-information case, the tax has virtually no impact on volatility, despite a similar fall in aggregate wealth.

The effect of the wealth tax on household inequality in our benchmark economy is even more surprising: A one percent tax increases the Gini coefficient on wealth by over 3 percent, whereas in the full-information case the tax meaningfully decreases inequality (by over 5 percent). The reason inequality increases is because of the increased volatility, leading to larger over-accumulation 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 this lack of increase in inequality. However, our results suggest that the effects of dynamic, heterogenous information choice may substantially alter the relative costs and benefits of macroeconomic policies in unexpected directions. Thus, our findings, on balance, imply a Lucas-style critique (Lucas Jr, 1976) of policy evaluations in full-information, heterogenous-agent 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 choices (e.g., Bordalo *et al.*, 2016; Bordalo *et al.*, 2017; and Gabaix, 2019). Notwithstanding such alternative drivers, we show that households' rational incentives to systematically acquire different information fundamentally alter the dynamics and consequences of redistributive macroeconomic policies. We conjecture that behavioral heuristics, salience effects, and other behavioral drivers of information choices 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 states of the economy. 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 macroeconomic 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 solution and 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 additional results and analysis.

2 Motivating Evidence

We present new evidence on the relationship between household wealth and the accuracy of household expectations. We use micro data on household expectations from the *Survey of Consumer Expectations.* The SCE is a monthly panel of point and density forecasts for several macroeconomic and financial variables. In addition, the survey contains detailed data on household economic characteristics.[4] We link the monthly SCE expectation survey with the SCE's supplemental survey of household finances, which includes detailed data on household wealth and its composition. The merged SCE sample covers the period 2013M8-2020M1. Appendix A provides more information on the sample construction.

We explore the relationship between the accuracy of households' expectations and their wealth. To do so, we first focus on household forecasts of the one-year ahead unemployment rate, as unemployment represents the main source of income risk for many households. As such, perceived unemployment risk is a main driver of households' consumption and savings choices. We later include unemployment into our structural framework. We define a respondent's forecast error as the difference between the actual outcome and the respondent's forecast. A negative forecast error thus corresponds to an over-estimate of the variable. The SCE ask respondents for the "probability that the unemployment rate is higher 12-months from now". Unlike other variables (e.g., inflation) for which we can observe realized outcomes, the probability of unemployment rising is not objectively known. We proxy the true-but-unobserved probability of rising unemployment with the average probability computed from the *Survey of Professional Forecasters.* We make this choice because professional forecasters often provide more accurate predictions than even those from modern statistical and economic models.[5] We later show how our results are robust to other proxies of the probability of rising unemployment and extend to variables for which realized outcomes are objectively observed.

We begin by documenting a systematic correlation between household forecast errors and household wealth. Panel a in Figure 1 shows a marked, non-monotone relationship between household wealth and the accuracy of household expectations in the raw data. All else equal, wealthier households produce more accurate forecasts; however, in contrast to the results in Carroll (2003) and Vissing-Jorgensen (2003), this pattern is only discernible for households that are above the 20th percentile of the wealth distribution. The poorest households—those

[4] Armantier *et al.* (2017) provide an overview of the construction and scope of the Survey of Consumer Expectations, administered monthly by the Federal Reserve Bank of New York.

[5] See, for example, Stark *et al.* (2010), Faust and Wright (2013), and Bhandari *et al.* (2021). For interpretability reasons, we also scale the value of unemployment errors in the data with the average proxied probability of rising unemployment, to approximate the "Brier score" (see Appendix A).

Figure 1: Unemployment Expectations Across the Wealth Distribution

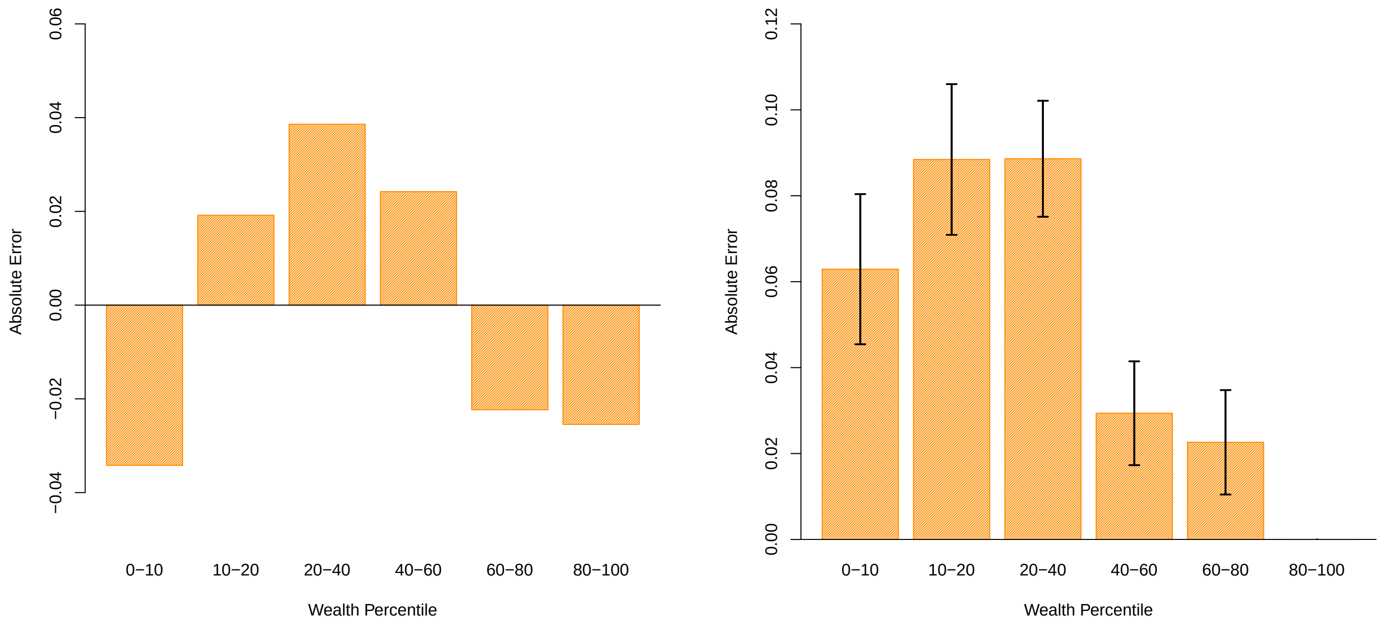

Panel a: Relative Accuracy

Panel b: Coefficient on Wealth

Note: Panel a plots the difference between the average one-year ahead accuracy of unemployment forecasts within wealth deciles/quintiles and the overall average taken across all wealth levels. Accuracy is measured by the absolute value of unemployment errors. Panel b plots the coefficient estimates on wealth from a regression of the absolute value of individual unemployment errors on the wealth decile/quintile the respondent belongs to, controlling for the age, education level, labor market status, and sex of the respondent, as well as time fixed effects. Estimates are relative to the wealthiest households, those in the 80-100 percentile of the wealth distribution. Whisker-intervals correspond to one-standard deviation robust confidence bounds (Table A.5). Sample: 2013M10-2020M1.

between the 0-10th percentile of the wealth distribution—produce unemployment forecasts that are of comparable accuracy to those from the wealthiest households. All else equal, this suggests that household expectations are *heterogenous* across the wealth distribution.

The relationship in Panel a in Figure 1 may be contaminated by other factors, such as labor-market status, that can simultaneously affect household wealth and the accuracy of expectations. To address this issue, Panel b in Figure 1 plots the coefficient estimates from a regression of the accuracy of individual expectations on the household wealth-decile/quintile controlling for household characteristics and time fixed effects. The figure confirms the relationship in the raw data. The accuracy of household expectations exhibits an inverse-u shape in wealth even when controlling for household characteristics. All else equal, wealthier households make more accurate unemployment forecasts, but the accuracy of households in the bottom decile is higher than those between the 20-40th percentile. The magnitudes are also meaningful: Moving a household from the 30th percentile of the wealth distribution to the 90th percentile, all else equal, increases the accuracy of the household's expectations by around 9 percent. To benchmark the magnitude, having a university degree is estimated to increase the accuracy of household expectations by only 7-8 percent (Table A.5). The evidence in Figure 1 is at odds with the common expectation assumption embedded in the full-information rational expectations framework, showing instead that the state of household finances matter for households' economic expectations.

We show that the systematic relationship between household wealth and the accuracy of expectations extends to other macroeconomic variables. We perform the same analysis for household forecasts of one-year ahead inflation and the growth rate of house prices. We use real-time data to measure the realizations of inflation and house prices, to capture the precise definition of the variable being forecasted. Figure 2 summarizes the estimates. For both variables, Figure 2 also includes the *perceived accuracy* of individual forecasts, as measured by respondents' interquartile range of the stated probability distribution of future outcomes.

All estimates show that wealthier households make more accurate forecasts and perceive themselves to be less uncertain. Apart from forecasts of future inflation, all measures of uncertainty are also higher for households close to the 20-40th percentile of the wealth distribution than for the poorest households. Furthermore, Figure A.10 in the Appendix shows that our results also extend to cases where we proxy the true probability of rising unemployment with those computed from standard forecasting VARs (Christiano *et al.*, 2005; Del Negro *et al.*, 2007). Finally, in Appendix A.2 we perform additional robustness exercises. There, we moreover show that, compared to professional forecasters, household expectations are systematically more dispersed, less accurate, and perceived to be more uncertain. We will leverage those additional moments to discipline our quantitative framework.

In summary, the results in this section provide evidence for systematic *heterogeneity* in the accuracy of household expectations. The data clearly reject the common-expectation assumption embedded in the full-information rational-expectation framework. Motivated by these findings, in the next section, we extend a workhorse incomplete-markets economy to allow for heterogeneity in the accuracy of household expectations. We then proceed by quantifying the impact of heterogenous expectations for positive and normative questions.

3 Model Framework

In this section, we describe a workhorse incomplete-markets model with idiosyncratic and 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 information about the state of the economy.

Figure 2: Inflation and House Prices Expectations Across the Wealth Distribution

Panel a: Inflation Forecasts

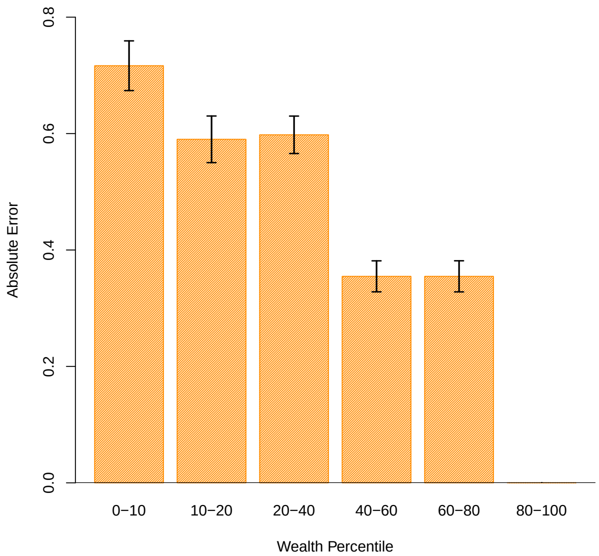

Actual Accuracy

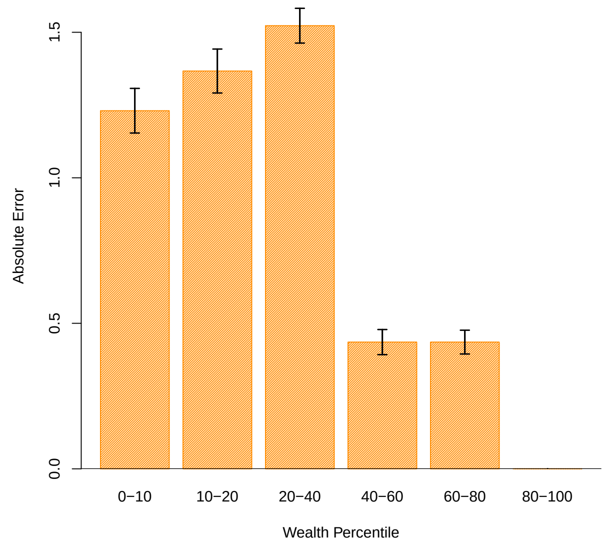

Perceived Accuracy

Panel b: House Price Forecasts

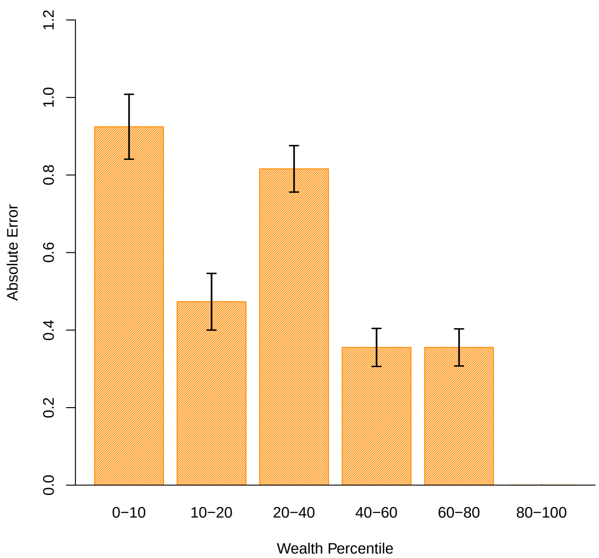

Actual Accuracy

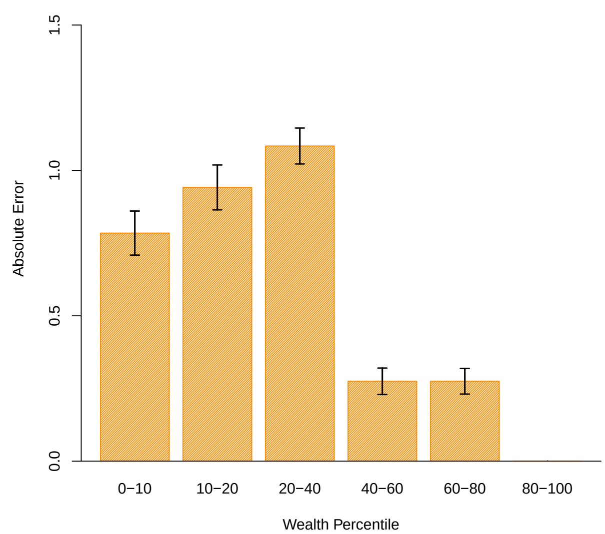

Perceived Accuracy

Note: Panel a and b consider the *actual* and *perceived* accuracy of individual forecasts of one-year ahead CPI inflation and the annual growth rate of US house prices, respectively. Both panels plot estimates on wealth from regressions of individual accuracy on the wealth decile/quintile the respondent belongs to, controlling for the age, education level, labor market status, and sex of the respondent, as well as time fixed effects. All estimates are relative to the wealthiest households, those in the 80-100 percentile of the wealth distribution. "Actual accuracy" corresponds to the absolute value of individual forecast errors, while "perceived accuracy" corresponds to the interquartile range of the reported probability distribution of the future outcome. Whisker-intervals are one-standard deviation robust confidence bounds. Sample: 2013M10-2020M1.

3.1 Households

The economy consists of a continuum of households $i \in [0, 1]$. There is only one consumption good per period $t = 0, 1, 2, ...$, and we assume that household preferences are described by the utility function:

$$\mathcal{U}_i = \mathbb{E}_{i0} \sum_{t=0}^{\infty} \beta^t \left[\frac{c_{it}^{1-\gamma} - 1}{1-\gamma} - \kappa_t \left(\mathcal{I}_{it} \right) \right], \tag{3.1}$$

where $\mathbb{E}_{i0} [\cdot] \equiv \mathbb{E} [\cdot \mid \Omega_{i0}]$ denotes household i's expectations conditional on its period-0 information set Ω_{i0}, $\beta \in (0, 1)$ the discount factor, c_{it} non-durable consumption, κ_t the utility cost of acquiring information $\mathcal{I}_{it}$, and $\gamma > 0$. The household's information set Ω_{it} accumulates according to $\Omega_{it} = \{\mathcal{I}_{it}, \Omega_{it-1}\}$.[6] The utility cost κ_t is distributed according to a type-I extreme value distribution with parameter α_k, and is i.i.d. across individuals and time. We introduce the extreme value shocks to account for unobserved heterogeneity in the survey data.

Each household is endowed with $\bar{l}$ units of time, which it supplies inelastically to the labor market. Labor productivity ϵ_{it} is stochastic and can take on two values $\epsilon_{it} \in \{0, 1\}$, which we interpret as unemployment and employment, respectively. We assume that ϵ_{it} follows a two-state, first-order Markov process $\Pi_{z_{t+1}, \epsilon_{it+1} | z_t, \epsilon_{it}}$, which depends on ϵ_{it} and aggregate total factor productivity z_t (described below). A household earns wage w_t when employed and receives unemployment benefits μw_t when unemployed, where the replacement rate $\mu \in (0, 1)$. 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 depreciation rate on capital. In addition to the borrowing constraint and a non-negativity constraint on consumption, household consumption choices are restricted by the per-period budget constraint:

$$c_{it} + k_{it+1} + \nu(\mathcal{I}_{it}) = r_t k_{it} + (1 - \tau_t) \epsilon_{it} w_t \bar{l} + \mu (1 - \epsilon_{it}) w_t + (1 - \delta) k_{it}, \tag{3.2}$$

where $\nu(\mathcal{I}_{it})$ denotes the monetary cost of acquiring information $\mathcal{I}_{it}$, and τ_t the tax rate on labor income. We refer to the right-hand side of (3.2) as a household's *cash-at-hand before choosing to acquire information*, and denote it by m_{it}. A household maximizes utility (3.1) subject to the budget constraint (3.2).

[6] We ultimately will analyze the model beginning at $t >> 0$, such that the economy has settled into its ergodic distribution and any effects of initial conditions wash out.

3.2 Technology and Markets

The production sector consists of a representative competitive firm, which maximizes profits. Output Y_t is produced using a Cobb-Douglas technology,

$$Y_t = z_t K_t^{\alpha} \left(L_t \bar{l}\right)^{1-\alpha}, \tag{3.3}$$

where K_t and L_t denote economy-wide capital and labor in period t, respectively. Aggregate total factor productivity z_t is stochastic and follows a first-order Markov process that takes two values $z_t \in \{z_l, z_h\}$ with $z_h > z_l$. The firm rents capital and labor in competitive markets, so that factor prices for labor w_t and capital r_t are given by their respective marginal products:

$$w_t = (1-\alpha) z_t \left(\frac{K_t}{\bar{l} L_t}\right)^{\alpha}, \quad r_t = \alpha z_t \left(\frac{K_t}{\bar{l} L_t}\right)^{\alpha-1}. \tag{3.4}$$

Finally, we assume that the share of households in a given idiosyncratic employment state only depends on the current value of total factor productivity z_t. Hence, the unemployment rate u_t is a function only of z_t, and thus only takes on two values u_h and u_l with $u_h < u_l$.

3.3 Government Policy

In our baseline analysis, the government runs a balanced-budget unemployment insurance scheme, such that $\tau_t = \frac{\mu u_t}{\bar{l} L_t}$. We consider a tax on household wealth in section 6.

3.4 Timeline and Information Structure

At the start of each period, idiosyncratic shocks $(\epsilon_{it}, \kappa_{it})_i$ and aggregate shocks z_t realize. Households then choose which signals $\mathcal{I}_{it}$ 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 does include elements of the state-space relevant for future prices (see below).[7] Next, firms rent capital and labor, production takes place, and factor payments are made. Finally, conditional on information choices, households make consumption and savings choices (c_{it} and k_{it+1}, respectively).

[7] An alternative approach is to instead allow households to flexibly design their optimal signal subject to a utility cost (e.g., Maćkowiak *et al.*, 2018). Such an optimal signal can, however, always be reduced to a signal of some combination of state variables, which the above approach, in principle, allows for. Furthermore, although the information-design approach has several advantages, it is computationally intractable for the non-linear, non-quadratic model that we study above (see also Section 3.5).

3.5 Recursive Formulation of 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 $\mathcal{P}_i(S)$.[8] Household i's second-order belief about household $j \neq i$'s belief is referred to as $\mathcal{P}_{ij}(S)$, and so on *ad infinitum*. Individual household beliefs are summarized by the object p_i, which includes the infinite-set of household (higher-order) beliefs. Let $\mathcal{P}$ denote the cross-sectional distribution of all such beliefs.[9] The aggregate state of the economy can then be described by $\Sigma = (S, \mathcal{P})$, while the individual state variables are described by $\sigma_i = (m_i, \epsilon_i, p_i)$, where m_i denotes household i's cash-at-hand before choosing to acquire information. We denote next period's realization of variable x by x' and previous period's realization by x_{-1}.

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

$$V\left(m_i, \epsilon_i, p_{i,-1}, \Sigma_{-1}\right) = \max_{\mathcal{I}_i} \mathbb{E}\left[W\left(m_i - \nu\left(\mathcal{I}_i\right), \epsilon_i, p_i, \Sigma\right) - \kappa(\mathcal{I}_i) \mid \Omega_{i,-1}\right], \tag{3.5}$$

where $V(\cdot)$ and $W(\cdot)$ denote a household's 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 the information choice $\mathcal{I}_i$. We note that households' expectations in the first stage are computed using previous period's posterior beliefs $p_{i,-1}$, 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., $\Sigma = H(\Sigma_{-1}, z, (\mathcal{I}_i)_i)$), and the exogenous transition matrix for z, Π_z, to form a prior about today's state variables from yesterday's posterior.

The assumption of type-I extreme value shocks for the utility cost of information implies a parsimonious logistic choice function for the probability of acquiring $\mathcal{I}_i \in \emptyset \cup \mathcal{I}^{\max}$:

$$\text{Prob}\left(\mathcal{I}_i \mid m_i, \epsilon_i, p_{i,-1}, \Sigma_{-1}\right) = \frac{e^{\mathbb{E}\left[W(m_i - \nu(\mathcal{I}_i), \epsilon_i, p_i, \Sigma) - \kappa(\mathcal{I}_i) \mid \Omega_{i,-1}\right]}}{\sum_{\tilde{\mathcal{I}} \in \emptyset \cup \mathcal{I}^{\max}} e^{\mathbb{E}\left[W\left(m_i - \nu\left(\tilde{\mathcal{I}}\right), \epsilon_i, p_i, \Sigma\right) - \kappa(\tilde{\mathcal{I}}) \mid \Omega_{i,-1}\right]}}, \tag{3.6}$$

yielding the standard value function (McFadden *et al.*, 1973):

$$V\left(m_i, \epsilon_i, p_{i,-1}, \Sigma_{-1}\right) = \frac{\gamma^E}{\alpha^\kappa} + \frac{1}{\alpha^\kappa} \log\left(\sum_{\tilde{\mathcal{I}} \in \emptyset \cup \mathcal{I}^{\max}} e^{\mathbb{E}\left[W(m_i - v(\tilde{\mathcal{I}}), \epsilon, p, \Sigma) - \kappa(\tilde{\mathcal{I}}) \mid \Omega_{i,-1}\right]}\right), \tag{3.7}$$

[8] Not to be confused with the Powerset, $\mathcal{P}_i$ here has a distribution with $\hat{\Gamma}_i$ and $\hat{z}_i$ as its typical elements, representing household i's first-order belief about the mass of capital and employment status at some point, as well as the household's beliefs about productivity, respectively. $\mathcal{P}_i$ is hence a distribution over distributions.

[9] More formally, we can describe $p = \left\{\mathcal{P}_i, (\mathcal{P}_{ij})_{j \in [0,1]}, ..., (\mathcal{P}_{ij...k})_{j,...,k \in [0,1]^{n-1}}, ...\right\}$, while the cross-sectional distribution of such beliefs $\mathcal{P} = \left\{(\mathcal{P}_i)_{i \in [0,1]}, (\mathcal{P}_{ij})_{i,j \in [0,1]^2}, ..., (\mathcal{P}_{ij...k})_{i,j,...,k \in [0,1]^n}, ...\right\}$.

where γ^E is the Euler-Mascheroni constant.

After deciding on information choices, households choose consumption c_i and savings k_i' out of cash-at-hand net of information acquisition costs, y_i:

$$
\begin{aligned}
W(y_i, \epsilon_i, p_i, \Sigma) &= \max_{c_i, k_i' \geq 0} \frac{c_i^{1-\gamma} - 1}{1-\gamma} + \beta \mathbb{E}\left[V(m_i', \epsilon_i', p_i, \Sigma) \mid \Omega_i\right] \\
\text{subj. } &\text{ to} \\
c_i + k_i' &= y_i \\
m_i' &= r(\Sigma') k_i' + (1-\tau) \epsilon_i' w(\Sigma') \bar{l} + \mu (1-\epsilon_i') w(\Sigma') + (1-\delta) k_i',
\end{aligned}
\tag{3.8}
$$

where the expectation is taken using today's updated information set Ω_i. We let $g(\cdot)$ denote the function that characterizes a household's savings choice $k_i' = g(\sigma_i, \Sigma)$, and $\iota(\cdot)$ the function that characterizes its information choice $\mathcal{I}_i = \iota(\sigma_i)$. Finally, today's posterior beliefs p_i are linked to yesterday's $p_{i,-1}$ through Bayes' Rule and the information choice $\mathcal{I}_i$.

3.6 Recursive Incomplete Information Competitive Equilibrium

The definition of a *Recursive Competitive Incomplete Information Equilibrium* (RIICE) straightforwardly extends the standard definition of a Recursive Competitive Equilibrium to the case with incomplete information: A RIICE is a law of motion $H(\cdot)$, a pair of individual value functions $V(\cdot)$, $W(\cdot)$, policy functions $\iota(\cdot)$ and $g(\cdot)$, as well as pricing functions $(r(\Sigma), w(\Sigma))$ such that: (i) $V(\cdot), W(\cdot), \iota(\cdot), g(\cdot)$ solve the household's optimization problem, (ii) $r(\cdot)$ and $w(\cdot)$ satisfy firm maximization, (iii) $H(\cdot)$ is generated by policy functions $\iota(\cdot)$ and $g(\cdot)$, the Markov processes $\Pi_{z',\epsilon'|z,\varepsilon}$ and Π_z, as well as Bayes' Rule, using the information contained in $(\mathcal{I}_i)_i$ and current beliefs described in $\mathcal{P}$.

4 Solution Method and Calibration

In this section, we first outline our procedure for computing RIICE equilibria. Our description here is non-technical, and we only include it in the main text because it is intimately linked to the two-way feedback mechanism between information choice and the macroeconomy that is at the heart of our analysis. We then proceed to discuss the specification of the information structure and the calibration of the model.

4.1 Computational Strategy

The endogenous aggregate state variables of the economy, Γ and $\mathcal{P}$, are infinite-dimensional objects. Even the full-information version of our incomplete-markets economy therefore presents a computational challenge, because of the high-dimensionality of Γ (the endogenous state variable of that model). Our incomplete-information framework has a *double-infinity problem*—the additional complexity arising from the entire set of (higher-order) beliefs $\mathcal{P}$, in principle, mattering for equilibrium dynamics, depending on the specification of $\mathcal{I}^{\max}$.

The standard strategy for computing incomplete-markets models *without* incomplete information involves approximating the distribution Γ with a finite set of moments $\mathbf{m} \equiv (m_1, m_2, \cdots, m_n)$ (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 the Krusell-Smith solution method is one of "boundedly rational" expectations, as households only keep track of a limited set of moments of the distribution. Importantly, in this solution method, the information that households use to base their expectations on is *exogenously* predetermined by the researcher—containing productivity z and the moments in $\mathbf{m}$. By contrast, in our model, households *optimally choose* the information on which to form their '"boundedly rational'" expectations. Thus, one can interpret our model as a natural framework to study incomplete-markets models with aggregate risk, since we provide a micro-foundation for the boundedly-rational solution based on costly information choice. In particular, the Krusell-Smith solution can be seen as the special case in which the cost of information is zero, and $\mathcal{I}_{it} = \mathcal{I}_t^{\max} = \{z_t, \mathbf{m}_t\}$ for all i and t, as a consequence.

In addition to exogenous information, notice that the Krusell-Smith framework also imposes *common knowledge* over both z and Γ, as all households' form expectations using the same information. Our framework relaxes the common knowledge assumption over the state variables by allowing for heterogenous information choices.[10] The RIICE framework therefore allows for the study of the three-way interaction between incomplete common knowledge, aggregate dynamics, and inequality on which we focus on below—in contrast to Krusell-Smith.

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

[10] Both our framework and the Krusell-Smith framework assume common knowledge over the underlying structure of the economy (preferences, technology, etc). In particular, the law of motion $H(\cdot)$ is common knowledge, such that if households knew the current state of the economy, they would correctly forecast tomorrow's state. That is why the Krusell-Smith solution can be see as a special case of our equilibrium concept when the cost of information is zero.

4.2 The Specification of Moments

For the set of moments that households can choose to use to forecast future wages and returns, we follow Krusell and Smith (1998)) and consider only the first moment $\mathbf{m} = \int g(\sigma_i)\,\Gamma(d\sigma_i) = K_t$. Even with this restricted set, the model, in principle, suffers from the problem of "*infinite regress of expectations*", described in e.g., Townsend (1983), which is induced by the public observation of an endogenous market-outcome. To solve this problem, we exploit a feature of Krusell-Smith-like economies: The sequence of shocks $\{z_s\}_{s=0}^{t}$ alone allows for very accurate predictions about the future capital stock K_{t+h}, $h \geq 1$(Den Haan *et al.*, 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 value of productivity z_t. Importantly, we check *ex post* that this assumption allows households to form accurate posteriors about K_t, and thereby make accurate predictions about future wages and rates of returns (Appendix B.2).

A main result in the literature on optimal signal design is that optimal signals can be reduced to signals about (some combination of) state variables (e.g., Maćkowiak *et al.*, 2018). Because the history of aggregate productivity z_t accurately approximates the relevant state variables for prices in our economy (Appendix B.2), our assumption can be viewed as allowing households to choose to observe elements of the optimal signal.[11] Finally, in previous work (Broer *et al.*, 2022), we show that the utility benefits of acquiring information about the capital stock K_t (or equivalent past productivity) conditional on current productivity, z_t, assuming a Markov process for z_t, are small—in the order of $3-30 at 2020 prices.

4.3 Approximated Problem and Equilibrium

Given our assumptions, we can state the approximated household problem that households solve: Households enter the period with cash-at-hand, m_i, their employment status, ϵ_i, their prior over whether the economy is in the high productivity state, $p_{i,-1}^z \equiv \mathrm{Prob}\,(z = z_h \mid \Omega_{i,-1})$, and their prior over the capital stock $p_{i,-1}^K$. Households can then choose whether to observe contemporaneous productivity z. The two-stage optimization problem can now be stated as:

Stage 1: Information choice

$$\begin{aligned} \tilde{V}\left(m_i, \epsilon_i, p_{i,-1}^z, p_i^K\right) &= \max_{\mathcal{I}_i \in \{\emptyset, z\}} \mathbb{E}\left[\tilde{W}\left(m_i - \nu(\mathcal{I}_i), \epsilon_i, p_i^z, p_i^K\right) - \kappa(\mathcal{I}_i) \mid p_{i,-1}^z, p_i^K\right] \\ p_i^z &= \mathrm{Prob}\,(z = z_h \mid \Omega_i), \quad \Omega_i = \{\mathcal{I}_i, \Omega_{-1}\}. \end{aligned} \tag{4.1}$$

[11] We thank Mirko Wiederholt for this comment.

Stage 2: Consumption-savings choice

$$
\begin{aligned}
\tilde{W}\left(y_i, \epsilon_i, p_i^z, p_i^K\right) &= \max_{c_i, k_i' \geq 0} \frac{c_i^{1-\gamma}-1}{1-\gamma} + \beta \mathbb{E}\left[\tilde{V}\left(m_i', \epsilon_i', p_i^z, p_i^{K\prime}\right) \mid p_i^z, p_i^{K\prime}\right] \qquad (4.2)\\
\text{subj. to} & \\
c_i + k_i' &= y_i \\
m_i' &= r\left(z', K'\right) k_i' + (1-\tau)\, \epsilon_i' w\left(z', K'\right) \bar{l} + \mu\left(1-\epsilon_i'\right) w\left(z', K'\right) + (1-\delta) k_i' \\
K' &= \tilde{H}(z, K) \\
p_i^{K\prime} &= \mathbb{E}[\tilde{H}(z, p_i^K) | \Omega_i],
\end{aligned}
$$

where $\tilde{H}(z, K)$ is the law of motion for the aggregate capital stock, which replaces $H(\cdot)$ as the aggregate law of motion in the approximated problem. Thus, in addition to their individual cash-at-hand and employment status, households forecast the probability of being in the high productivity state and the aggregate capital stock. They update their priors after their information acquisition decision using Bayes' rule.

We provide a brief overview of the numerical procedure that we use to solve for the (approximated) RIICE equilibrium. To compute the equilibrium, we use an iterative procedure to solve for the equilibrium fixed point: First, we postulate a law of motion $\tilde{H}(\cdot)$ for the aggregate state variables. Second, we solve the household's two-stage problem in (4.1) and (4.2) conditional on $\tilde{H}(\cdot)$ 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 and K, and estimate a new law of motion $\tilde{H}'(\cdot)$. We iterate until convergence on $\tilde{H}(\cdot)$.

4.4 Calibration

The aim of our calibration exercise is to ensure that the model can account for salient business cycle facts, as well as capture the rich heterogeneity in household expectations documented in Section 2. We assume that a model period corresponds 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, as those more closely align to our model framework than traditional NBER-dated recessions. We define a boom as a period with a below-trend unemployment rate.[12] 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 $z_h/z_l = 1.027$. The individual transition probabilities in labor productivity ϵ_{it} are set to match US labor market transitions calculated from the Current Population Survey. We choose an unemployment rate in booms and busts equal to 6 and 10 percent, respectively. Monthly job-finding rates are set to match unemployment to employment flows in the CPS, and are 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 z_t. Finally, we set the UI replacement rate μ to 0.40.

Internally Calibrated Parameters 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 γ and the information cost parameters α_κ and ν to quantitatively capture key features of the micro-data on expectations discussed in Section 2. We set $\gamma = 5$ and the monetary cost per signal equal to $\nu = 0.0012$ (equivalent roughly to 0.05 percent of pre-tax wages) to match our empirical finding that forecast accuracy increases in wealth for richer households (Section 5.3).[13] We set the scale parameter α_κ equal to $1/3 \times 10^{-8}$ to capture the dispersion in unemployment expectations observed in the SCE (even for households with similar observable characteristics). To see how household expectations compare to those in the SCE, we concentrate on expectations of future unemployment. Table I compares the accuracy and standard deviation of households' one-year ahead unemployment rate errors in the model and the data. Recall that 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". For households in the model, we therefore compute the difference between a household's perceived probability conditional on its current information $\text{Prob}(u_{t+4} > u_t|\Omega_{it})$ and the true probability $\text{Prob}(u_{t+4} > u_t|z_t)$, which depends on current productivity z_t. We then compare the resulting errors with the corresponding errors in the survey data. Table I shows that the dispersion of errors is somewhat smaller than in the data, but overall the model replicates both the accuracy and volatility of expectation errors in the data well. Table B.1 in the Appendix summarizes the parameters.

[12] We use an HP filter with smoothing parameter λ equal to 14.400 to construct the trend in the unemployment rate from monthly unemployment data.

[13] The 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: Unemployment Expectations: Model vs. Data

	Mean Abs. Error	Std. Dev. of Abs. Error
Survey of Consumer Expectations	0.63	0.52
Model Simulated Data	0.62	0.45

Note: The table shows the mean and standard deviation of the absolute value of errors in the probability that the unemployment rate four-quarters ahead is higher than at time t. The table compares the simulated moments from the calibrated model to those from the Survey of Consumer Expectations (see Section 2 for a description). For interpretability reasons, we scale the absolute value of unemployment errors in the model and in the data with the average true probability of rising unemployment, proxied in the data with average probability of rising unemployment from the Survey of Professional Forecasters (Section 2).

5 Quantitative Results

In order to understand the consequences and drivers of households' information choices, we proceed in four steps. First, we show how different information choices affect households' savings decision, households' inter-temporal decision variable in the model. Second, we characterize how household information choices depend on individual state-variables. Third, we combine the insights from the first two steps to show how the interaction between information and savings choices allow us to match the micro data on expectations. Finally, we explain the impact of the wealth-expectations nexus on aggregate dynamics and inequality.

5.1 Savings Choices and Information

We start by analyzing how the acquisition of information about productivity affects households' savings decisions. In the model, households have two reasons to save: to smooth consumption in the face of volatile income; and to intertemporally substitute consumption in response to movements in rates of return. The state of the economy that households can acquire information about affects these both through exogenous fluctuations in productivity and employment, as well as through endogenous fluctuations in the capital stock.

We consider the polar cases of a household who has just acquired information ("informed") and a household who has a 50-50 prior ("uninformed") over productivity. We assume both have the same prior over the capital stock. In Figure 3, we plot the difference in savings choices ("informed minus uninformed") as a function of a household's cash-at-hand and its employment status. We interpret this measure as the "static cost" of not being informed, as we keep the prior over capital fixed. Dynamically, however, not being informed will lead to mean-biased expectations of the capital stock, which we consider below.

Figure 3: Household Saving Choices and Wealth

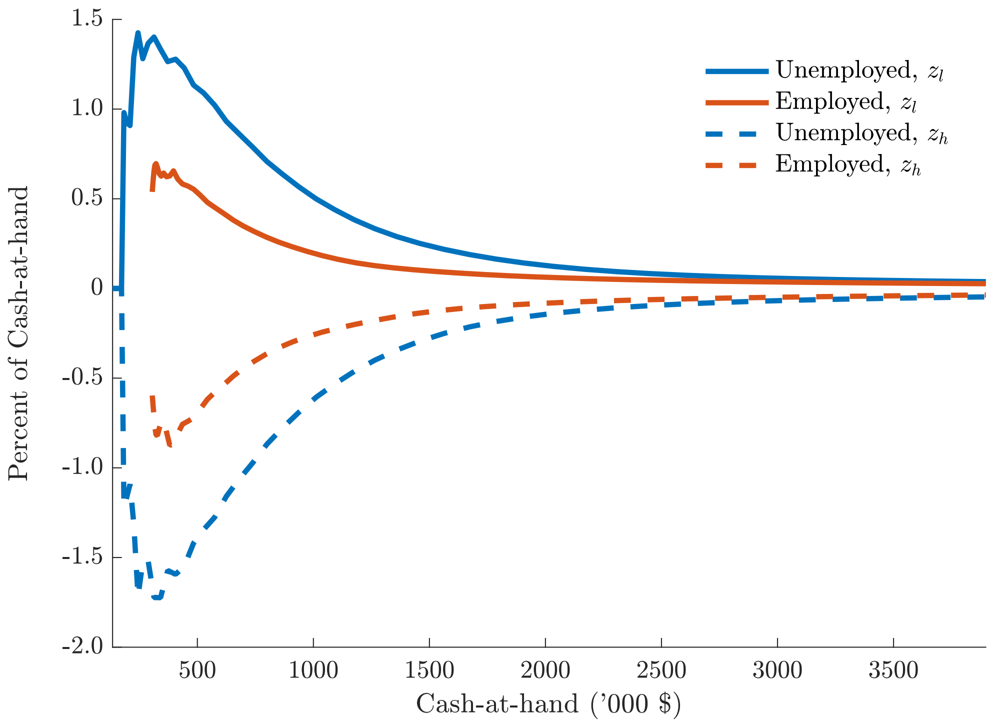

Note: The figure plots the difference between savings choices by "informed households" and "uninformed households" at a mean prior for aggregate capital K_t. We plot household savings choices as a function of individual cash-at-hand, using our benchmark parameterization of the model. We use 2020 values of US household income to convert cash-at-hand in our model to $ amounts.

All else equal, informed households save more than uninformed in recessions (z_l), as they know that the probability of becoming unemployed (or staying unemployed) is higher. Conversely, informed households save less in booms (z_h). As Figure 3 shows, the percentage difference in savings is, however, strongly non-linear 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 cash-at-hand, where the precautionary-savings motive is strongest. 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: Being better able to predict the difference in idiosyncratic risk between booms and recessions is less valuable when the household is wealthy.

Notice that in the thought experiment we conduct in Figure 3 informed and uninformed households have the same prior over the capital stock. Thus, the perceived difference in wages and rates of return from being informed stem only from the perceived differences in

Figure 4: Household Saving Choices and the Prior over Capital

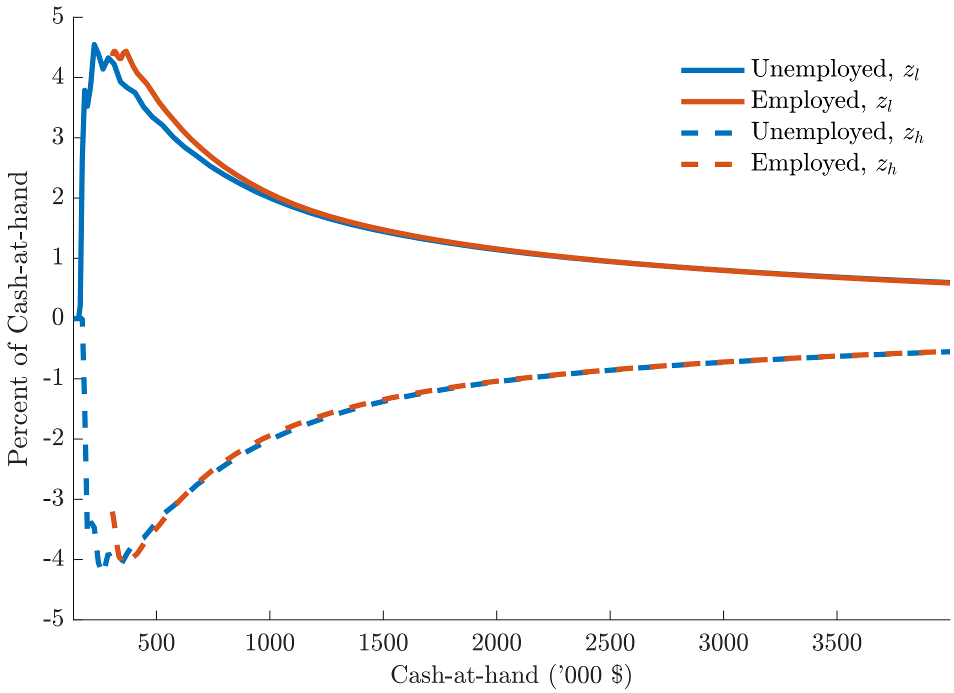

Note: The figure plots the difference between savings choices by "informed households" and "uninformed households". Specifically, in a boom (z_h), we plot the difference between the savings policy functions of an informed household that has a prior over the capital stock K_t that is one standard deviation higher than the mean (as over time in booms the capital stock rises) and an uninformed household with a prior at the mean capital stock (solid lines). We do the same for a recession (z_l), but where the informed household has a prior that is one standard deviation below the mean (dashed lines). We plot household savings choices as a function of individual cash-at-hand, using our benchmark parameterization of the model. We use 2020 values of US household income to convert cash-at-hand in our model to $ amounts.

productivity and labor supply. With incomplete information, however, uninformed households also perceive less accurately the dynamics of the capital stock (that rises in booms and falls in recessions). Their expectations about future capital, returns and wages, are less cyclical and biased towards the unconditional mean. Over time, households that are worse informed will hence have less accurate priors of the capital stock, which leads to less accurate savings decisions. We illustrate this mechanism in Figure 4. Specifically, in a boom we plot the difference between (i) the savings policy function of an informed household that has a prior over the capital stock that is one standard deviation higher than the mean (as over time in booms the capital stock rises); and (ii) an uninformed household with a prior at the mean capital stock (solid lines). We do the same for a recession, but where the informed household has a prior that is one standard deviation below the mean (dashed lines).

Compared to Figure 3, having mean-biased expectations increases the magnitude of savings mistakes that all households make (except for the borrowing constrained). The effect of biased expectations are most pronounced for the low-wealth, employed households, whose savings mistakes are amplified by almost an order of magnitude. This is because the errors in household expectations of future wages now combines the effects of productivity and capital accumulation. In addition, having mean-biased capital expectations also leads to biased expectations of the real interest rate. That leads to significant savings mistakes even at high-wealth levels, as uninformed households overpredict returns when the capital stock is high and actual returns low, and vice versa. Combined, Figure 3 and 4 illustrate how acquiring information "statically and dynamically" affects households' savings decisions. Next, we explore how those two forces interact to shape households' information acquisition decision. This will cast further light on the wealth-expectations nexus.

5.2 Household Information Choices

Having characterized how information affects households' savings decision, we proceed to illustrate how the incentives to improve savings decisions shape information choices. Decisions are most easily described by the probabilities of information acquisition. 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 5. Combined, the results in Figure 5 showcase the rich heterogeneity that exists in the incentives to acquire information, depending on the precise realization of individual state-variables. Unsurprisingly, households with less informative prior expectations (closer to one-half) are more likely to acquire information for all wealth levels (except those at the borrowing constraint).

Employed households—the main 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 do not outweigh the benefits. The potential savings mistakes are small, as illustrated in the previous section. They also know that in the event of future job losses 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 cost of acquiring information relative to wealth falls and the benefit of predicting returns on increasing financial wealth rises, increasing the information acquisition probability.

Now, consider instead *unemployed households*. The unemployed are dissavers in the model, as they attempt to smooth consumption, and when their cash-at-hand falls low enough they end up at the borrowing constraint. Hence, at low values of cash-at-hand, unemployed households almost never acquire information. Those households choose to be at the borrowing constraint in all future states of the world, hence have no benefit from acquiring information. However, as their cash-at-hand begins to rise they rapidly start to acquire information. The

Figure 5: Information Acquisition Probabilities

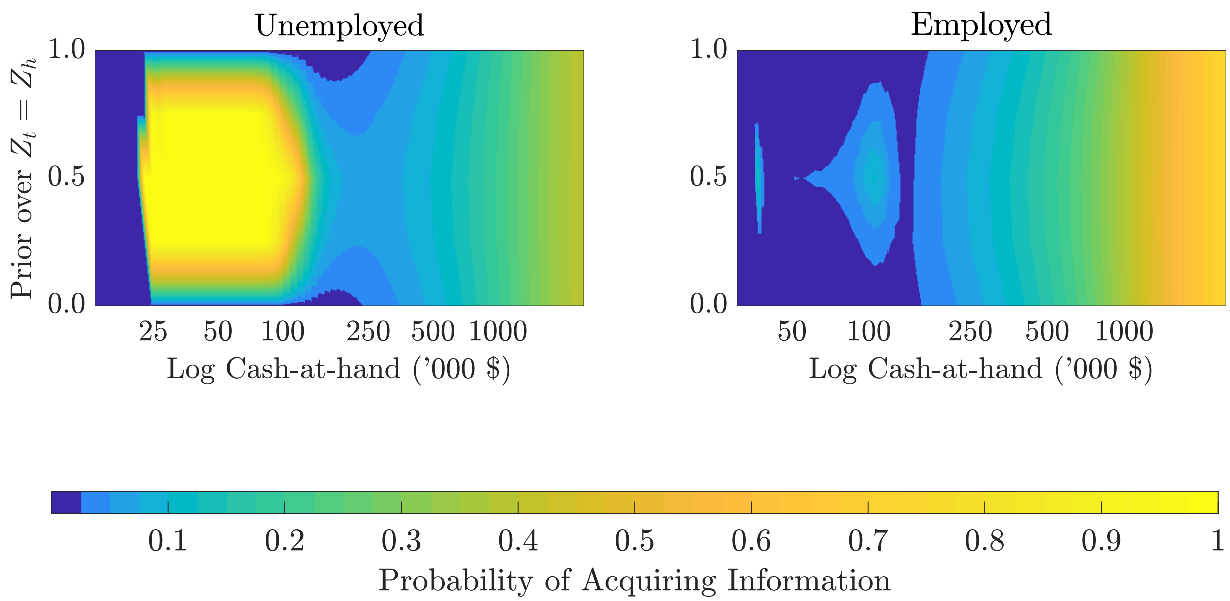

Note: The figure plots the probability of information acquisition for different values of the prior belief that current productivity z_t is high and for different values of quarterly household cash-at-hand. The figure uses our baseline calibration to depict the probabilities and evaluates them at the mean prior about the aggregate capital stock K_t. The left-hand (right-hand) panel shows the probabilities for an unemployment (employed) household. We use 2020 values of US household income to convert cash-at-hand in our model to $ amounts.

cost of making a savings mistake close to the borrowing constraint is high (due to marginal utility of consumption being high), so those households almost uniformly choose to acquire information. As wealth rises further, marginal utility falls, the savings policy function becomes approximately linear, and the value of information initially drops; the household is no longer at risk of imminently hitting the borrowing constraint due to a savings mistake. The value of information, however, then slowly starts to rise again with wealth for the same reasons as discussed in the case of the employed.

In summary, the left-hand and right-hand panels of Figure 5 illustrate the complex incentives that households face when deciding to acquire information. We next turn to how these forces interact to shape the joint distribution of wealth and expectations.

5.3 Accuracy of Expectations

We have described how wealth and employment status affect a household's decision to acquire information, and how a household's savings decision is, in turn, affected by the accuracy of its information. Before we turn to the macroeconomic consequences of the two-sided interaction between household heterogeneity and information choice, we study how these forces interact

Figure 6: Accuracy of Forecasts: Model vs. Data

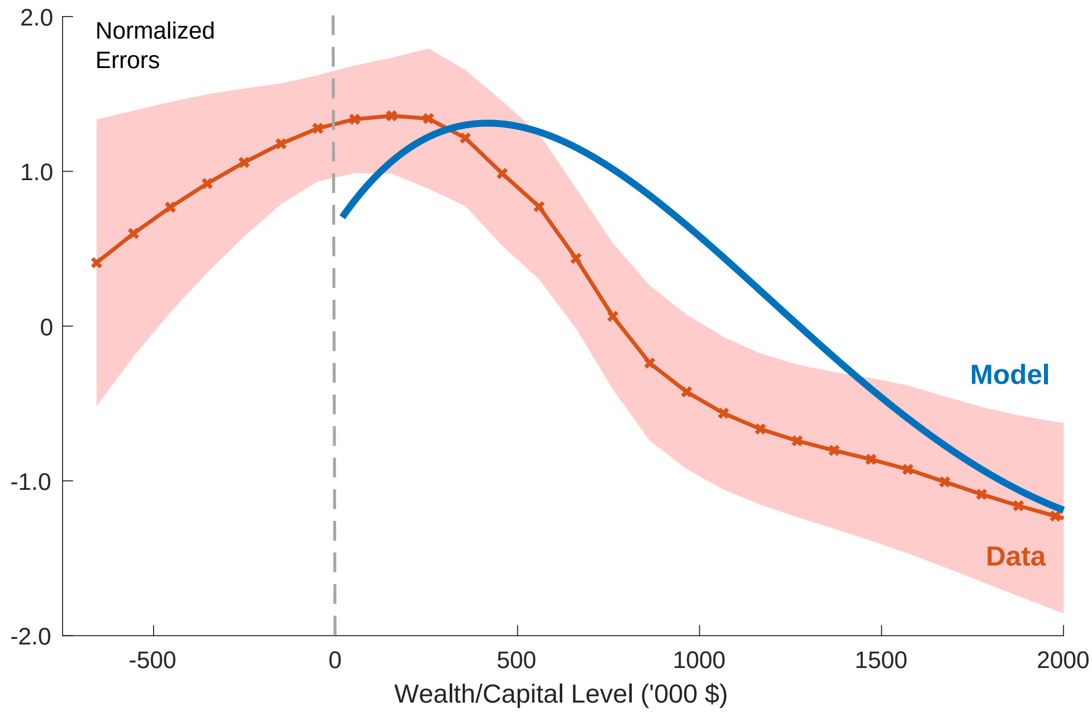

Note: The figure depicts the estimated relationship between (the absolute value of) normalized errors of the one-year ahead probability of the unemployment rate increasing and household wealth. We plot this relationship both in the SCE data and in the calibrated model (see also Section 2). We use a local polynomial regression (the LOESS regression) to estimate the non-linear relationship between the accuracy of household expectations and household wealth. Error bands correspond to one-standard deviation confidence bounds. We use 2020 values of US household income to convert values in the data and in the model to $ amounts.

to shape the accuracy of household expectations across the wealth distribution. This will allow us to confront our model with the empirical evidence discussed in Section 2.

Figure 6 shows how households' information acquisition probabilities in equilibrium translate into a systematic relationship between the accuracy of households' expectations and their wealth level. Because the model matches the mean and standard deviation of absolute errors, Figure 6 plots normalized errors in both the model and in the data. Although our model cannot speak to the positive slope that exists in the data for households with negative wealth—recall that we assume a simple no-borrowing limit $k' \geq 0$ for households—the model generates an inverse-u shape, which, on balance, resembles that in survey data.

The inverse-u shape in the model is a result of two opposing 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 observable decline in forecast accuracy (increase in the average absolute error) at low-levels of wealth. Second, as wealth increases, the probability of acquiring information eventually becomes monotonic in wealth (Figure 5). These are in effect the employed households that we discussed in Section 5.2. As wealth increases, the costs of making savings mistakes rise, and households' information acquisition probabilities increase. This then leads to the eventual rise in forecast accuracy (decline in average absolute errors) visible in Figure 6. We conclude that the model, on balance, matches the salient features of the relationship between the accuracy of household expectations and household wealth, making it a suitable laboratory to explore the effect of expectation heterogeneity on the macroeconomy.

5.4 Aggregate Implications

Our analysis so far has centered on the dynamics of savings and information for an individual household. In this subsection, we discuss how these decisions impact business-cycle dynamics.

In Table II, we contrast the aggregate dynamics of our benchmark economy with those that arise in an economy in which all households acquire information every period. We henceforth refer to this counterfactual economy as the "full-information economy". For comparison, on average, only around 13 percent of households choose to acquire information in any given period under our benchmark calibration.[14] We also compare the dynamics from our model with those that arise from an economy in which households *exogenously* face a 13 percent probability of information acquisition in every period à la Mankiw and Reis (2002). We refer to this economy as the "exogenous-information economy".

Relative to the full-information counterpart, fluctuations in all aggregate variables are substantially more pronounced in our baseline economy with incomplete information. The standard deviation of the capital stock is around 2/3 higher, and output and consumption are, as a result, 11-16 percent more volatile. The stark difference is caused by uninformed households' savings choices—the main driver of our economy. Dynamically, as discussed in the Subsection 5.1, households who choose *not* to acquire information have expectations about the level of the capital stock that are more tilted towards the long-run average level of capital, the unconditional mean of the capital stock. In booms, such households systematically underpredict the capital stock, and vice versa in recessions.[15] Because uninformed households systematically underpredict the capital stock in booms, they overpredict the return on savings, and hence save more than if they had full information, as illustrated previously in Figure 4. The converse is true in recessions. In equilibrium, the economy, therefore, systematically

[14] The probabilities of information acquisition in the benchmark model are 0.1437 and 0.1295 for the unemployed and employed, respectively. The average rate of unemployment is 8 percent in our model.

[15] We should note, however, that a hypothetical household that acquires information in every period would have an accurate estimate of the true capital stock and make negligible forecast errors (Appendix B.2).

Table II: Business Cycle Moments

	Panel a: Level of Moments					
	$\sigma(K)$	$\sigma(Y)$	$\sigma(I)$	$\sigma(C)$	$Cor(C,Y)$	$Cor(I,Y)$
Benchmark Model	6.04	3.65	11.19	1.45	0.72	0.97
Exogenous Information	5.07	3.43	11.47	1.37	0.69	0.96
Full Information	3.61	3.15	10.54	1.31	0.71	0.96
	Panel b: Percent Difference to Full Information					
	$\sigma(K)$	$\sigma(Y)$	$\sigma(I)$	$\sigma(C)$	$Cor(C,Y)$	$Cor(I,Y)$
Benchmark Model	67.31	15.87	6.17	10.69	1.41	0.46
Exogenous Information	40.44	8.88	8.82	4.58	-2.82	0.19

Note: The table shows the standard deviation σ of the logarithm of economy-wide capital (K), output (Y), investment (I), and consumption (C). In addition, the table shows that correlation between aggregate consumption, investment, and output, respectively (e.g., $Cor(I,Y)$). The table computes these moments in the calibrated model ("Benchmark Model"), the associated full-information economy ("Full Information"), as well as in a model with an exogenously specified probability of acquiring information ("Exogenous Information"). The probabilities of information acquisition in the benchmark model are 0.1437 and 0.1295, respectively, for the unemployed and employed. This probability is set equal to 0.1306 in the exogenous information case.

"overaccumulates" capital in booms and "underaccumulates" in recessions, leading to much larger fluctuations in output, consumption, and investment.

Compared to the exogenous-information model, where all households have the same probability of acquiring information, the endogeneity of information choice that is a feature of our benchmark economy amplifies the increase in volatility (Table II). All else equal, under our baseline calibration, the uninformed households are the middle-to-rich households (Section 5.2 and 5.3), who combine to hold most of the capital stock in our model—around 64 percent of the capital stock is held by households with a wealth level between \$400,000 and \$1,000,000. These are the households, on average, for which the absence of information leads to mean-biased expectations. As a consequence, the weaker mean-version of the capital stock is heightened in our baseline economy relative to an economy in which all households have the same probability of acquiring information. The heterogeneity of information amplifies the consequences of incomplete information.

We conclude that the presence of heterogeneous, incomplete information serves as an amplifying force—it induces weaker mean-reversion of the capital stock relative to the full-information economy. This creates additional endogenous propagation of shocks. These dynamics elucidate a more general feature of our framework: Information acquisition choices are *strategic substitutes.* The individual benefits of information rise with the volatility of the capital stock. But, when the average share of information in the economy increases, the volatility of the capital stock falls, and so does the incentive to acquire information. In previous work, Broer *et al.* (2022), we show how this may imply non-existence of homogeneous-information (representative-agent) equilibria in neoclassical economies.

5.5 Distributional Implications

The weakening of the mean-reversion of the capital stock interacts with households' information choices to also affect wealth inequality in the economy. Figure 7 and Table III contrast the wealth distribution in our calibrated benchmark with that from the equivalent full-information economy. Table III also shows the moments of the wealth distribution for the economy in which we exogenously fix the probability of information acquisition.

On balance, the introduction of heterogenous, incomplete information increases the dispersion of the wealth distribution—with more mass placed at the bottom and near the top of the wealth distribution (Figure 7). Conversely, there are fewer middle-income as well as extremely wealthy households. The latter, in this case, being defined as households with more than \$2.5 million in wealth. Consistent with the widening of the wealth distribution, the 90/10-ratio of the wealth distribution increases with the introduction of heterogenous, incomplete information (Table III). Yet, because the Gini coefficient is sensitive to the behavior of

Figure 7: Information and the Wealth Distribution

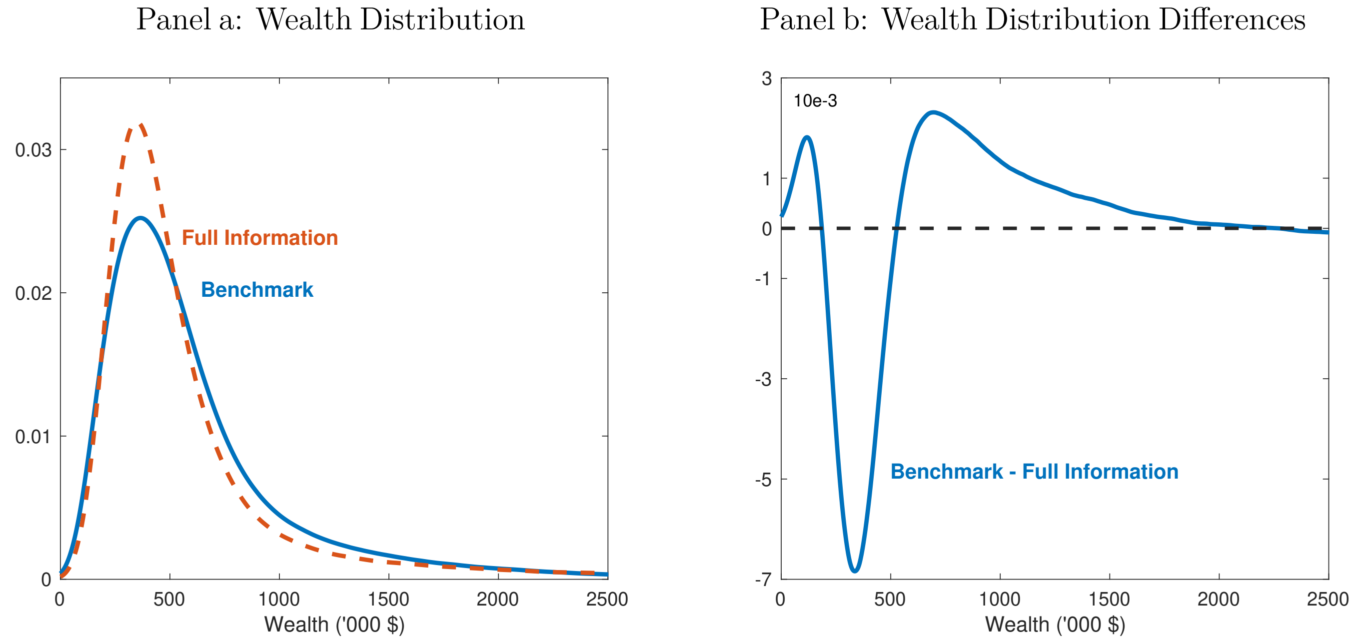

Note: The figure illustrates changes in the average wealth distribution relative to the full-information version of the benchmark economy. Panel a depicts the average probability density function of the wealth distribution in the two cases, while Panel b depicts the difference between the two distributions. The horizontal axis in both panels is household wealth (capital levels) in '000s \$. We use 2020 values of US household income to convert values in the model to \$ amounts. Probability density functions are estimated from a simulated panel of households, using a kernel density estimator with the Epanechnikov kernel.

the extreme right-tail (Cowell and Flachaire, 2002), the Gini coefficient actually falls by a small amount (see also the 99/1-ratio in Table III). This is, however, solely an artifact of the decline in the share at the extreme right-tail of the distribution; inequality among the rest of the distribution is increased. This provides one example of the pitfalls that exists when summarizing distributions by their Gini coefficient.

Compared to the exogenous-information economy, the widening-out of the wealth distribution is dampened in our benchmark economy. Households' optimal information choices lessen the widening of the wealth distribution (see also the decomposition in Figure 8). Finally, for both the benchmark economy and its exogenous-information counterpart, incomplete information causes inequality to become more cyclical. The correlation between the Gini coefficient (as well as other measures of dispersion) and output increases by around 18 percent—busts (booms) become periods with relatively more (less) inequality.

Decomposing the drivers of the change in the wealth distribution is challenging, as the distribution is the equilibrium outcome of a model in which the dynamics also change when we modify the informational assumptions. To make progress, we break-down three separate

Table III: Wealth Distribution

	Panel a: Level of Moments					
	Mean K	Gini G	90/10	99/1	$Cor(K,Y)$	$Cor(G,Y)$
Benchmark Model	39.34	0.36	5.51	24.74	0.70	-0.35
Exogenous Information	39.34	0.41	5.80	32.33	0.65	-0.35
Full Information	39.38	0.40	5.23	30.89	0.57	-0.29
	Panel b: Percent Difference to Full Information					
	Mean K	Gini G	90/10	99/1	$Cor(K,Y)$	$Cor(G,Y)$
Benchmark Model	-0.10	-9.26	5.32	-19.92	22.34	17.98
Exogenous Information	-0.10	1.17	10.83	4.68	14.09	18.42

Note: The table shows the mean of the logarithm of capital (K), the Gini coefficient of the capital distribution (G), as well as the 90/10 and 99/1 percentile ratios of the wealth distribution. In addition, the table shows that correlation between the logarithm of capital, the Gini coefficient, and output (Y) (e.g., $Cor(G,Y)$). The table computes these moments in the calibrated model ("Benchmark Model"), the associated full-information economy ("Full Information"), as well as in the model with an exogenously specified probability of acquiring information ("Exogenous Information"). The probabilities of information acquisition in the benchmark model are 0.1437 and 0.1295 for the employed and unemployed, respectively. This probability is set equal to 0.1306 in the exogenous information case for all households and for all moments in time.

forces that shape the wealth distribution: (i) the change in the equilibrium law-of-motion for capital; (ii) the presence of incomplete information; and (iii) the heterogeneity in the extent of incomplete information. The first captures the change in the wealth distribution that occurs due to the different general equilibrium behavior of economy-wide aggregates (as there are significantly larger business-cycle fluctuations). The second measures the change that arises due to the incompleteness of information itself. The last measures the added effect from the endogenous heterogeneity of information. Combined, these forces capture the partial and general equilibrium effects by which heterogenous, incomplete information alters the wealth distribution. Figure 8 provides a breakdown of the overall change into these three components.

General-equilibrium Effects:

To isolate the general-equilibrium effects, we conduct the following experiment: We solve for household policy functions in the full-information economy taking the *benchmark* economy's law of motion for capital, $\tilde{H}(z,K)$, as given. We then simulate the economy with the same sequence of shocks as in the full-information case and compare the two wealth distributions.

Figure 8: Decomposition of Changes to the Wealth Distribution

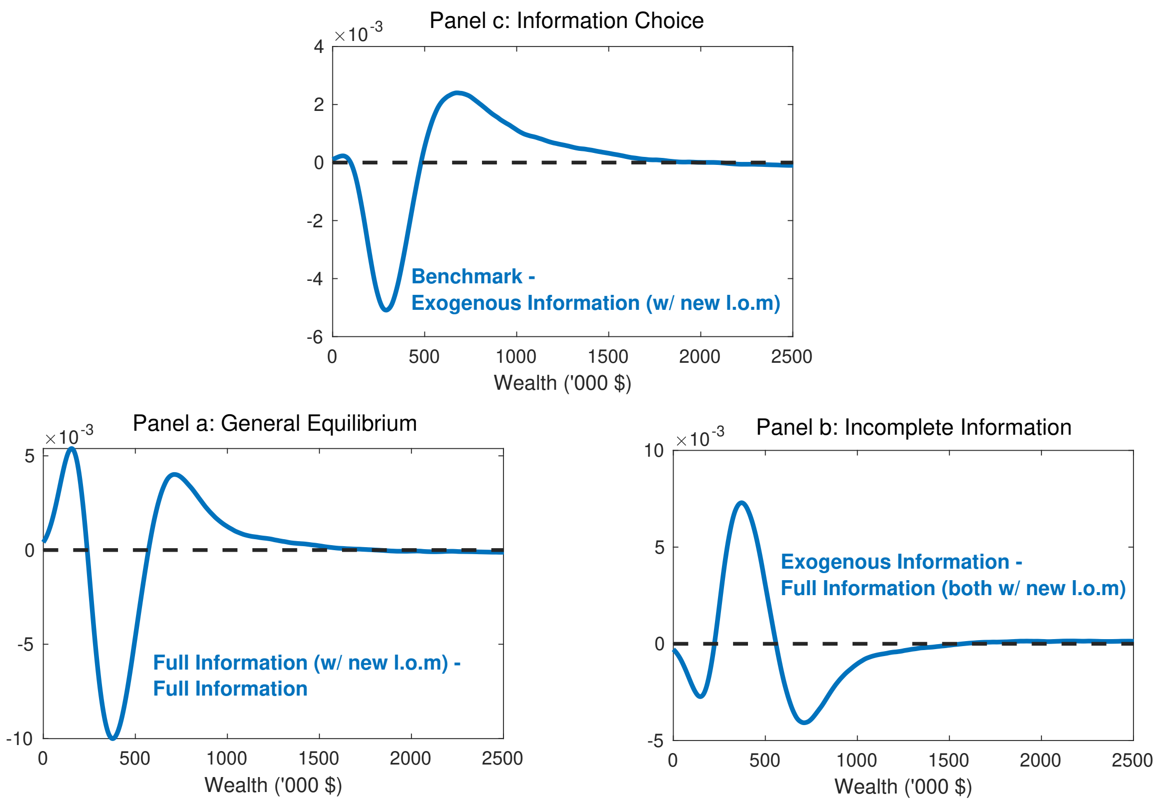

Note: The figure decomposes the changes in the average wealth distribution relative to the full-information version of the benchmark economy. Panel a shows the changes in the average probability density function of the wealth distribution between the benchmark economy and an exogenous incomplete-information economy, where the probability of information acquisition is fixed at the average value in the benchmark economy. We equip the latter economy with a law-of-motion for the capital stock equal to that in our benchmark economy. Panel b depicts the changes between the full-information economy and the full-information economy in which the law of motion for the capital stock equals that in our benchmark economy. Finally, Panel c shows the differences between the economy with exogenous incomplete-information and the full-information economy, where the law-of-motion in both cases equals that from our baseline economy. The horizontal axis in all panels is household wealth (capital levels) in ’000s \$. We use 2020 values of US household income to convert values in the model to \$ amounts. Probability density functions are estimated from a simulated panel of household wealth-levels, using a kernel density estimator with the Epanechnikov kernel.

This experiment isolates the channel by which the weakening of the mean-reversion of capital affects the wealth distribution.[16] Panel a in Figure 8 shows that the change in the *equilibrium law-of-motion for capital is a* large contributor to the overall change in the wealth distribution.

All else equal, the more volatile, persistent capital stock causes a widening of the difference between the income of employed household and unemployed households.[17] The cross-sectional dispersion of labor-market income increases with the volatility of wages, and hence in the volatility of the aggregate capital stock. This increase in the dispersion of labor-market income, all else equal, causes the wealth distribution to widen. Additionally, savings out of labor market income also become more cyclical, due to the increase in the persistence and volatility of wages. This further increases inequality, on average, for most of the distribution.

That said, an offsetting effect, which can be seen from the behavior of the extreme right-tail in Panel a, arises from the increased persistence of the rate of return on capital. This follows from the increased persistence of the capital stock. The right-tail features households for which labor income comprises a small share of total resources while capital income comprises the bulk. All else equal, the increased persistence of the return on capital increases the strength of the income effect and makes households less willing to exploit expected changes in returns. This, in turn, makes savings and returns less positively correlated, which dampens high-wealth households' asset accumulation and decreases the right–most tail of the wealth distribution.

Incomplete Information:

The next experiment isolates the effects that the presence of incomplete information itself has on the average wealth distribution. We now solve for the policy function and simulate the distribution when households have exogenously incomplete information, but believe that the law-of-motion for capital is the one from the benchmark economy. By comparing the distribution under this experiment to the previous one (full-information with benchmark law-of-motion), we quantify the effect of household incomplete information. Panel b in Figure 8 plots the difference between the two distributions.

The presence of incomplete information causes households at the bottom of the wealth distribution to increase precautionary savings. These are the households that are sensitive to the additional risk. As a result, the mass at the lower-end of distribution falls and the mass of lower-middle-class households rises. The mass of household near the top-end of distribution, moreover, declines somewhat, as households near the top-end of distribution make erroneous savings decisions; they are unable to effectively exploit periods of high returns on capital.

[16] To be clear, households still optimize in this experiment and markets clear in every period. Household expectations are simply inconsistent with the aggregate dynamics; they are "surprised" every period.

[17] In equilibrium, the difference between wages and unemployment benefits equal $(1-\mu)w_t = (1-\mu)(1-\alpha)z_t K_t^{\alpha}\bar{l}^{-\alpha}$, which is increasing in the volatility of the capital stock K_t.

At the extreme right-tail of the distribution in Panel b, however, this effect is counteracted by the increased randomness by which households make savings choices. Since even extreme-wealth households only infrequently update information, their information about the expected rate of return is stochastic. Furthermore, these households have roughly linear policy functions, as they are far away from the borrowing constraint. Combined, this makes extreme-wealth households behavior analogous to that described in Piketty and Saez (2003), which show how exogenous random savings rates and linear policy functions can generate Pareto tails in the wealth distribution. Here, however, incomplete information provides a microfoundation for this type of "random savings behavior", as opposed to other models that either assume exogenously stochastic savings rates or random returns on savings.

Information Choice:

Our final decomposition compares the distribution from the previous experiment (exogenous, incomplete information with benchmark law-of-motion) with the benchmark distribution. This difference isolates the additional effects that *heterogeneous information choices* have on the wealth distribution. We plot the difference in Panel c in Figure 8.

All else equal, the more informed households in our baseline economy are the rich households with substantial amounts of wealth (Figure 6). These households, who are better informed than their exogenous information counterpart, make fewer savings mistakes—they are, on average, better able to exploit differences in rates of return on capital. As a result, the right-end of the distribution increases somewhat in size. That said, the extreme right-tail of the distribution—those with more than \$2.5 million in wealth– conversely moves in somewhat, due to the decline in "random savings" caused by the additional information.

Lastly, notice the small increase in the mass of wealth-poor households, those that hold between \$0-50,000. These are the households that are on the upward-sloping part of Figure 6. These households have slightly more accurate information than their exogenous information counterpart under the baseline calibration. Hence, these households engage in less precautionary savings, which, all else equal, increases the mass of wealth-poor households.

Summary of Effects:

In sum, the presence of heterogeneous, incomplete information leads to rich and complex changes in the wealth distribution. On the one hand, the presence of incomplete information widens the wealth distribution by increasing the volatility and persistence of fluctuations, and by leading to "random savings behavior". The former increases the expected difference between employed and unemployed households and leads to more cyclical savings. The existence of heterogeneity in information further exacerbates inequality by allowing informed, wealthy households to make better savings decisions. On the other hand, incomplete information also increases pre-cautionary savings, all else equal, leading to a tighter wealth distribution. On balance, we find that former effects dominate the latter. For most of the distribution, heterogeneous, incomplete information causes inequality to increase. Yet, the combined effects are subtle and affect different parts of the distribution differentially.

The results in this section point to a broader conclusion. Because the presence of incomplete information alters households' economic choices, heterogeneity in information acquisition among households naturally contributes to economic differences between households. Such differences, in turn, feed back and spill-over onto the aggregate economy, which itself affects economic inequality. The next section shows how the interplay between information choice, pre-cautionary savings, and the aggregate economy also modifies the predictions of a simple economic policy that directly targets the expectation-wealth nexus.

6 A Policy Experiment

The foregoing section highlighted the interaction between heterogenous information choice, the aggregate economy, and precautionary savings. In this section, we illustrate how policy reforms that affect households' savings decisions modify information acquisition choices, potentially changing the properties and dynamics of the economy in unexpected ways.

For our policy counterfactual, we consider a wealth tax. Such a tax has been hotly debated by policymakers and academics in recent years and introduced as a policy proposal in the U.S. Congress in 2021 (e.g., Guvenen *et al.*, 2019; Saez and Zucman, 2022).[18] One of the main arguments of the proponents of the tax is that it will reduce inequality and be an equitable 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.[19] Household cash-at-hand m_i is therefore given by the expression:

$$m_i = rk_i + (1-\tau)\,\epsilon_i w\bar{l} + \mu\,(1-\epsilon)\,w + (1-\delta-\tau_k)k_i. \tag{6.1}$$

We model the size of the wealth tax based on the wealth tax that was in effect in France, which had an wealth tax in place from 2011 to 2017. The magnitude of the tax is also consistent with the size of the recent proposal in the U.S. Congress. In particular, we set $\tau_k = 0.0025$,

[18] The "Warren 2021 proposal" can be found here: https://www.congress.gov/bill/senate-bill/510

[19] For 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 government consumption G.

Table IV: Quantitative Effects of a Wealth Tax

	Percent Difference (%)						
	Mean K	$\sigma(Y)$	Gini	90/10	99/1	Info. emp.	Info. unemp.
Benchmark Model	-10.22	4.38	2.68	-2.29	16.51	-24.46	-30.27
Full Information	-10.24	1.11	-5.13	-9.67	-8.13	.	.

Note: The table shows the effects of a one percent per annum wealth tax on moments of the average cross-sectional capital distribution and the logarithm of economy-wide output. The table computes the moments for the both calibrated model ("Benchmark Model") and the associated full-information economy ("Full Information"). The final two columns measure the percent difference in the average number of households who acquire information. We compute this probability separately for the employed and unemployed.

corresponding to a one percent per annum wealth tax. Table IV and Figure 9 report the macroeconomic effects of the wealth tax.[20]

The direct effect of the wealth tax is, unsurprisingly, to reduces aggregate savings as seen by the 10 percent drop in the capital stock reported in Table IV. The wealth tax reduces a household's incentive to save—for a given level of wealth and information—relative to the no-tax benchmark. Information acquisition policies, on the other hand, for *a given* wealth level are approximately unaffected by the tax. However, because the probability of information acquisition is heterogeneous across the wealth distribution (Figure 5), the wealth tax changes the average level of information in the economy. The introduction of a wealth tax moves the mean of the distribution of individual capital-holdings to the left, resulting in a shift in the distribution of household expectations. Since lower-wealth households acquire information less frequently, the fraction of households that acquire information every period falls.

The consequences of this change in household information choices matter quantitatively: 24-30 percent fewer households acquire information every period after the introduction of the tax. The drop for the unemployed is somewhat more pronounced, as their information acquisition probability has a noticeable hump shape at low asset levels (Figure 5). Households within this region are therefore more sensitive to the reduction in wealth that is caused by the tax. In turn, this reduction in information acquisition probabilities decreases the accuracy of household expectations, on average. The increased dispersion in expectations dampens the mean-reversion of capital, leading to more volatility in the capital stock. As a result, the volatility of output also increases by about 5 percent after the introduction of the wealth tax.

In addition to changing the aggregate dynamics of the economy, the introduction of the

[20] The French wealth tax was called the "Impôt de solidarité sur la fortune" (ISF). The IFS was an annual tax, with rates from 0.5 percent to 1.0 percent per annum, depending on your wealth. We abstract below from the progressivity of the wealth tax.

wealth tax has complex distributional consequences. We plot the change in the wealth distribution after the introduction of the tax as the blue line in Figure 9. The number of super-rich households (with greater than \$3.5 million in wealth) actually *increases* with the introduction of the wealth tax. The economy experiences a fall in the share of merely rich households—those between \$600,000 and \$3.5 million in wealth— and a rise in households with wealth below \$600,000. Although the 90/10 percentile ratio of the wealth distribution falls, the Gini coefficient and the 99/1 percentile ratio increase (Table IV). The increase in the Gini coefficient is explained by its sensitivity to the extreme right-tail of the wealth distribution (Cowell and Flachaire, 2002). Thus, counterintuitively, introducing a wealth tax actually increases "wealth inequality" in our benchmark model, as measured by the Gini coefficient.

There are three contrasting forces that explain our results. First, the direct effects of the wealth tax disproportionally impacts high-wealth households—those above around \$500,000 in wealth of which we have more under heterogeneous, incomplete information (Figure 9)—lowering their wealth share. Second, poorer households are, on average, less informed, and hence make worse savings decisions. Combined, the first two forces lead to a powerful narrowing of the wealth distribution in our benchmark economy.

Finally, the decrease in information acquisition probabilities make wealthy households' savings more stochastic. Their knowledge about the rate of return on capital becomes more random. That, combined with approximately linear policy rules (as wealthy households are far from the borrowing constraint), generates approximately "random savings rates". Similar to the results in Piketty and Saez (2003), and for the same reasons as discussed in Section 5.4, this generate a wider tail in the wealth distribution. The mass of super wealthy households with more than \$3.5 million in wealth, as a consequence, increases, despite the wealth tax.

To illustrate the importance of the wealth-expectations nexus for understanding policy counterfactuals, we perform the same experiment in the full-information version of our economy. As shown in Table IV, the direct impact of the tax on average wealth is nearly identical—both capital stocks fall by the same amount. But, that is where the similarities end. The volatility of output only increases marginally, as compared to the benchmark case. The dampened mean-reversion of capital under incomplete information is absent in the full-information case, leading to little change in endogenous volatility through capital accumulation.

Turning to the distributional implications of the tax, the difference in the effect on inequality between the two cases is even more pronounced. Under full-information, standard measures of wealth inequality—the Gini coefficient, the 90/10 percentile ratio, and the 99/1 percentile ratio—all fall. Inequality is uniformly decreased, as also shown by the red line in Figure 9. The mass of households with less than \$450,000 increases, while the mass of households with greater wealth decreases. Thus, in the full-information economy the wealth tax

Figure 9: Wealth Taxes and Changes in the Wealth Distribution

Panel a: Low-to-high Range

Panel b: High-end Range

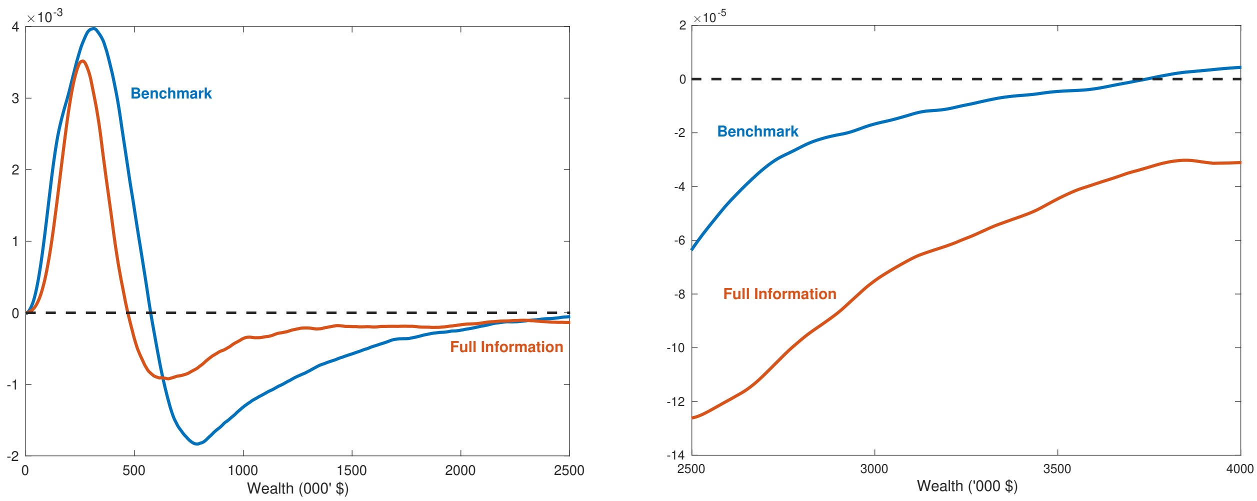

Note: The figure illustrates changes in the average wealth distribution relative to the zero-wealth tax case. We illustrate these changes for both our calibrated model ("Benchmark Model") and the associated full-information economy ("Full Information"). We use 2020 values of US household income to convert values of capital-holdings in the model to \$ amounts. Probability density functions are estimated from a simulated panel of households, using a kernel density estimator with the Epanechnikov kernel.

uniformly reduces wealth inequality. This is in contrast to the benchmark case, where, although inequality is reduced more for middle-wealth households, it is increased relative to the top-end of the wealth distribution. This provides one example of how the general equilibrium forces that arise from heterogeneous information choices can counteract, and in this case even sometimes dominate, the direct effects of a policy reform.

While we abstain from making welfare statements about the desirability of the wealth tax that we analyze, our positive findings indicate that policymakers should proceed with caution when evaluating the consequences of tax policy. The effects on households' information choices may lead to implications which run counter to the stated objectives of the policy—in this case, lead to an increase in the mass of extremely wealthy households and a more volatile economy.

More generally, the above policy experiment illustrates that macroeconomic policies may have important additional effects in environments with heterogenous, endogenous information. By changing the distribution of agents' information, and hence their expectations, macroeconomic policies fundamentally alter an economy's responsiveness to shocks, as well as individual agents' decision rules. These additional effects may be quantitatively important—both from a positive and a normative perspective.

7 Conclusion

The frontier of macroeconomic research continues to incorporate salient dimensions of household and firm heterogeneity to provide a more complete and accurate description of the macroeconomy. In this paper, we illustrated how the interaction between two important dimensions of household heterogeneity—heterogeneity in expectations and heterogeneity in wealth—gives rise to new qualitative and quantitative insights about macroeconomic dynamics and the effects of macroeconomic policies. In particular, we demonstrated how the wealth-expectation nexus increases the endogenous propagations of shocks and partially accounts for the lack of wealth inequality in standard frameworks with incomplete markets. We showed how the wealth-expectation nexus further fundamentally alters the predictions of the consequences of a wealth tax—and in unexpected ways.

Our findings have important implications for both the heterogenous-agent macro literature and the literature on models with dispersed information. For the former, our policy experiment provides a "Lucas-style" criticism (Lucas Jr, 1976) to policy analysis in incomplete-markets models: Any policy that has a substantial impact on the wealth distribution will systematically affect household information choices and their expectations, with associated implications for macroeconomics dynamics and the cross-section.[21] For the latter, studying the consequence of dispersed information in models with linear policy rules misses the important two-way interaction between the distribution of agent wealth and the non-linearity of the value of additional information. 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 choices have 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 or alter the manner in which they condition their policy instruments? And how should such subsidies or policy instruments target a particular subset of the population? We leave these exciting questions for future research.

[21] In this sense, our results provide a Lucas-style criticism (Lucas Jr, 1976) of Lucas' own comments about the response of an economy with incomplete information to shocks; that "It seems safe and, for my purposes, sensible to abstract here from the fact that in reality this situation can be slightly mitigated by the purchase of additional information" (p. 1121, Lucas Jr, 1975).

References

Andrade, P., Coibion, O., Gautier, E. and Gorodnichenko, Y. (2022). No firm is an island? how industry conditions shape firms expectations. *Journal of Monetary Economics*, **125**, 40–56.

Armantier, O., Topa, G., Van der Klaauw, W. and Zafar, B. (2017). An overview of the survey of consumer expectations. *Economic Policy Review*, (23-2), 51–72.

Auclert, A., Rognlie, M. and Straub, L. (2020). Micro jumps, macro humps: Monetary policy and business cycles in an estimated hank model. *National Bureau of Economic Research*.

Bhandari, A., Borovička, J. and Ho, P. (2021). Survey data and subjective beliefs in business cycle models. *Available at SSRN 2763942*.

Bordalo, P., Coffman, K., Gennaioli, N. and Shleifer, A. (2016). Stereotypes. *The Quarterly Journal of Economics*, **131** (4), 1753–1794.

—, Gennaioli, N. and Shleifer, A. (2017). Memory, attention, and choice. *The Quarterly journal of economics*.

Broer, T., Kohlhas, A., Mitman, K. and Schlafmann, K. (2022). On the possibility of krusell-smith equilibria. *Journal of Economic Dynamics and Control*.

Carroll, C., Slacalek, J., Tokuoka, K. and White, M. N. (2017). The distribution of wealth and the marginal propensity to consume. *Quantitative Economics*, **8** (3), 977–1020.

Carroll, C. D. (1997). Buffer-stock saving and the life cycle/permanent income hypothesis. *The Quarterly journal of economics*, **112** (1), 1–55.

— (2003). Macroeconomic expectations of households and professional forecasters. *the Quarterly Journal of economics*, **118** (1), 269–298.

—, Crawley, E., Slacalek, J., Tokuoka, K. and White, M. N. (2020). Sticky expectations and consumption dynamics. *American economic journal: macroeconomics*, **12** (3), 40–76.

Christiano, L. J., Eichenbaum, M. and Evans, C. L. (2005). Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of political Economy*, **113** (1), 1–45.

—, Eichenbaum, M. S. and Trabandt, M. (2016). Unemployment and business cycles. *Econometrica*, **84** (4), 1523–1569.

—, Trabandt, M. and Walentin, K. (2010). Dsge models for monetary policy analysis. **3**, 285–367.

Coibion, O., Gorodnichenko, Y. and Kamdar, R. (2018). The formation of expectations, inflation, and the phillips curve. *Journal of Economic Literature*, **56** (4), 1447–91.

—, — and Ropele, T. (2020). Inflation expectations and firm decisions: New causal evidence. *The Quarterly Journal of Economics*, **135** (1), 165–219.

Cowell, F. and Flachaire, E. (2002). Sensitivity of inequality measures to extreme values. *LSE STICERD Research Paper*, (60).

Del Negro, M., Schorfheide, F., Smets, F. and Wouters, R. (2007). On the fit of new keynesian models. *Journal of Business & Economic Statistics*, **25** (2), 123–143.

Den Haan, W. J., Judd, K. L. and Juillard, M. (2010). Computational suite of models with heterogeneous agents: Incomplete markets and aggregate uncertainty. *Journal of Economic Dynamics and Control*, **34** (1), 1 – 3, computational Suite of Models with Heterogeneous Agents: Incomplete Markets and Aggregate Uncertainty.

Faust, J. and Wright, J. H. (2013). Forecasting inflation. *Handbook of economic forecasting*, **2**, 2–56.

Gabaix, X. (2019). Behavioral inattention. In *Handbook of Behavioral Economics: Applications and Foundations 1*, vol. 2, Elsevier, pp. 261–343.

Grossman, S. J. and Stiglitz, J. E. (1980). On the impossibility of informationally efficient markets. *The American economic review*, **70** (3), 393–408.

Guvenen, F., Kambourov, G., Kuruscu, B., Ocampo-Diaz, S. and Chen, D. (2019). Use it or lose it: Efficiency gains from wealth taxation.

Hellwig, C. and Veldkamp, L. (2009). Knowing what others know: Coordination motives in information acquisition. *The Review of Economic Studies*, **76** (1), 223–251.

Jakobsen, K., Jakobsen, K., Kleven, H. and Zucman, G. (2020). Wealth taxation and wealth accumulation: Theory and evidence from denmark. *The Quarterly Journal of Economics*, **135** (1), 329–388.

Kaplan, G. and Violante, G. L. (2018). Microeconomic heterogeneity and macroeconomic shocks. *Journal of Economic Perspectives*, **32** (3), 167–94.

Krueger, D., Mitman, K. and Perri, F. (2016). Macroeconomics and household heterogeneity. In *Handbook of Macroeconomics*, vol. 2, Elsevier, pp. 843–921.

Krusell, P. and Smith, A. (1998). Income and Wealth Heterogeneity in the Macroeconomy. *Journal of Political Economy*, **106** (5).

Lucas Jr, R. E. (1975). An equilibrium model of the business cycle. *Journal of political economy*, **83** (6), 1113–1144.

— (1976). Econometric policy evaluation: a critique, in k brunner and a meltzer (eds.), the phillips curve and labour markets, carnegie-rochester conference series on public policy, vol. 1,.

Macaulay, A. and Moberly, J. (2022). Heterogeneity in imperfect inflation expectations: theory and evidence from a novel survey. *Oxford University mimeo.*

Maćkowiak, B., Matějka, F. and Wiederholt, M. (2018). Dynamic rational inattention: Analytical results. *Journal of Economic Theory*, **176**, 650–692.

—, — and Wiederholt, M. (2021). Rational inattention: A review. *Journal of Economic Perspectives*.

Mankiw, N. G. and Reis, R. (2002). Sticky information versus sticky prices: a proposal to replace the new keynesian phillips curve. *The Quarterly Journal of Economics*, **117** (4), 1295–1328.

—, — and Wolfers, J. (2003). Disagreement about inflation expectations. *NBER macroeconomics annual*, **18**, 209–248.

McFadden, D. *et al.* (1973). Conditional logit analysis of qualitative choice behavior. *Institute of Urban and Regional Development, University of California Berkeley*.

Muth, J. F. (1961). Rational expectations and the theory of price movements. *Econometrica: Journal of the Econometric Society*, pp. 315–335.

Piketty, T. and Saez, E. (2003). Income inequality in the united states, 1913–1998. *The Quarterly journal of economics*, **118** (1), 1–41.

Reis, R. *et al.* (2020). The people versus the markets: A parsimonious model of inflation expectations.

Saez, E. and Zucman, G. (2022). Wealth taxation: Lessons from history and recent developments. **112**, 58–62.

Sims, C. A. (2003). Implications of rational inattention. *Journal of monetary Economics*, **50** (3), 665–690.

Stark, T. *et al.* (2010). Realistic evaluation of real-time forecasts in the survey of professional forecasters. *Federal Reserve Bank of Philadelphia Research Rap, Special Report*, **1**.

Townsend, R. M. (1983). Forecasting the forecasts of others. *Journal of Political Economy*, **91** (4), 546–588.

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.

Weber, M., DAcunto, F., Gorodnichenko, Y. and Coibion, O. (2022). The subjective inflation expectations of households and firms: Measurement, determinants, and implications. *National Bureau of Economic Research*.

A Motivating Evidence

A.1 Additional Estimates

Table A.5: Unemployment Expectations Across the Wealth Distribution

	Absolute Error		
	(1)	(2)	(3)
Wealth Share (0-10 percent)	-0.009	0.063***	0.060***
	(0.020)	(0.017)	(0.017)
Wealth Share (10-20 percent)	0.045**	0.088***	0.082***
	(0.020)	(0.018)	(0.017)
Wealth Share (20-40 percent)	0.064***	0.089***	0.084***
	(0.016)	(0.013)	(0.013)
Wealth Share (40-60 percent)	0.050***	0.029**	0.026**
	(0.015)	(0.012)	(0.012)
Wealth Share (60-80 percent)	0.003	0.023*	0.020
	(0.015)	(0.012)	(0.012)
Wealth Share (80-100 percent)	–	–	–
Controls	–	✓	✓
Time Fixed Effects	–	✓	✓
Pre-2020Q1	–	–	✓
Observations	40,998	37,163	36,408
F Statistic	6.12	409.57	355.05
R^2	0.01	0.44	0.40

Note: Column (1) shows estimates from a regression of the absolute value of individual unemployment errors on the wealth bucket (decile/quintile) that the individual respondent belongs to. Estimates are relative to the wealthiest households, those in the 80-100 percentile of the wealth distribution. Column (2) adds controls to the regression specification: the age, education level, labor market status, and sex of the respondent, as well as time fixed effects. Column (3) considers estimates "pre-covid"; that is, only before January 2020. Robust standard errors in parentheses. Sample: 2013M10-2020M3. $*$ p<.1, $**$ p<.05, $***$ p<.01

Figure A.10: Unemployment Expectations Across the Wealth Distribution (VAR)

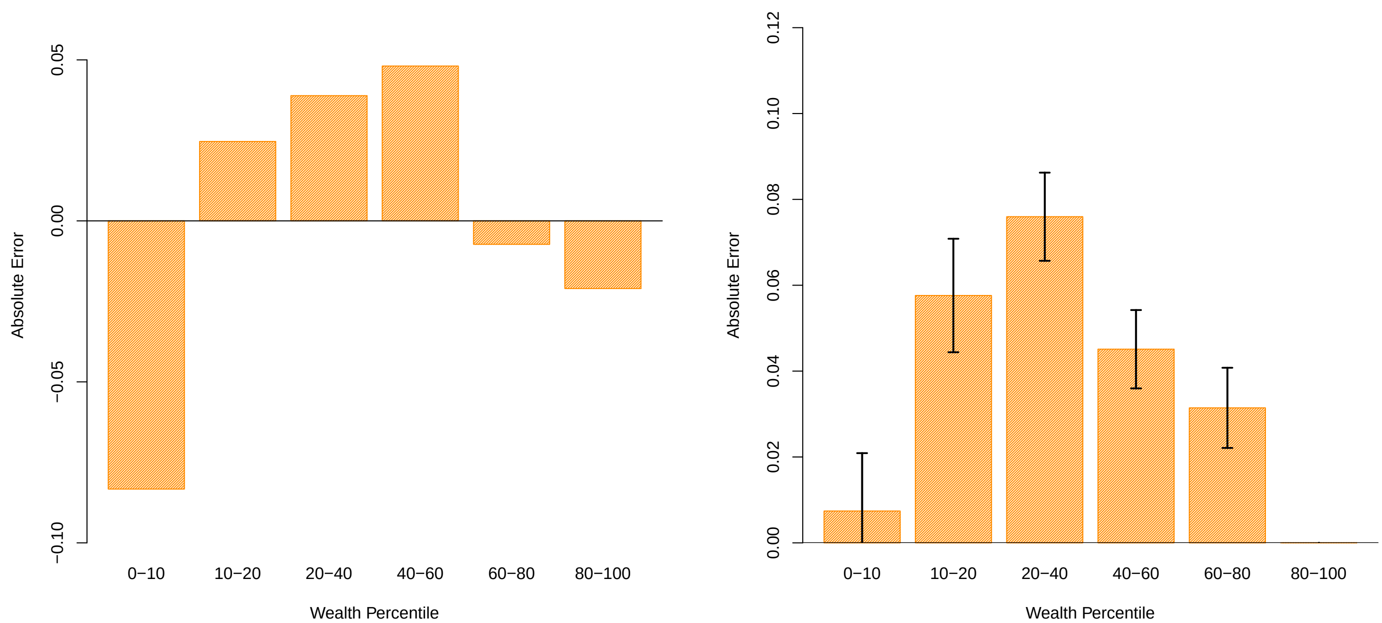

Panel a: Relative Accuracy

Panel b: Coef. Estimates on Wealth

Note: Panel a plots the difference between the average one-year ahead accuracy of unemployment forecasts within wealth deciles/quintiles and the overall average taken across all wealth levels. Accuracy is measured by the absolute value of unemployment errors. We proxy the true probability of rising unemployment with that from a standard Bayesian VAR (Online Appendix A.3). Panel b plots the coefficient estimates on wealth from a regression of the absolute value of individual errors on wealth deciles/quintiles, controlling for the age, education level, labor market status, and sex of the respondent, as well as time fixed effects. Estimates are relative to the wealthiest households, those in the 80-100 percentile of the wealth distribution. Whisker-intervals correspond to one-standard deviation robust confidence bounds. Sample: 2013M10-2020M1.

A.2 Data Construction

The SCE is a monthly internet survey of c. 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. Importantly, a yearly module also asks the survey respondents for key financial variables, including their financial wealth.

A.2.1 Variable Definitions

We focus on expectations of three variables: inflation, house prices, and the unemployment rate. The former two ask respondents for their best guess of a variable's outcome, in addition to the probability of it falling into a number of bins. The exact questions are:

- Inflation:
  "*What do you expect the rate of (CPI) 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 prices:
  "*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 estimates 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 $P_{it}(u_{t+12} > u_t)$, we would ideally compare household $i \in [0, 1]$'s response to the true-but-unobserved probability $P_t(u_{t+12} > u_t)$. Consistent with ample evidence that professional forecasters provide more accurate predictions than those from modern statistical and economic models (Stark *et al.*, 2010; Faust and Wright, 2013; and Bhandari *et al.*, 2021), we proxy the true probability by the consensus forecast from the SPF, which we denote $P_{SPF,t}(u_{t+12} > u_t)$. 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 also multiply our measure by 2 to make it consistent with the "Brier score". We thus compute the errors in unemployment forecasts as

$$\nu_{it} = 2 \times \frac{P_{it}(u_{t+12} > u_t) - P_{SPF,t}(u_{t+12} > u_t)}{T^{-1}\sum_t P_{SPF,t}(u_{t+12} > u_t)}, \tag{A1}$$

where the average is computed across all observations in our sample.

In addition to survey estimates, we use the following household characteristics: sex, age, dummies that take values of one 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.[22] We construct wealth deciles/quintiles based on the initial two-years of data (2013 and 2014). We deflate the resulting quantities by the level of the US consumer price index.

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

A.2.2 Summary Statistics

Table A.6 illustrates that households' 12-month unemployment and inflation expectations from the SCE are on average less accurate than professional forecasts. Households attach on average a higher probability to rising unemployment than professional forecasters, implying larger forecast errors during a sample period where unemployment 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 A.6 demonstrates that household expectations are

[22] The 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*?". The question about mortgage debt is *"Approximately, what is the total amount of outstanding loans against your home(s), including all mortgages and home equity loans?"*, while that for total debt is *"Approximately, what is the total amount of your [and your spouses/partners] current outstanding debt?"*.

Table A.6: Macroeconomic Expectations in the SCE and SPF

	Panel a: Unemployment Rate	
	Median Forecast	Std. Dev. of Forecast
SCE	39.00	22.98
SPF	32.13	17.80

	Panel b: Inflation			
	Median Abs. Error	Std. Dev. of Error	Median IQR	Std. Dev. of IQR
SCE	1.61	2.93	2.00	4.48
SPF	0.72	0.65	0.56	0.25

Note: The table shows moments of the individual probability distributions from the Survey of Consumer Expectations (SCE) and the Survey of Professional Forecasters (SPF). Panel a shows the median and standard deviation of individual unemployment forecasts. Panel b shows the median error of individual inflation forecasts (column 2), the standard deviation of these errors (column 3), the median interquartile ranges derived from individual distributions (column 4), and their standard deviation (column 5).

substantially more uncertain than professional forecasts. 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 A.6 also shows that household expectations are substantially more *heterogeneous* than SPF forecasts. Specifically, household unemployment expectations and point forecasts for CPI inflation have a substantially higher cross-sectional 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 professional forecasters.

A.3 Forecasting VAR

We use a standard quarterly forecasting VAR to compute the forecasts of the probability of a rising unemployment rate under the data-generating measure. All time series are downloaded from FRED for the period 1960Q1– 2019Q4: CPI inflation (CPIAUCSL, percentage change from a year ago), real GDP (GDPC1, percentage from a year ago), unemployment rate (UNRATE), log hours worked per capita (average hours per worker PRS85006023 multiplied by the employment-population ratio CE16OV/CNP16OV), and the federal funds rate (FEDFUNDS). The VAR is estimated with two lags and we use a AR(1)-Minnestota prior for all variables. These choices for the VAR are similar to those made in Christiano *et al.* (2005), Del Negro *et al.* (2007), Christiano *et al.* (2010), or Christiano *et al.* (2016). We sample

100,000 observations at each moment in time from the posterior distribution, to estimate the probability of a rising unemployment rate. We experimented with increasing the number of lags used and including additional forecasting variables (e.g., consumption of non-durables, wages, and capacity utilization), without materially affecting our results. In all cases, the characteristics of the documented heterogeneity in unemployment expectations is similar to that in Section 2 and Figure A.10: an inverse u-shape in accuracy across wealth.

B Calibration and Model Fit

B.1 Calibration Parameters

Table B.1: Parameterization

Parameter	Value
Externally calibrated parameters	
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
Internally calibrated parameters	
Discount factor (β)	0.99
Relative risk aversion (γ)	5.00
Monetary cost of information (ν)	0.0012
Scale parameter of utility cost of information (α^{κ})	$1/3e^{-8}$

B.2 Time-series for Capital and Priors

Incomplete information makes individual prior expectations about the current capital stock move more slowly than the actual capital stock. In particular, households who choose not to acquire information will have priors ($\hat{K}$) about 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

Figure B.1: Mean Capital K_t: Realization and Priors

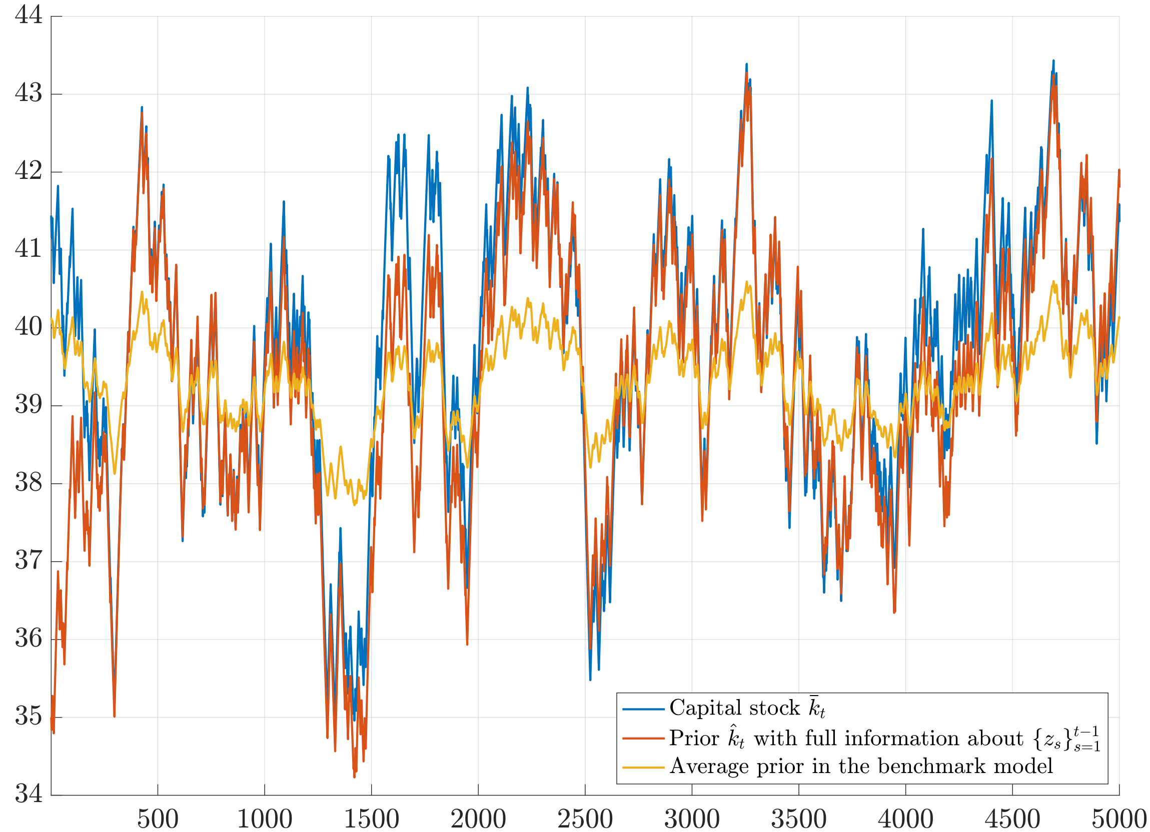

Note: Based on a simulation of the calibrated benchmark model, the figure shows time series of the mean (aggregate) capital stock $K_t = \bar{k}_t$ (blue line), the prior about current aggregate capital $\hat{K}_t = \hat{k}_t$ of households who acquire information about the current productivity state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

about productivity. In fact, for economies with full information, Den Haan *et al.* (2010) show that the history of shocks z^t alone allows for very accurate predictions about the future capital stock K_{t+h}, $h \geq 1$. 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 has an arbitrary belief about capital in period 0 but then 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).[23]

[23]In 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