Expectation and Wealth Heterogeneity in the Macroeconomy*

Tobias Broer Alexandre Kohlhas Kurt Mitman Kathrin Schlafmann

September 2025

Abstract

We document systematic differences in macroeconomic expectations across US 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. In comparison to a full-information counterpart, our model suggests substantially increased macroeconomic volatility and inequality. We show through a series of examples that neglecting the information channel can lead to erroneous conclusions about the effects of macroeconomic policies. While in the model without information choice, a wealth tax, for example, reduces wealth inequality, in our framework it reduces information acquired, leading to increased volatility and higher top-end inequality in equilibrium.

JEL codes: D84, E21, E27, E63

Keywords: Expectations, heterogeneity, information

*Broer: PSE (tobias.broer@psemail.eu). Kohlhas: University of Oxford (alexandre.kohlhas@economics.ox.ac.uk). Mitman: IIES, Stockholm University, CEMFI (kurt.mitman@iies.su.se). Schlafmann: Copenhagen Business School (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 various locations. A previous version of this paper was circulated under the title "Heterogeneous Information Choice in General Equilibrium". This research received financial support from Ragnar Soderberg Stiftelsen and the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant agreement No. 759482) . Broer and Kohlhas thank the Institut Universitaire de France and the John Felt Fund for financial support, respectively. Mitman gratefully acknowledges the support of the Riksbankens Jubileumsfond, the Ministerio de Ciencia, Innovación y Universidades Grant IneqEx, and the Swedish Research Council Project Grant 2023-01635. Schlafmann acknowledges financial support from the Danish Finance Institute (DFI).

1 Introduction

Expectations have been part of 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—posits that all households, at all moments in time, have the same expectations about the macroeconomy. Building 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 characteristics. This creates systematic heterogeneity in both the level and accuracy of expectations across the distribution of households (e.g., Carroll, 2003; Lusardi and Mitchell, 2014; Coibion *et al.*, 2018; Weber *et al.*, 2022). Given the importance of expectations to macroeconomics, it seems 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 heterogeneous-agent model with aggregate risk. Our main contribution is to provide the first, to our knowledge, heterogenous-agent model that allows for the study of the macroeconomic consequences of systematic differences in expectations. Our quantitative-theoretical 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 (e.g., Krueger *et al.*, 2016; Kaplan and Violante, 2018). We use our framework—disciplined by survey data—to assess the consequences of the expectation-wealth nexus for understanding aggregate fluctuations and the distribution of wealth. We then explore how the presence of heterogeneous expectations can also modify the efficacy of macroeconomic policies.

We first provide new evidence on the heterogeneity in household expectations using US micro-level data. We show that in a leading household survey (the New York Fed's Survey of Consumer Expectations) both the mean and self-reported uncertainty of stated forecasts of key macroeconomic variables differ substantially across households. Importantly, we document that the accuracy of household expectations is systematically related to household wealth: 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 clearly monotone—especially at the lower-end of the wealth distribution, where the accuracy of expectations appears to decline with financial wealth.[1]

[1]A burgeoning literature has begun to document the various ways in which household and firm expectations differ from one another and within-groups (e.g., Coibion *et al.*, 2018; Coibion *et al.*, 2020; Reis, 2020; Andrade *et al.*, 2022; and Macaulay and Moberly, 2022). We contribute to this line of research by providing new evidence on the systematic (non-monotone) relationship between the accuracy of expectations and household wealth. Our estimates of the effects of other characteristics (e.g., education and sex) are consistent with those

Next, we embed dynamic information choice into an otherwise standard business-cycle model with idiosyncratic risk and incomplete markets, to explore households' heterogeneous incentives to acquire information. 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 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.

While prior work has examined the financial and macroeconomic consequences of costly information choices, it has primarily focused on once-and-for-all information choices that are identical across time and decision-makers (Grossman and Stiglitz, 1980; Sims, 2003; Hellwig and Veldkamp, 2009; Veldkamp, 2011; and Maćkowiak *et al.*, 2021). In contrast, in our model, households make *dynamic information choices* that depend on *individual characteristics* at any point in time. We refer to our synthesis of a heterogenous-agent economy with a model of dynamic information choice—and hence heterogeneous expectations—as `HetExp`.

We show how to adapt results from the heterogenous-agent literature to provide a novel solution method for GE frameworks with dynamic information choices and non-linear decision rules. Solving heterogenous-agent models with aggregate risk and non-linear decision rules is challenging. Our framework adds a further layer of complexity by allowing for heterogeneity also in expectations. We develop a tractable method to tackle these challenges. In closely-related work, Auclert *et al.* (2020) and Carroll *et al.* (2020) analyze a heterogenous-agent economy with *exogenous* information, based on Mankiw and Reis (2002) and Carroll (2003), and (partially) linearized policy rules. We document how the endogeneity of information and a fully non-linear approach profoundly alter the macro consequences of incomplete information.

Using our solution method, we calibrate our model framework to match key features of US macro and micro data. The model-generated distribution of household expectations rationalizes the survey evidence. Heterogeneous information choices—consistent with the data—naturally arise from differences in wealth and employment status that are the fundamental characteristics of heterogenous-agent economies. To understand households' heterogenous incentives to acquire information—and to highlight how they can profoundly shape macroeconomic outcomes—it is instructive to understand households' savings decisions.

Consider first *unemployed households,* who dissave to smooth consumption. The poorest—those at the borrowing constraint for all states in the next period—have no benefit from acquiring information. By contrast, unemployed households with modest wealth have substantial benefits, as savings mistakes are costly near the borrowing constraint. This is due to

found in the literature (e.g., Lusardi and Mitchell, 2014), lending further support to our findings.

the curvature of both the utility and policy functions being high. As wealth then increases, the value of information initially declines. In this region, errors about future labor-market prospects and capital returns push saving in opposite directions: Optimistic expectations about the job-finding rate reduce precautionary savings, while expecting a higher return on capital raise savings through intertemporal substitution. In the middle of the distribution, these two effects partially offset each other, temporarily lowering the value of information. At even higher wealth levels, financial assets eventually comprise the bulk of household resources and return risk, as a result, dominates, leading to increasing information acquisition.

The value of information for *employed households* is similar to that of wealthy unemployed households: Employed households always have (relatively) higher income, and thus cash-at-hand, compared to unemployed, and separation rates are small. The value of additional information about the state of the economy, as a result, starts low and then rises with wealth.

We show that such heterogeneity in information choices substantially alters 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 level, heterogeneous information choices feed back into wealth and income inequality, as differently informed households make disparate savings choices. Indeed, the introduction of heterogeneous information moderately exacerbates inequality. In particular, poor households with incomplete information are unable to exploit periods of good labor market prospects and high returns to build up financial wealth; wealthy households with incomplete information are likewise unable to effectively increase consumption and run down savings when higher future returns increase their permanent income. Both groups in response acquire information at higher rates than average, yet still face substantial information frictions in equilibrium. The introduction of heterogeneous, incomplete information, as such, modestly (but non-trivially) mitigates the lack of wealth inequality that exists in standard frameworks.

On the macro level, 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 pro-cyclicality of the savings response. By contrast, uninformed households' expectations about returns are sluggish to adjust, which makes household savings more pro-cyclical and the economy more volatile. This mechanism is itself somewhat dampened by increased information acquisition, due to the benefits of information about the economy being higher when the economy is more volatile. In equilibrium, not all households acquire information in every period, leading to 5-11 percent larger fluctuations in consumption and output relative to the full-information case.

Importantly, we show that these micro and macro results—and the mechanisms behind them—extend to alternative versions of our model that, for example, consider alternative information cost structures and features of US labor-income taxes. A prominent issue with the Aiyagari-Bewley-Huggett-Imrohoroglu class of models that we depart from is that it does not generate realistic wealth heterogeneity: The data display significantly more skewness in wealth than the models. Using an extension that builds on Bayer *et al.* (2024), in which a fixed share of wealthy entrepreneurs receive all pure profits in the economy, we show that our results nevertheless carry over to a framework that better matches the wealth distribution.

In the final part of the paper, we illustrate that the consequences of macroeconomic policies may also be altered once one accounts for households' heterogeneous information choices. To demonstrate this, we consider two popular policies to reduce inequality that directly affect the wealth-expectation nexus: (i) a wealth tax, modeled on the French tax system and the recent proposal in the US Congress;[2] and (ii) an increase in unemployment benefits.[3] Both policies disproportionately affect the resources available to one of two groups—the rich and the poor—that we find acquire information at higher rates than average.[4]

The direct impact of *the wealth tax* is to decrease household wealth; its indirect effect is to reduce information acquisition, as information on average rises with wealth. By reducing the information in the economy, economic volatility rises by 6 percent. In contrast, in the full-information case, the wealth tax has virtually no impact on aggregate volatility, despite a similar fall in aggregate wealth. The effect of the wealth tax on inequality is similarly surprising: A one percent tax hardly changes the Gini coefficient. As in the full-information case, the lump-sum rebate of the wealth tax, and the disincentive effect from lower after-tax returns, reduce savings across the wealth distribution—and hence inequality. However, in the `HetExp`-economy, this decline is countered by the increased "randomness" of savings, as information acquisitions fall. This, in turn, leads to larger over-accumulation of savings for uninformed, high-wealth households and more savings mistakes by the poor. Our framework, thus, provides one rationale for why several countries did not see increases in wealth inequality following the abolition of wealth taxes (e.g., Jakobsen *et al.*, 2020), although clearly there may be multiple, alternative drivers not captured by our framework.

The impact of *increased unemployment benefits* on the aggregate economy is akin to the effect of the wealth tax: Average wealth in the economy decreases—caused, in this case, by a

[2]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`.

[3]Specifically, we consider a 1 percent per annum wealth tax and an increase in the replacement rate from 40 to 50 percent of current wages.

[4]Notice that the policy reforms considered here are stylized and we focus more on the qualitative impacts than the quantitative numbers, as a proper treatment of these first-order public finance questions would require a more comprehensive model framework. Our focus is mainly on illustrating how taking into account heterogeneity in information choice can fundamentally alter, or even reverse, the consequences of policies that have a first-order impact on the wealth distribution.

reduced need of precautionary savings for poor households—and information acquisitions fall. In equilibrium, this fall in information once more substantially raises aggregate volatility (by 4 percent) and inequality also rises.

Overall, both policy experiments suggest that the consequences of dynamic, heterogenous information choices may substantially alter the relative costs and benefits of macroeconomic policies in unexpected directions. Our findings, thus, imply a Lucas-style critique (Lucas, 1976) of policy evaluations in full-information, heterogenous-agent economies.

Finally, three wider implications of our framework 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), as well as any relationship between, for example, education, gender, and household information (e.g., Lusardi and Mitchell, 2014 and Reiche, 2025). Notwithstanding such alternative drivers, we show that households' rational incentives to systematically acquire different information, depending on their wealth and employment status, fundamentally alter the dynamics and consequences of redistributive macroeconomic policies. We conjecture that behavioral heuristics, salience effects, and other household drivers of information choices would only increase the gap between the predictions of standard models and those relevant for macroeconomic policy.

Second, our emphasis on households' rational information choices, highlighting that wealth-poor households produce forecasts of close-to commensurable accuracy to wealth-rich, connects with the small-scale surveys in Harrington (1997), Shipler (2005), Newman (2009), Morduch and Schneider (2017), and others, showing that poor and working-class households are often substantially more aware of local conditions, prices, and opportunities, and devote more mental resources to tracking economic conditions, than wealthier households. We surmise that accounting for the full scale of such effects would only lead to a richer relationship between household informativeness and wealth than that detailed below. Our work is similarly closely related to that in finance, documenting that informed, wealthier investors often earn higher returns, which contributes to inequality among traders (e.g., Peress, 2004; Kacperczyk *et al.*, 2019; and Mihet, 2022). Our work, in part, accounts for a similar mechanism.

Lastly, because of the complexity of computing rational expectations equilibria in neoclassical heterogeneous-agent economies, several authors have proposed dimensionality reduction methods (e.g., Moll, 2024). 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 households themselves to decide which variables (or moments) to use to forecast the future state of the economy. In this sense, our framework presents a natural evolution of the Krusell and Smith (1998) computational approach.

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.[5] 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 household 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 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.[6] 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 findings of 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 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

[5] 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.

[6] See, for example, Stark *et al.* (2010), Faust and Wright (2013), and Bhandari *et al.* (2025). 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" (Appendix A).

Figure 1: Unemployment Expectations Across the Wealth Distribution

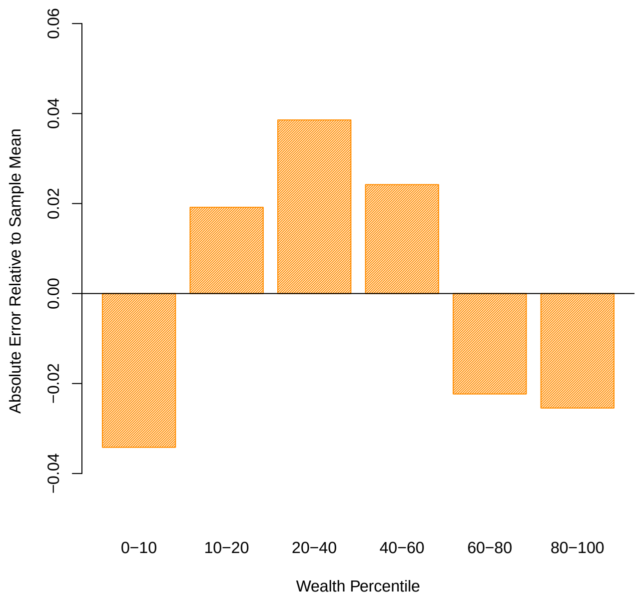

Panel (a): Relative Accuracy (SPF)

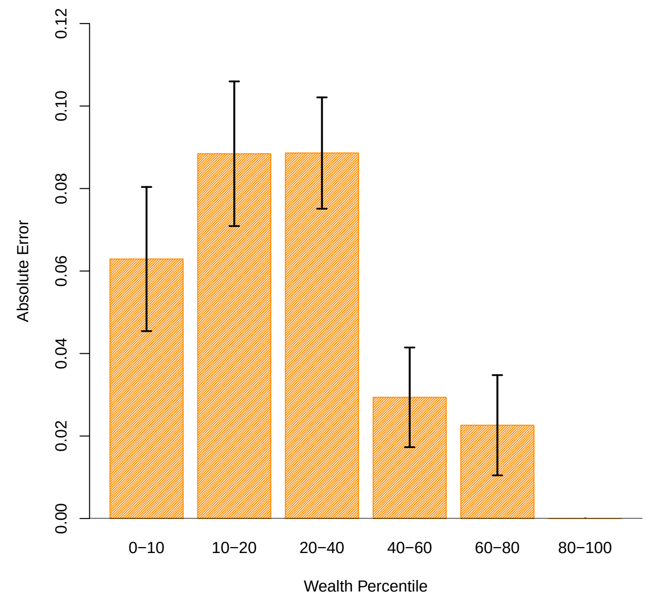

Panel (b): Coefficient on Wealth (SPF)

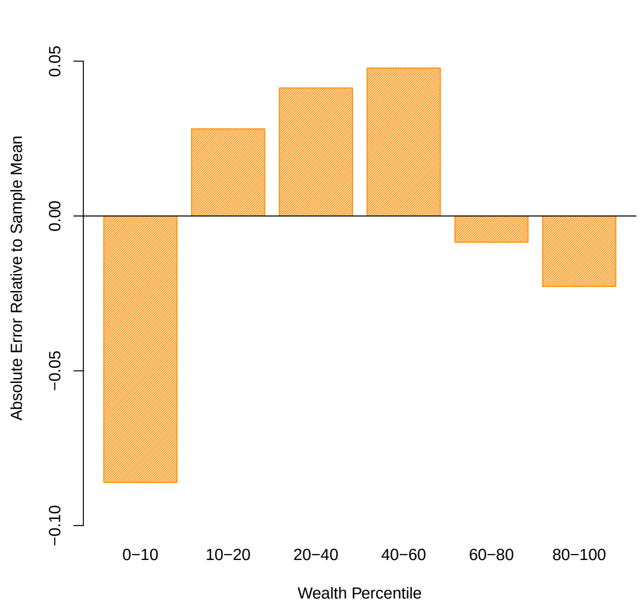

Panel (c): Relative Accuracy (BVAR)

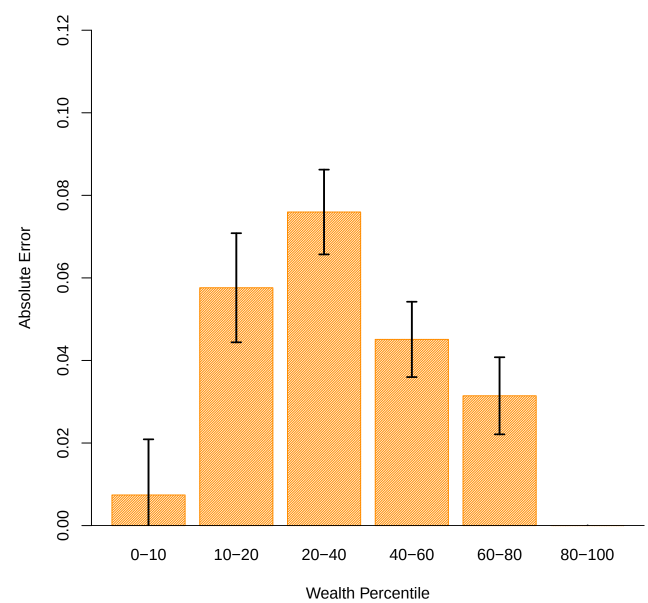

Panel (d): Coefficient on Wealth (BVAR)

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. The true probability is proxied by the average probability of rising unemployment from the SPF (Appendix A.2). Panel (b) plots the coefficient estimates on wealth from a regression of the absolute value of individual 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.7). Panel (c) and (d) proxy the true probability of rising unemployment with that from a standard forecasting VAR (Online Appendix A.3). Sample: 2013M10-2020M1.

suggests that household expectations are *heterogenous* across the wealth distribution.

The relationship in Panel (a) in Figure 1 may be contaminated by idiosyncratic factors, such as labor-market status, or aggregate shocks 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 regression coefficients exhibit a similar non-monotonic relationship to that in the raw data. All else equal, wealthier households make more accurate unemployment forecasts; yet the accuracy of households in the bottom decile is higher than those between the 20-40th percentile, although the difference is not statistically significant at conventional levels. The magnitudes are also meaningful: Considering a household in the 30th percentile of the wealth distribution instead of the 90th percentile, all else equal, decreases the accuracy of the household's expectations by around 9 percent. To benchmark the magnitude, having a university degree is estimated to only increase accuracy by 7 percent (Table A.7).[7] Crucially, Panel (c) and (d) in Figure 1 show that our results also extend to cases where we proxy the probability of rising unemployment with that computed from a standard forecasting VAR (Christiano *et al.*, 2005; Del Negro *et al.*, 2007), while Table A.9 in the Appendix shows that our results further extend to the case in which we directly control for the percentile rank of respondents in the wealth distribution. This illustrates the robustness of our findings.

Clearly, the estimates in Figure 1 cannot be interpreted as causal, as a household's wealth and the accuracy of its expectations are determined jointly (e.g., Section 3). That said, the evidence is clearly at odds with the assumption of common expectations embedded at the heart of the full-information rational expectations framework, showing instead that the state of household finances is closely tied to 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 their 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 inflation, all measures of uncertainty are also higher

[7]Our results on education are consistent with the findings of Lusardi and Mitchell (2014), among others, who show that education, in part, through its impact on financial literacy improves households' forecast accuracy. The focus of our analysis is on the three-way relationship between wealth, unemployment, and information, which is why we abstract from the influence of education in what follows.

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

Panel (a): Inflation Forecasts

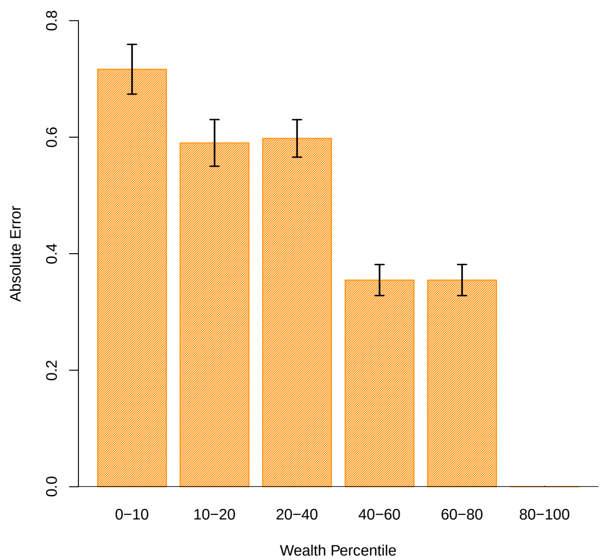

Actual Accuracy

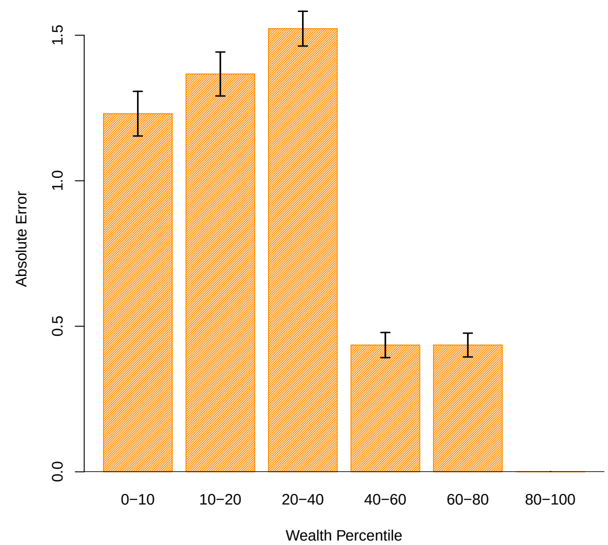

Perceived Accuracy

Panel (b): House Price Forecasts

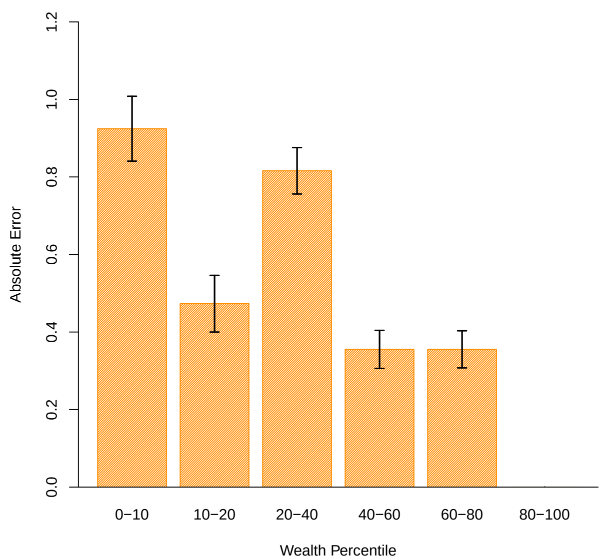

Actual Accuracy

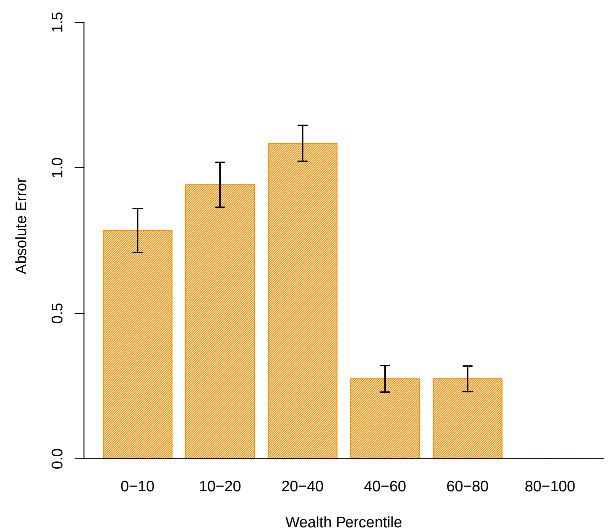

Perceived Accuracy

Note: Panels (a) to (d) show the *actual* and *perceived* accuracy of individual forecasts of one-year ahead CPI inflation and the annual growth rate of US house prices, respectively. All panels plot estimates from regressions of individual (actual or perceived) 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.

for households near the 20-40th percentile of the distribution than for poorest households.

Table A.8 in the Appendix uses the "worst-case" Likelihood Ratio test developed by Silvapulle and Sen (2005) to explicitly test for the monotonicity of the regression coefficients in Figure 1 and 2.[8] As evident, 6 of the 8 tests reject the null hypothesis at the 10 percent level, providing instead evidence in favor of an inverse-u-shaped relationship between wealth and the accuracy of expectations. Finally, in Appendix A.2 we perform additional robustness exercises. There, we furthermore show that, compared to professional forecasters, household expectations are more dispersed, less accurate, and perceived to be more uncertain. We will later leverage these additional moments to discipline our structural 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 extends the environment in Krusell and Smith (1998) 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.

3.1 Households

The economy consists of a continuum of heterogeneous households $i \in [0, 1]$, whose preferences over streams of consumption and information are described by the utility function:

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

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}, β the effective discount factor between two periods, c_{it} non-durable consumption at time $t = \{0, 1, 2, \dots\}$, κ_{it} the utility cost of acquiring information $\mathcal{I}_{it}$, and $\gamma > 0$. The

[8]Note that, because the regression coefficients are only set-identified under the null hypothesis, the LR test corresponds to a "worst-case test"—it uses the coefficients most in line with the null hypothesis. This makes any rejection of the null difficult, but conversely also implies that any rejection provides strong evidence against the null, which we interpret as evidence in favor of a humped-shaped relationship.

effective discount factor $\beta = \rho b$, where $1-\rho \in (0,1)$ and $b \in (0,1)$ are the per-period probability of death and the per-period discount factor, respectively.[9] The household's information set Ω_{it} accumulates according to $\Omega_{it} = \{\mathcal{I}_{it}, \Omega_{it-1}\}$.[10] The utility cost κ_{it} follows a type-I extreme value distribution with parameter α_κ, and is i.i.d. across households, time, and signals acquired. We introduce the extreme value shocks to account for unobserved heterogeneity in the survey data and to account for any utility costs of information processing.

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 equals $\mu \in (0,1)$. We assume that households cannot borrow but can save in physical capital k_{it}, whose net return equals $r_t - \delta$, where r_t denotes the stochastic rental rate and $\delta \in (0,1)$ the depreciation rate of capital. Households are assumed have access to perfect annuity markets.[11]

In addition to the borrowing constraint and a non-negativity constraint on consumption, household consumption-saving choices are restricted by the per-period budget constraint:

$$c_{it} + k_{it+1} = (1 + r_t - \delta)\,\rho^{-1} k_{it} + (1-\tau_t)\left[\epsilon_{it} w_t \bar{l} + (1-\epsilon_{it})\,\mu w_t\right] - \eta(\mathcal{I}_{it}), \tag{3.2}$$

where $\eta(\mathcal{I}_{it})$ denotes the resource cost of acquiring $\mathcal{I}_{it}$, and τ_t is the tax rate on income. We refer to the right-hand side of (3.2) as a household's *cash-at-hand after choosing to acquire information*, and denote it by m_{it} in what follows. A household maximizes utility (3.1) subject to the budget constraint (3.2) and non-negativity constraints on c_{it} and k_{it+1}.

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 aggregate capital and employment in period t, respectively. Aggregate productivity z_t is stochastic and follows a first-order Markov process that takes two values,

[9] Consistent with this, a fraction $1-\rho$ of households are born every period with zero initial wealth.

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

[11] The capital of the deceased is, thus, used to pay an extra return on capital of ρ^{-1} .

$z_t \in \{z_l, z_h\}$ 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 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 $\tau_t = \frac{\mu u_t}{\bar{l} L_t}$. Appendix C.3 and Section 5.4 discuss an alternative setup with acyclical income taxes. We consider a tax on household wealth and the response of the economy to changes in the replacement rate μ in Section 6.

3.4 Timeline and Information Structure

At the start of each period, idiosyncratic $(\epsilon_{it}, \kappa_{it})_i$ and aggregate shocks (z_t) realize. Firms rent capital and labor, production takes place, and factor prices are determined. Households, who do not observe the realization of the aggregate shock but know the savings they brought into the period and their employment status, then choose which signals $\mathcal{I}_{it}$ to acquire about the current state of the economy from a maximum signal set $\mathcal{I}^{\max}$.[12] We assume that $\mathcal{I}^{\max}$ 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).[13] Finally, conditional on information choices and factor payments, households make consumption and savings choices (c_{it} and k_{it+1}, respectively) and the period ends.

3.5 Recursive Formulation of the Household Problem

Given the timeline and informational assumptions, we develop a recursive formulation of the household problem. Let $S = (\Gamma, z)$, where Γ denotes the cross-sectional distribution of

[12] We assume a household that is born in period t inherits the information set of its "parent" (i.e., the corresponding household that died in period t). We experimented with different assumptions about the informativeness of recently born households and found that it makes little difference to our results below.

[13] 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 (see also Sections 3.5 and 3.7).

capital and employment status. We denote an individual household's first-order belief about S by $\mathcal{P}_i(S)$.[14] 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.[15] The *aggregate state* of the economy can then be described by $\Sigma = (S, \mathcal{P})$, while the *individual state* variables at the stage in which consumption and savings choices are made (*Stage 2*) are described by $\sigma_{i,2} = (m_i, \epsilon_i, p_i)$, where m_i is household i's cash-at-hand *net* of information costs. We denote next period's realization of variable x by x' and previous period's realization by x_{-1}. The individual state variables at the stage where information choices are made (*Stage 1*) are $\sigma_{i,1} = (k_i, \epsilon_i, p_{i,-1})$.

Stage 2: At the end of the period, after making its information choices, a household chooses consumption c_i and savings k_i' out of cash-at-hand net of information acquisition costs:

$$\begin{aligned}
\mathcal{W}(m_i, \epsilon_i, p_i) &= \max_{c_i, k_i' \geq 0} \frac{c_i^{1-\gamma} - 1}{1-\gamma} + \beta \mathbb{E}\left[\mathcal{V}(k_i', \epsilon_i', p_i) \mid \Omega_i\right] \\
\text{subj. } &\text{ to} \\
c_i + k_i' &= m_i
\end{aligned} \tag{3.5}$$

where $\mathcal{W}(\sigma_{i,2})$ and $\mathcal{V}(\sigma_{i,1})$ are a household's value functions *after* and *before* information acquisition, respectively, and 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 (i.e., $k_i' = g(\sigma_{i,2})$), while $h(\cdot)$ characterizes its consumption choice (i.e., $c_i = h(\sigma_{i,2})$). We assume that households rationally use the equilibrium law of motion for the aggregate state, which we denote by H (i.e., $\Sigma' = H(\Sigma)$), and the exogenous transition matrix for z, Π_z, to form their prior expectation about tomorrow's state—and hence wealth k_i' embedded in $\sigma_{i,1}'$—from today's posterior beliefs. *Stage 2* posterior beliefs p_i are linked to *Stage 1* priors $p_{i,-1}$ through Bayes' Rule and the information choice $\mathcal{I}_i$ the household makes in *Stage 1*.

[14] 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.

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

Stage 1: At the beginning of the period, households choose what information to acquire:

$$\begin{aligned}
\mathcal{V}\left(k_i, \epsilon_i, p_{i,-1}\right) &= \max_{\mathcal{I}_i \subseteq \{\emptyset \cup \mathcal{I}^{\max}\}} \mathbb{E}\left[\mathcal{W}\left(m_i, \epsilon_i, p_i\left(\mathcal{I}_i\right)\right) - \kappa_i(\mathcal{I}_i) \mid \Omega_{i,-1}\right] \qquad (3.6)\\
\text{subj. to} & \\
m_i &= \left[1 + r\left(\Sigma\right) - \delta\right] \rho^{-1} k_i + (1-\tau)\left[\epsilon_i w\left(\Sigma\right) \bar{l} + (1-\epsilon_i)\, \mu w\left(\Sigma\right)\right] - \eta\left(\mathcal{I}_i\right).
\end{aligned}$$

Recall that information acquisition entails *both* a utility cost $\kappa\left(\mathcal{I}_i\right)$ and a resource cost $\eta\left(\mathcal{I}_i\right)$—and that individual states and expectations themselves also depend on the information choice $\mathcal{I}_i$. A household's expectation in the first stage is computed using previous period's posterior belief $p_{i,-1}$ (and hence information). We let $\iota\left(\cdot\right)$ denote the function that characterizes its optimal information choice (i.e., $\mathcal{I}_i = \iota(\sigma_{i,1})$).

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 $\mathcal{I} \subseteq \{\emptyset \cup \mathcal{I}^{\max}\}$:

$$\text{Prob}\left(\mathcal{I} \mid \sigma_{i,1}\right) = \frac{e^{\mathbb{E}\left[\alpha_\kappa \mathcal{W}\left(m_i(\mathcal{I}_i), \epsilon_i, p_i(\mathcal{I}_i)\right) - \kappa_i(\mathcal{I}) \mid \Omega_{i,-1}\right]}}{\sum_{\tilde{\mathcal{I}} \subseteq \{\emptyset \cup \mathcal{I}^{\max}\}} e^{\mathbb{E}\left[\alpha_\kappa \mathcal{W}\left(m_i(\tilde{\mathcal{I}}), \epsilon_i, p_i(\tilde{\mathcal{I}})\right) - \kappa_i(\tilde{\mathcal{I}}) \mid \Omega_{i,-1}\right]}}, \qquad (3.7)$$

where we suppress arguments other than $\mathcal{I}$ in the policy functions, yielding the value function:

$$\mathcal{V}\left(\sigma_{i,1}\right) = \frac{\gamma_E}{\alpha_\kappa} + \frac{1}{\alpha_\kappa} \log\left(\sum_{\tilde{\mathcal{I}} \subseteq \{\emptyset \cup \mathcal{I}^{\max}\}} e^{\mathbb{E}\left[\alpha_\kappa \mathcal{W}\left(m_i(\tilde{\mathcal{I}}), \epsilon_i, p_i(\tilde{\mathcal{I}})\right) - \kappa_i(\tilde{\mathcal{I}}) \mid \Omega_{i,-1}\right]}\right), \qquad (3.8)$$

where γ_E is the Euler-Mascheroni constant (McFadden *et al.*, 1973).

3.6 Recursive Incomplete Information Competitive Equilibrium

The definition of a *Recursive Competitive Incomplete Information Equilibrium* (RIICE) extends the standard definition of a Recursive Competitive Equilibrium to the case with incomplete information: A RIICE is a law of motion $H\left(\cdot\right)$, a pair of individual value functions $\{\mathcal{V}\left(\cdot\right), \mathcal{W}\left(\cdot\right)\}$, policy functions $\{g\left(\cdot\right), h\left(\cdot\right), \iota\left(\cdot\right),\}$, as well as pricing functions $\{r\left(\cdot\right), w\left(\cdot\right)\}$ such that: (i) $\{\mathcal{V}\left(\cdot\right), \mathcal{W}\left(\cdot\right), g\left(\cdot\right), h\left(\cdot\right), \iota\left(\cdot\right)\}$ solve the household's optimization problem (i.e., Equations 3.5-3.6) given $H\left(\cdot\right)$; (ii) $r\left(\cdot\right)$ and $w\left(\cdot\right)$ satisfy firm maximization (i.e. Equation 3.4), (iii) $H\left(\cdot\right)$ is generated by policy functions $g\left(\cdot\right)$, $h\left(\cdot\right)$, and $\iota\left(\cdot\right)$, the Markov process $\Pi_{z',\epsilon'|z,\varepsilon}$, as well as Bayes' Rule, using the information contained in $\left(\mathcal{I}_i\right)_i$ and current beliefs described in $\mathcal{P}$;[16] and (iv) market-clearing conditions hold for capital and goods markets (e.g., $Y(\Sigma) + (1-\delta)K = \int \left(g(\sigma_{i,2}) + h(\sigma_{i,2}) + \eta\left[\iota\left(\sigma_{i,1}\right)\right]\right) di$).[17]

3.7 Remarks on Modeling Assumptions

The above recursive formulation helps clarify the nature of the two-sided relationship between wealth and information that exists in the model. On the one hand, a household's capital holdings and employment status, in addition to its prior expectation, help determine the extent of the household's information acquisition at the first stage of any period (Equation 3.6). Yet, on the other hand, a household's information choice also helps determine the household's consumption-savings choice in the second stage, through its expectations, and hence the household's future wealth level (Equation 3.5). This two-sided interaction, which we label the *expectation-wealth nexus*, will be central for the aggregate consequences of the wealth-accuracy relationship documented in the data (Section 2).

The recursive formulation further illustrates that our framework falls within the broader class of "noisy rational expectations" models. For example, although households are uncertain about the current state of the economy, they rationally use the law of motion to form their expectations about future productivity conditional on current information—and hence the likelihood of, for instance, future information purchases. As such, our framework is closely related to the work on "costly information acquisition" within rational expectations frameworks (e.g., Grossman and Stiglitz, 1980; Veldkamp, 2011) and its extension to "rational inattention" (e.g., Sims, 2003; 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 framework, in this context, is to highlight the macroeconomic consequences of *dynamic, heterogenous information choices.*[18]

Finally, notice that our benchmark model allows for both a resource cost of information

[16] For example, we can construct the part of $H(\cdot)$ that pertains to the marginal distribution of capital, $H_k(\cdot)$, conditional on employment states and productivity, as follows. For all measurable sets Δk:

$$H_k(\Sigma)(\Delta k, \epsilon, z) = \sum_{\tilde{\epsilon}} \Pi_{\epsilon, z|\tilde{\epsilon}, \tilde{z}} \cdot \int \mathbf{1}\left\{g\left[m\left(k, \tilde{\epsilon}; \Sigma\right), \tilde{\epsilon}, p\right] \in \Delta k\right\} \cdot d\Phi\left(k, \tilde{\epsilon}, p\right),$$

where $\Phi(\cdot)$ denotes the joint cross-sectional distribution of (k, ϵ, p).

[17] Notice that the information cost $\eta(\mathcal{I})$ is paid in terms of the numéraire. Thus, the goods-market clearing condition accounts for total information expenditures, $\int \eta(\mathcal{I}_i)\, di$. Baley and Veldkamp (2025) discuss the different assumptions possible for information costs (i.e., whether costs are paid in goods, labor, or capital).

[18] In our framework, we deliberately abstract from the "non-instrumental value" of information and focus only on the consequences that information has on utility through improving decision-making. Brunnermeier and Parker (2005) and Caplin and Leahy (2019), among others, consider the consumption-saving biases that arise when households engage in "wishful thinking", in which expectations maximize average felicity, optimally balancing the benefit of optimism against the costs of worse decision-making.

$\eta(\mathcal{I}_i)$ (Equation 3.6) as well as a utility cost $\kappa_i(\mathcal{I}_i)$ (Equation 3.6). The former captures the physical costs associated with the acquisition of information, while the latter captures the (cognitive) utility costs associated with its processing (e.g., Veldkamp, 2011), in addition to accounting for unobserved heterogeneity in the survey data.[19] Crucially, the latter, for example, also captures the utility costs associated with the processing of market-based information (see below for further discussion). Our calibrated framework in Section 5 highlights the distinctive footprints that each of these information costs has at the aggregate level.

4 Solution Method and Calibration

In this section, we first outline our procedure for computing RIICE equilibria. Our description here is deliberately non-technical. We include it in the main text because our description is intimately linked to the three-way relationship between information choices, consumption-savings decisions, and the macroeconomy that is at the heart of our analysis—and because it illustrates the underlying methodology behind our framework. 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 our 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 view our model as

[19] Veldkamp (2011) discusses how a resource cost of information, for example, also captures the costs associated with "getting advice from experts" or out-sourcing decisions.

a natural extension of the Krusell-Smith approach 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 predetermined information, notice that the Krusell-Smith approach also imposes *common knowledge* over both z and $\mathbf{m}_t$, as all households form expectations using the same information. Our framework, by contrast, relaxes the common knowledge assumption by allowing for heterogenous information choices.[20] The RIICE framework, as a result, allows for the study of the three-way interaction between incomplete common knowledge, aggregate dynamics, and inequality that 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).[21]

4.2 The Specification of Moments

For the set of moments households can choose between, we follow Krusell and Smith (1998) and consider only the first moment of the cross-sectional distribution of capital, $\mathbf{m} = \int g(\sigma_{i,2})\, d\sigma_{i,2} = K_t$. 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 extremely

[20] Both our framework and the Krusell-Smith framework assume common knowledge over the underlying structure of the economy (preferences, technology, and so on). 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.

[21] Our assumption that households know the law of motion $H(\cdot)$ and use it to form expectations is in line with a large literature on noisy rational expectations, going back to at least Lucas (1972). In all of this work, agents observe signals about unobserved fundamentals, characterizing the state of the economy, but crucially know the mapping between the signals they observe and the underlying fundamentals, as well as all other (equilibrium) relationships (e.g., Woodford, 2002; Sims, 2003; Lorenzoni, 2009; Maćkowiak and Wiederholt, 2009; Coibion and Gorodnichenko, 2012, 2015; and Angeletos *et al.*, 2021 among others). This allows us to focus on the implications of uncertainty about the state of the economy and not conflate any effects with those that also arise from uncertainty about the structure of the economy.

accurate predictions about the future capital stock K_{t+1} (Den Haan *et al.*, 2010). We therefore set $\mathcal{I}^{\max} = \{z_t\}$, 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.3).

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.3), our assumption can be viewed as allowing households to choose to observe elements of the optimal signal.[22] Finally, in previous work (Broer *et al.*, 2022), we have shown 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 capital, k_i, their employment status, ϵ_i, their prior over whether the economy is in the high productivity state, $p^z_{i,-1} \equiv \text{Prob}\,(z = z_h \mid \Omega_{i,-1})$, and their prior over the capital stock $p^K_{i,-1}$. Households can then choose whether to observe contemporaneous productivity z.[23] In line with the above notation, let $\tilde{\sigma}_{i,2} \equiv \left(m_i, \epsilon_i, p^z_i, p^K_i\right)$ and $\tilde{\sigma}_{i,1} \equiv \left(k_i, \epsilon_i, p^z_{i,-1}, p^K_{i,-1}\right)$. The two-stage optimization problem can then be stated as:

Stage 2: Consumption-savings choice

$$\begin{aligned}
\tilde{\mathcal{W}}\left(\tilde{\sigma}_{i,2}\right) &= \max_{c_i, k_i' \geq 0} \frac{c_i^{1-\gamma} - 1}{1-\gamma} + \beta \mathbb{E}\left[\tilde{\mathcal{V}}\left(\tilde{\sigma}'_{i,1}\right) \mid p^z_i, p^{K\prime}_i\right] \qquad (4.1)\\
\text{subj. to} & \\
c_i + k_i' &= m_i \\
K' &= \tilde{H}(z, K) \\
p^z_i &= \text{Prob}\left(z = z_h \mid \Omega_i\right), \quad p^{K\prime}_i = \mathbb{E}\left[\tilde{H}(z, p^K_i) | \Omega_i\right].
\end{aligned}$$

[22] We thank Mirko Wiederholt for this comment.

[23] For computational tractability, we assume households have a point prior on capital, $p^{K\prime}_i = \mathbb{E}\left[\tilde{H}(z, p^K_i)|\Omega_i\right]$. Appendix C.6 shows that our results are extremely similar when households internalize uncertainty in their forecasts about capital, analogous to the findings in Den Haan *et al.* (2010) and Broer *et al.* (2022).

Stage 1: Information choice

$$\begin{aligned}
\tilde{\mathcal{V}}\left(\tilde{\sigma}_{i,1}\right) &= \max_{\mathcal{I}_i\in\{\emptyset,z\}} \mathbb{E}\left[\tilde{\mathcal{W}}\left(\tilde{\sigma}_{i,2}\right) - \kappa_i(\mathcal{I}_i) \mid p^z_{i,-1}, p^K_i\right] && (4.2)\\
\text{subj. to} & \\
m_i &= \left[1 + r\left(z, K\right) - \delta\right]\rho^{-1}k_i + (1-\tau)\left[\epsilon_i w\left(z, K\right)\bar{l} + (1-\epsilon_i)\,\mu w\left(z, K\right)\right] - \eta\left(\mathcal{I}_i\right)
\end{aligned}$$

Notice that $\tilde{H}(z, K)$ in Equation (4.1) 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. Households update their priors after their information acquisition decision using Bayes' rule. Consistent with the notation above, we let $\tilde{h}(\cdot)$, $\tilde{g}(\cdot)$, and $\tilde{\iota}(\cdot)$ denote the policy functions for consumption, capital, and information, respectively (i.e., $c_i = \tilde{h}(\tilde{\sigma}_{i,2})$, $k'_i = \tilde{g}(\tilde{\sigma}_{i,2})$, and $\mathcal{I}_i = \tilde{\iota}(\tilde{\sigma}_{i,1})$). The aggregate state variable is denoted by $\tilde{\Sigma} = (z, K)$ in what follows.

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. As in Krusell and Smith (1998), we assume a log-linear relationship between K' and K, whose coefficients depend on the (boom or bust) realization of z. Second, we solve the household's two-stage problem in (4.2) and (4.1) conditional on $\tilde{H}(\cdot)$ and the cross-sectional distribution of information, income, and wealth. We use Equation (3.7) and the EGM-algorithm from Carroll (2006) to solve for the policy functions and value function iteration to solve the value functions. 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 a time-series for K, and estimate a new law of motion $\tilde{H}'(\cdot)$ (i.e., new log-linear relationships). 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.[24] 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 $b = 0.987$ and the probability of death $1 - \rho = 0.005$ to generate a quarterly capital-output ratio of 10 (Carroll *et al.*, 2017) and to be consistent with an expected work-life of 45 years. We calibrate the degree of relative risk aversion γ and the information cost parameters α_κ and η 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 $\eta = 0.0028$ (equivalent roughly to 0.1 percent of quarterly pre-tax wages) to match our empirical finding that forecast accuracy increases in wealth for richer households (Section 5.3).[25] We set the scale parameter α_κ equal to $1/15 \times 10^{-4}$ to capture the dispersion in unemployment expectations observed in the SCE (even for households with similar observable characteristics). To check 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 US 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 larger in the data, but

[24] 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 (Ravn and Uhlig, 2002).

[25] The benefit of additional information for wealthy households arises mainly from improved predictions about the future rate of return on capital. Yet 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	1.24	0.68
Model Simulated Data	1.28	0.52

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).

that 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.

5 Quantitative Results

In order to understand the consequences and drivers of households' information choices, we proceed in four steps. First, we study how different information choices affect households' savings decisions—their intertemporal decision variable in the model. Second, we characterize how households' 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 allows us to match the micro data on expectations. Finally, we explain the impact of the wealth-expectations nexus on aggregate dynamics and inequality, and show that our results carry over to several alternative versions of our benchmark model.

5.1 Savings Choices and Information

We start by describing how information affects households' savings decisions. In the model, households save to smooth consumption in the face of volatile income, and to intertemporally substitute consumption in response to movements in the interest rate. The state of the economy that households can acquire information about affects these both through *exogenous fluctuations* in productivity and employment, and through *endogenous fluctuations* in capital—all of which matter for households' human income and future rates of return.

Figure 3 studies 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 households have the same prior expectation over the capital stock, and plot the differences in savings choices ("informed minus uninformed") as a function of a household's cash-at-hand

Figure 3: Household Saving Choices and Wealth

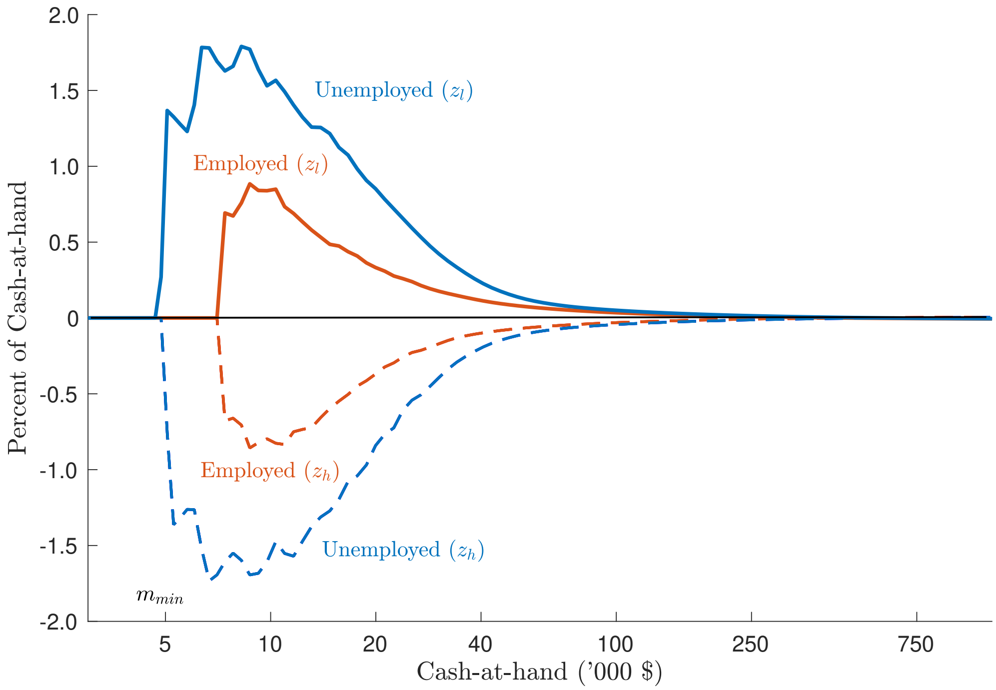

Note: The figure plots differences between the savings choice of an "informed household" and an "uninformed household" as a percentage of cash-at-hand at the mean prior p_i^K for aggregate capital K. We plot this difference in a recession (solid lines, $z = z_l$) and a boom (dashed lines, $z = z_h$) for both an unemployed (blue lines) and an employed (red lines) household. We plot household savings choices as a function of cash-at-hand and use 2020 values of US household income to convert cash-at-hand in our model to USD (\$) amounts.

and its employment status. We interpret this measure as the "static cost" of being uninformed, as we keep the prior over capital fixed. Dynamically, being uninformed will, however, also lead to errors in household expectations of capital, which we consider below.

As Figure 3 shows, informed households, all else equal, save more than uninformed in recessions (z_l), as they know the probability of becoming unemployed (or staying unemployed) is higher. Conversely, informed household save less in booms (z_h). However, the percentage difference in savings between the two cases—informed vs uninformed—varies strongly across the distribution of cash-at-hand, and between unemployed and employed households.

The poorest households optimally reduce their savings to zero independent of the current state of the economy. As such, information does not change their savings behavior. Savings rates of informed households, however, differ strongly from those of the uninformed at low-but-positive levels of cash-at-hand, where the non-linearity of the savings policy function is most pronounced near the kink at the borrowing constraint. Information about productivity helps such household better predict the probability of employment and future wages, and thus future

Figure 4: Household Saving Choices and the Expectation of Capital

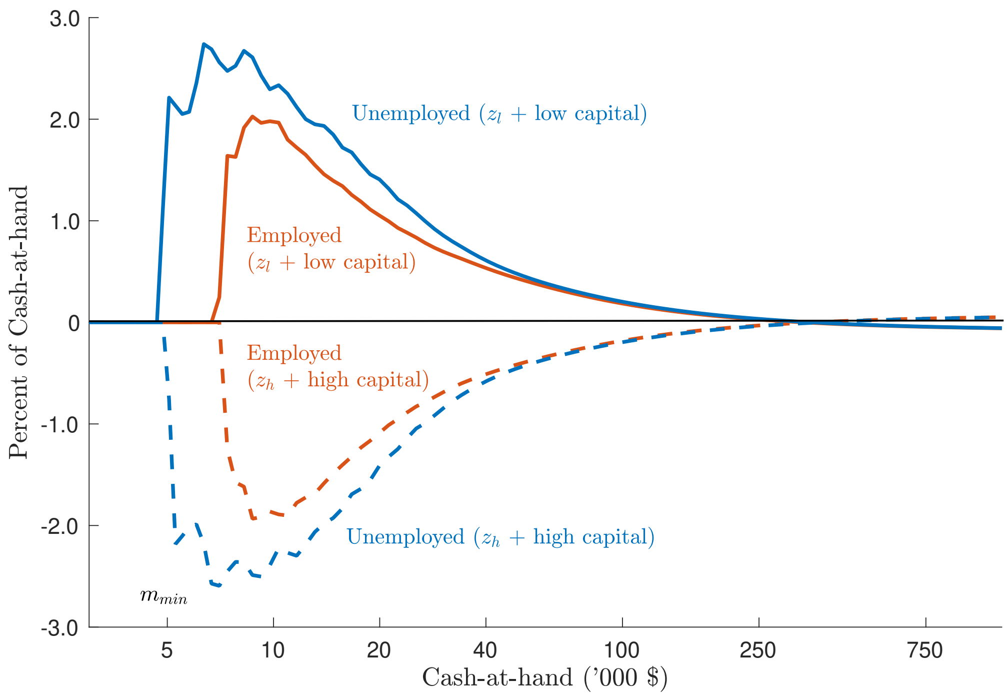

Note: The figure plots differences between the savings choice of an "informed household" and an "uninformed household" as a percentage of cash-at-hand. Specifically, in a boom (z_h), we plot the difference between the savings policy function of an informed household that has a prior p_i^K over the capital stock K_t that is one standard deviation higher than the mean (i.e., high capital) 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, low capital). We plot these differences for unemployed (blue lines) and employed (red lines) households. We plot savings choices as a function of cash-at-hand and use 2020 values of US household income to USD ($) amounts.

labor income, which is the main determinant of future consumption for low-wealth households. This difference between informed and uninformed savings is largest for the unemployed, whose job-finding rate differs strongly across booms and busts.

As cash-at-hand increases, the difference between informed and uninformed savings eventually falls for two reasons. First, because accurate predictions of future labor income have a smaller impact on savings when a household is wealthy. Second, although exogenous fluctuations in productivity alter rates of return, which are important for the wealthy, their effect on savings of the rich is small. This is due to the combination of a low elasticity of intertemporal substitution (and hence a weak substitution effect) and a moderate persistence of productivity shocks, and hence interest rates (and thus weak income effects). This decline is further reinforced by the savings function becoming increasingly linear at higher wealth levels.

While knowing the current productivity state is important, acquiring information also causes households to perceive the economy's endogenous dynamics more accurately—particularly the evolution of the capital stock. Notice that in Figure 3 informed and uninformed households have the same expectation of the capital stock. Any perceived differences in wages and rates of return from being informed stem only from perceived differences in productivity and employment. With imperfect information, however, uninformed households also perceive less accurately the dynamics of capital (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, thus, have less accurate expectations of capital—and because capital is persistent the resulting errors in expectations of future wages and returns induce strong income and substitution effects, resulting in less accurate savings decisions. We illustrate how such mean-biased capital expectations changes savings behavior in Figure 4: In a boom, we plot the difference between (i) the savings function of an informed household that has an expectation of capital that is one standard deviation higher than the mean; and (ii) an uninformed household with an expectation at the mean capital stock (solid lines). We also plot the same for a bust, 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 capital expectations increases the magnitude of savings mistakes that the majority of unconstrained households make. The additional effect is especially pronounced for low-wealth, employed households, whose savings mistakes are amplified by a factor of almost two. These savings mistakes are explained by errors in households' expectations of future wages—crucial for these households' savings decisions—which now combine the effects of errors due to mistaken *exogenous components* of labor income (i.e., productivity and employment transitions) and *endogenous components* (i.e., capital accumulation). Since a higher capital stock in booms provides an additional, persistent boost to future labor income through higher wages, both components here compound each other.

That said, although accounting for capital dynamics amplifies savings differences for the bulk of households, this is not true for the wealthy—the difference between informed and uninformed savings indeed flips sign at high wealth levels (above \$400,000-\$500,000). Although labor income co-moves positively with capital (as wages are high when capital is), financial income co-moves inversely (as rates of return decline with the capital stock). Since the persistence of the capital stock is higher than productivity, the resulting income effect on consumption is larger than that from productivity fluctuations. As a result, *informed households* in the upper-part of the wealth distribution (above the 90th percentile) *optimally save less* when the capital stock is low and interest rates are high. The additional return on their stock of wealth, nevertheless, ensures that total wealth accumulation by the rich remains elevated when the capital stock is low. As we will see in Section 5.6, such heterogenous effects of information—making savings more correlated with returns for the poor but less so for the rich—have important consequences for the cyclicality and the inequality in our economy.

Combined, Figure 3 and 4 illustrate how acquiring information "statically and dynamically" impacts households' savings choices. Next, we explore how these forces interact to shape households' information choice. This will cast further light on the wealth-expectations nexus.

5.2 Household Information Choices

Information choices are most easily described by the probability of acquiring of information. Figure 5 plots this probability as a function of a household's state variables—cash-at-hand and the prior over productivity—separately for the employed and unemployed, and for a low and a high prior over the capital stock. Combined, the results in Figure 5 showcase the rich heterogeneity that exists in the incentives to acquire information. 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 savers in the economy—are less likely to acquire information, especially at low levels of wealth. Because separation rates are only mildly cyclical (5.0 percent in recessions and 3.5 percent in booms), their optimal savings change little across boom-bust states (Figure 3). With low unemployment risk and a moderate persistence of boom-busts, they are further unlikely to hit the borrowing constraint soon. Thus, they expect to continue accumulate assets over time, so that the average savings error from remaining uninformed is small, as discussed in the previous section. The option to acquire information after a future job loss further lowers the value of doing so today. As wealth rises, however, the benefit of predicting returns on increasing financial wealth rises and the cost of acquiring information relative to wealth falls, increasing the probability of information acquisition.

Now, consider instead *unemployed households*. The unemployed are dissavers, as they attempt to smooth consumption. At low enough values of cash-at-hand, they consume all of their resources and end up at the borrowing constraint. Thus, at sufficiently low values of cash-at-hand, unemployed households almost never acquire information. These households are at the borrowing constraint for all states that realize tomorrow, and hence have no benefit from acquiring information today. However, at slightly higher values of cash-at-hand unemployed households start to acquire information with high probability. As job-finding rates are highly cyclical, the optimal rate of decummulation (or slope of the policy function) is higher in booms than busts. Savings mistakes from being uninformed are costly, especially near the borrowing constraint where the curvature of both the utility function and policy functions are high. In this range, unemployed households almost uniformly choose to acquire information.[26]

[26] The same incentives raise information-acquisition probabilities of uninformed employed households, who

Figure 5: Information Acquisition Probabilities

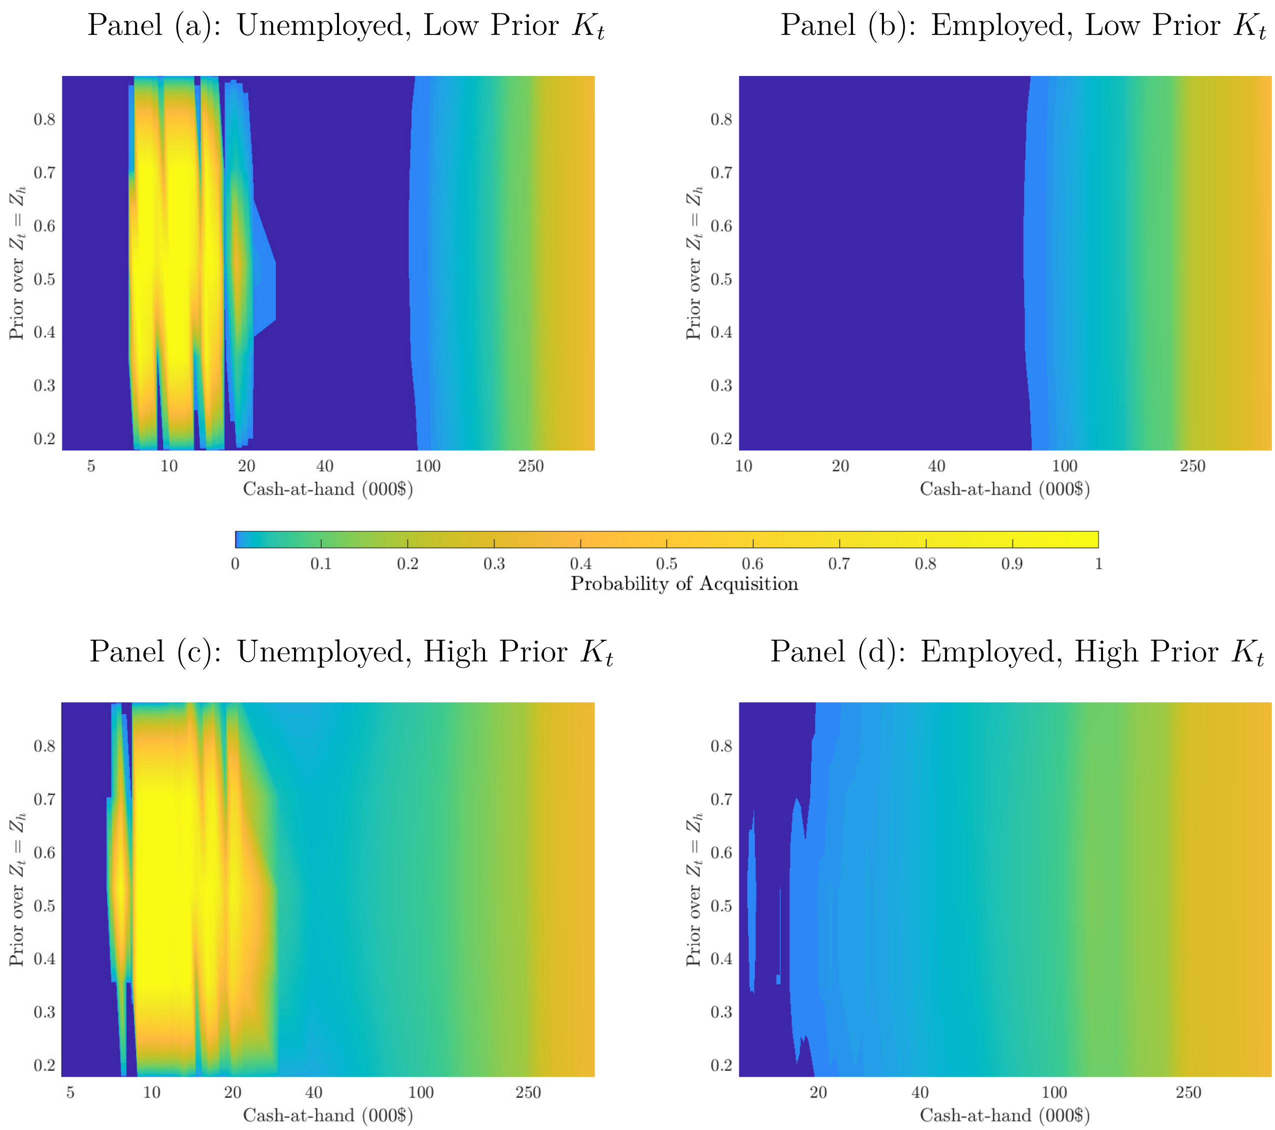

Note: The figure plots the probability of information acquisition (from 0-1) for different values of the prior belief that current productivity z_t is high and for different values of household cash-at-hand. The figure uses our baseline calibration to depict the probabilities and evaluates them at a "low" (i.e., 25% below the mean) and "high" (i.e. 25% above the mean) prior over the aggregate capital stock K_t. The left-hand (right-hand) panels show the probabilities for an unemployment (employed) household. We use 2020 values of US household income to convert cash-at-hand in our model to USD ($) amounts.

As wealth rises further, marginal utility eventually falls, and the value of information initially drops. The household is no longer at risk of imminently hitting the constraint, and expectation errors about the labor market and capital returns have opposing effects on savings: Misperceiving a higher job-finding rate, all else equal, reduces the need for precautionary savings, while misperceiving a higher return on capital, by contrast, increases savings through intertemporal substitution. For unemployed households in the middle of the distribution, these two effects partially offset each other, temporarily lowering the value of information. The value of information, nevertheless, then slowly starts to rise again with wealth for the same reasons as discussed in the case of the employed.

Finally, notice that households with low-levels of wealth are, on balance, more likely to acquire information when their prior about the current capital stock is high (Panel (c) and (d) in Figure 5). This is because higher wages implied by abundant capital increase the dispersion of incomes across employment states. This increases the value of accurately predicting employment prospects based on information about the current state of the economy, and thus the probability of information acquisition.

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 to shape the accuracy of household expectations across the wealth distribution. This will allow us to also 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 household 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 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

face a smaller but positive probability of becoming unemployed, at moderate levels of cash-at-hand, as seen in Panel (d) in Figure 5.

Figure 6: Accuracy of Forecasts: Model vs. Data

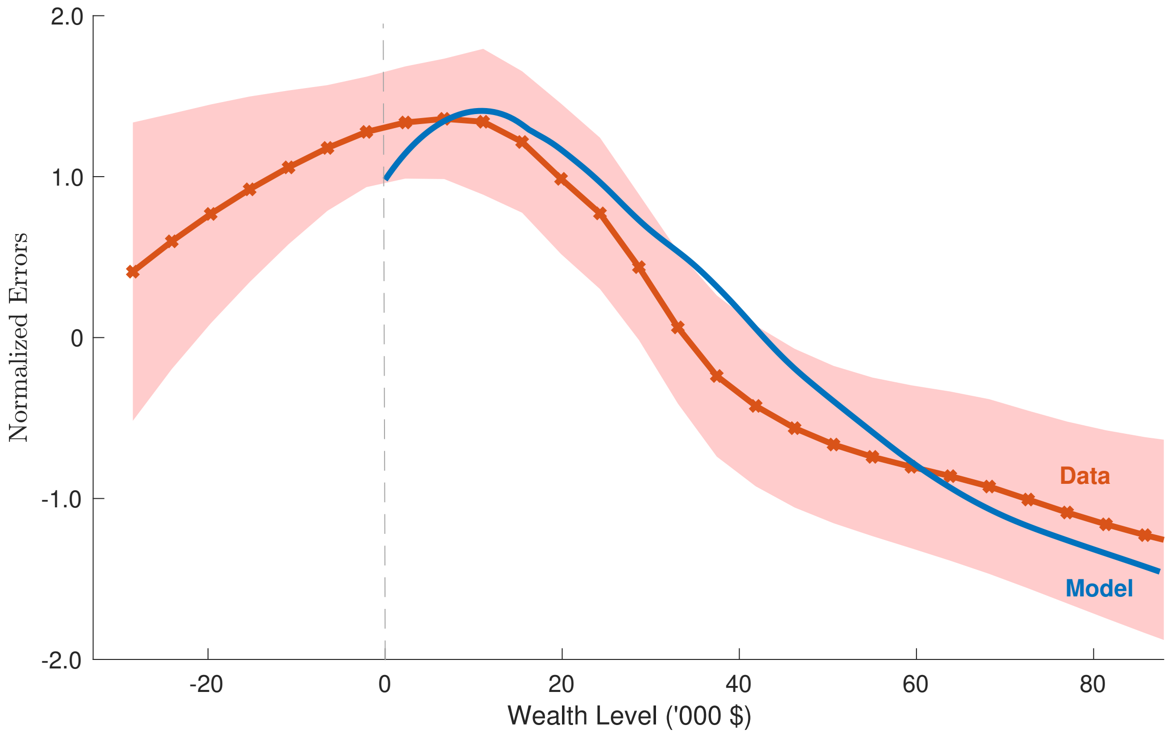

Note: The figure shows the estimated relationship between (the absolute value of) normalized errors of the one-year ahead probability of the unemployment rate increasing and household wealth. We plot this relationship both in the SCE data and in the calibrated model (see also Section 2). We use a local 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 quarterly US household income to convert values in the data and in the model to $ (USD) amounts.

to the observable decline in accuracy at low-levels of wealth. Second, as wealth increases, the probability of acquiring information eventually becomes monotonic in wealth (Figure 5). As wealth rises, the benefits of information about its return increases and the relative cost of information falls—both of which increases a household's information acquisition probability and thus accuracy. 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 effects 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 show how the differences in savings and information depicted in Figures 3 to 5 accumulate to create substantial differences in business-cycle dynamics.

In Table II, we contrast the dynamics of our benchmark economy with those that arise in an economy in which all households exogenously acquire information every period. We refer to this counterfactual economy as the "full-information economy". For comparison, on average, only around 10 percent of households choose to acquire information in any given period in our baseline calibration.[27] We also compare the dynamics from our model with those that arise from an economy in which households face an exogenous 10 percent probability of information acquisition in every period à la Mankiw and Reis (2002). In a slight abuse of terminology, we denote this economy as the "exogenous-information economy".[28] In both cases, we assume households pay the resource and utility cost of information upon acquisition. Section 5.7 compares our results to the case in which households do not pay these costs. Table B.2 in the Appendix further compares the business-cycle dynamics to those in the US.[29]

Relative to the full-information case, fluctuations in all aggregate variables are substantially more pronounced in our benchmark economy. The standard deviation of capital is around 2/3 higher, and output and consumption are, as a result, 5-11 percent more volatile than in the full-information economy. This stark difference is caused by uninformed, medium-wealth households' savings choices—the main driver of capital dynamics in our economy.[30] Dynamically, as discussed in the Section 5.1, households who choose *not* to acquire information have expectations that are tilted towards the long-run average level of capital. In booms, such households therefore systematically underpredict the capital stock, and hence future labor income, and vice versa in recessions.[31] Consequently, medium-wealth households save more in booms and less in recessions than if they had full information (as shown in Figure 4). The wealthy, in contrast, make the opposite savings mistake, as their total wealth is dominated by financial wealth, as opposed to human wealth. But, because these households buy information relatively frequently and mistakes are small (Figure 3 and 5), they do not make substantial savings errors on average. As a result, in equilibrium, the economy systematically "overaccumulates" capital in booms and "underaccumulates" in recessions, leading to much larger fluctuations in output, consumption, and investment than under full-information.

[27] The probabilities of information acquisition in the benchmark model are 0.122 and 0.097 for the unemployed and employed, respectively. The average rate of unemployment is 7.6 percent in our model.

[28] Notice that the full-information economy is merely an exogenous-information economy with the probability of information acquisition equaling 1 in every period. In what follows, we keep the "full-information'" and "exogenous-information" labels to clarify that any comparison between the two focuses on the *incompleteness* of information. A comparison between our benchmark economy and the exogenous-information economy, by contrast, focuses on the consequences of *heterogeneity* in information. We calibrate the discount factor in each case to target the same $K/Y-$ratio as in our benchmark economy.

[29] Compared to the data (Table B.2), our model features the same deficiencies synonymous with workhorse RBC or Krusell and Smith (1998) economies: Investment is too volatile, consumption too smooth, and output is, as a result of employment and investment dynamics, too volatile. We abstract from these deficiencies in what follows below, as our focus is squarely on the consequences of the wealth-expectation nexus.

[30] Around 2/3 of the capital stock is held by households with a wealth-level between \$33,000 and \$250,000.

[31] 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.3).

Table II: Business Cycle Moments

	Panel (a): Level of Moments				
	$\sigma(K)$	$\sigma(Y)$	$\sigma(I)$	$\sigma(C)$	Cor(C, Y)
Benchmark Model	5.05	3.38	11.98	1.25	0.70
Exogenous Information	4.19	3.22	11.56	1.19	0.67
Full Information	3.07	3.04	10.54	1.19	0.72
	Panel (b): Percent Difference w.r.t. Full Information				
	$\sigma(K)$	$\sigma(Y)$	$\sigma(I)$	$\sigma(C)$	Cor(C, Y)
Benchmark Model	64.24	11.03	13.61	5.08	-3.08
Exogenous Information	36.43	5.84	9.63	0.29	-7.24

Note: The table shows the standard deviation σ of the logarithms of economy-wide capital (K), output (Y), investment (I), and consumption (C). In addition, the table shows the correlation between aggregate consumption and output($Cor(I, Y)$). 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 0.122 and 0.097, respectively, for the unemployed and employed. This probability is set equal to 0.10 in the exogenous information case.

Compared to the exogenous-information case, 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 aggregate volatility (Table II). This is because the middle-wealth households that combine to hold most of the capital stock in our model buy information with less than average frequency (Figure 5). This further dampens the mean-reversion of capital in our benchmark economy relative to an economy in which all households have the *same probability* of acquiring information. The heterogeneity of information choices, in other words, here amplifies the consequences of incomplete information.

Finally, Appendix C.1 illustrates the non-linear relationship between aggregate volatility and the accuracy of expectations about capital. To do so, we compare the business-cycle dynamics of the benchmark economy with those from two additional economies, where households *exogenously* learn the current capital stock with a fixed probability every period, in addition to their *endogenous* decision to acquire information about productivity. Our results document strong non-linearities in the effects of incomplete information. In the case in which household learn capital every period—and thus for which only *static effects* of incomplete information about productivity exists—the volatility of capital and output increase by around

5 and 1 percent, respectively, relative to the full-information case. By contrast, if households learn capital every 10 quarters, this difference is more than 2.5 as large—while, in our benchmark model, where households are never exogenously told about the capital stock, volatilities increase by close to a factor of 4. Our results, thus, highlight the strong non-linearities inherent in the business-cycle effects of mean-biased capital expectations.

We conclude that the presence of endogenous, incomplete information serves as an amplifying force—it induces weaker mean-reversion of the capital stock relative to the full-information case. This amplification acts as a powerful, endogenous propagation mechanism, making the economy respond more to shocks than standard models would otherwise predict.[32]

5.5 Distributional Implications

We next turn to the implications of heterogenous information for household inequality.[33] As discussed in Section 5.1, the weakening of the mean-reversion of the capital stock interacts with households' information choices to alter households' savings behavior—and hence also wealth inequality in the economy. Figure 7 and Table III contrast the wealth distribution in our calibrated benchmark—averaged across a long simulation of the model—with that from the equivalent full-information and exogenous-information economies.

The introduction of heterogenous information increases wealth dispersion relative to the full-information case—with more mass placed at the bottom and at the top of the wealth distribution (Panel (b) and (c) in Figure 7). Conversely, there are fewer households with intermediate wealth levels (between $65,000-750,000). Consistent with the widening of the wealth distribution, all measures of inequality (Gini coefficient, the 90/10-ratio as well as the 99/1-ratio) increase modestly (around 2 percent) with the introduction of heterogenous information (Table III). Comparing instead to the exogenous-information case with a fixed probability of information acquisition, the widening of the wealth distribution is also somewhat amplified. Households' endogenous information choices increase the widening of the wealth distribution—especially near the top. Finally, the introduction of heterogeneous information causes inequality to also be more cyclical. We return to dynamics of inequality in Section 5.6.

Decomposing the change in the wealth distribution is challenging, as the distribution is the equilibrium outcome of a model in which the dynamics change when we modify the information structure. To make progress, we decompose the overall change into three separate forces:

[32] 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.

[33] Similar to the baseline Krusell and Smith (1998) model, our benchmark economy under-predicts wealth inequality vis-a-vis the data. We consider an extension that matches the Gini of wealth in Section 5.7

Figure 7: Information and the Wealth Distribution

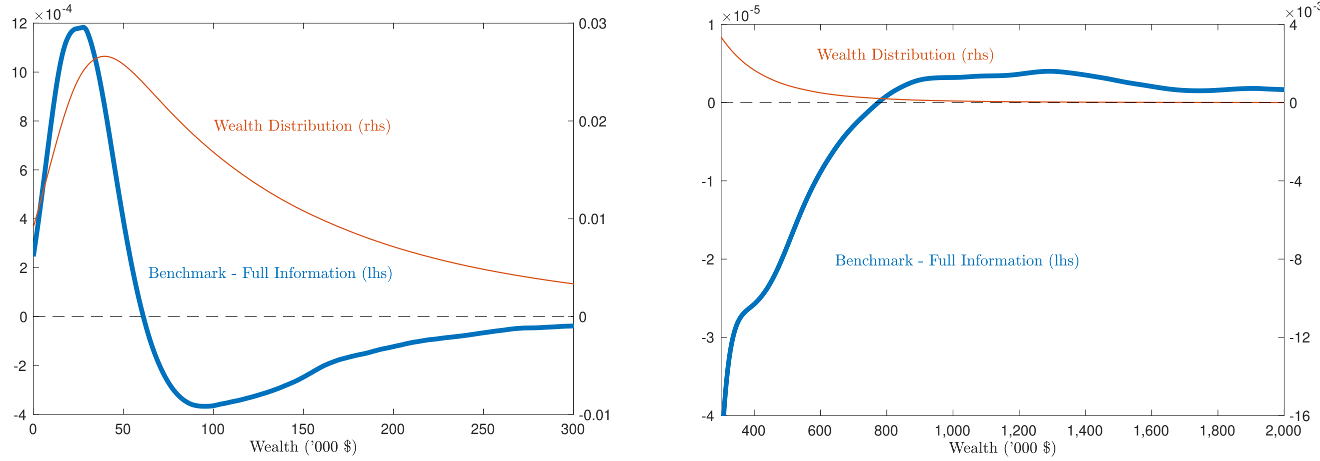

Panel (a): Distribution Differences (I/II)

Panel (b): Distribution Differences (II/II)

Note: The figure showcases differences in the wealth distribution relative to the full-information economy (left-hand side axis) and plots the underlying wealth distribution in the benchmark economy (right-hand side axis). The horizontal axis is household wealth (capital levels) in USD '000 ($). 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.

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

1. General Equilibrium: To isolate the general-equilibrium component, 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. This experiment isolates the channel by which the weakening of the mean-reversion of capital affects the wealth distribution.[34]

Panel (b) in Figure 8 shows that the change in the *equilibrium law of motion for capital* in-and-of itself *dampens inequality*. All else equal, the weaker mean-reversion in the benchmark

[34] To be clear, households still optimize in this experiment and markets clear in every period. Household expectations are simply inconsistent with the true aggregate dynamics.

Table III: Wealth Distribution

	Panel (a): Level of Moments				
	Gini G	90/10	99/1	Cor(K,Y)	Cor(G,Y)
Benchmark Model	0.51	14.53	345.3	0.62	-0.09
Exogenous Information	0.51	14.51	339.4	0.56	-0.07
Full Information	0.50	14.16	328.9	0.49	-0.05
	Panel (b): Percent Difference w.r.t. Full Information				
	Gini G	90/10	99/1	Cor(K,Y)	Cor(G,Y)
Benchmark Model	1.60	2.66	4.98	24.66	86.15
Exogenous Information	1.07	2.51	3.20	14.38	55.79

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 the correlation between the logarithm of capital, the Gini coefficient, and log. 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, on average, 0.122 and 0.097, respectively, for the unemployed and employed. This probability is set equal to 0.10 in the exogenous-information case for all households and for all moments in time.

economy causes a substantially more persistent law of motion for capital, and thus increases the persistence of wages and the rate of return on capital. This, in turn, causes stronger income effects on savings, all else equal. Relative to the full-information economy with its own law of motion, a stronger correlation between returns and savings at the bottom—and a lower correlation at the top—of the distribution thus arises, reducing wealth inequality. Recall that income effects reduce savings for low-to-medium wealth households when the capital stock (and hence wages) are high, while the opposite is the case for wealthy households as returns (and hence financial income) decline with capital.[35] General-equilibrium dynamics—through these mechanisms—partially offset the widening of the wealth distribution.

2. Incomplete Information: Our next experiment isolates the consequences that the *incompleteness of information* itself has on the wealth distribution. We solve for the policy

[35] This mechanism is further amplified by the widening of the difference between the income of employed and unemployed. The cross-sectional dispersion of labor-market income increases with the volatility of wages, and hence with the volatility of the capital stock. In equilibrium, the difference between wages and unemployment benefits equal $(1-\mu)w_t = (1-\mu)(1-\alpha)z_t K_t^{\alpha}\bar{l}^{-\alpha}$, whose average is increasing in the volatility of the capital stock K_t. This, in turn, causes pre-cautionary savings to also increase, especially near the bottom of the distribution, further reducing inequality.

Figure 8: Decomposition of Changes to the Wealth Distribution

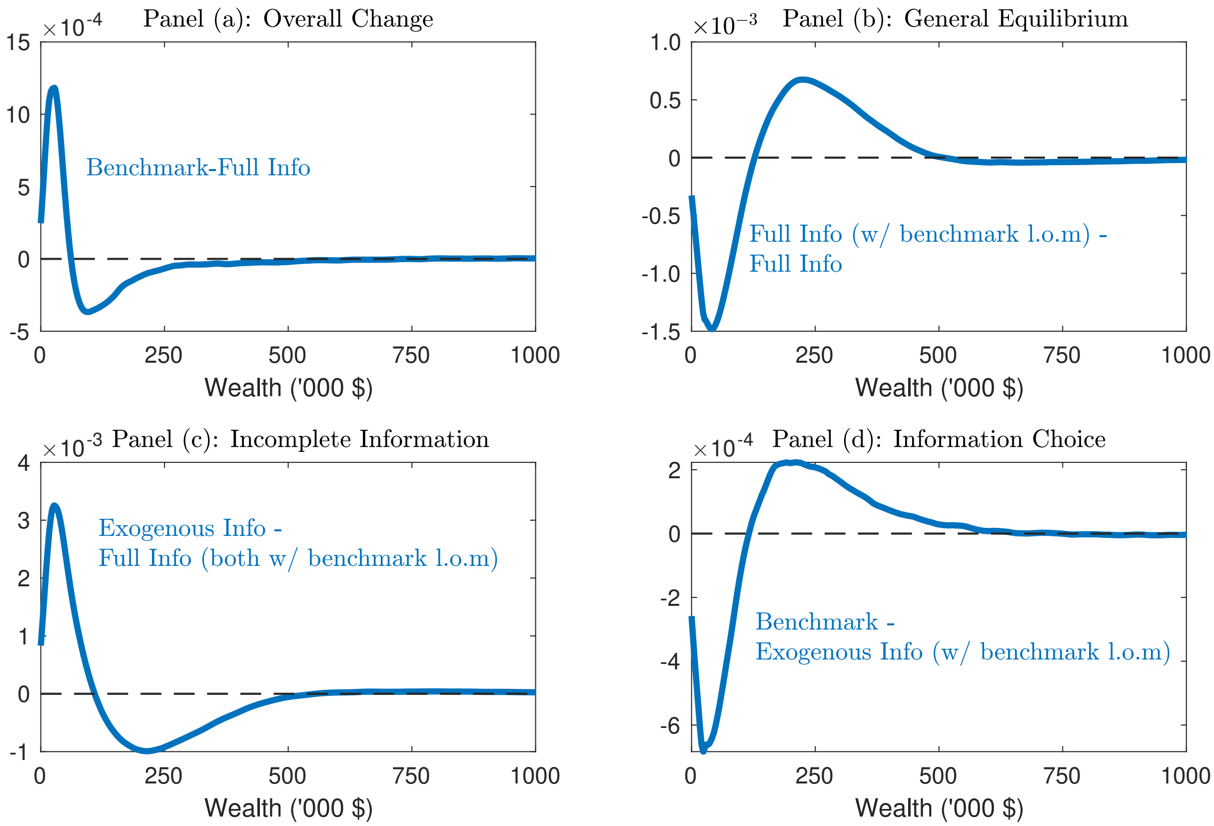

Note: The figure decomposes the change in the average wealth distribution relative to the full-information version of the benchmark economy (Panel (a)). Panel (b) shows the difference in the average probability density function of the wealth distribution between the full-information economy and the full-information economy in which the law of motion for the capital stock equals that in our benchmark economy. Panel (c) shows the difference between the exogenous-information economy and the full-information economy, where we equip both economies with a law of motion for the capital stock equal to that in our benchmark economy. Finally, Panel (d) shows the differences between our benchmark (endogenous-information) economy and exogenous-information economy equipped with the law of motion from our benchmark economy. The horizontal axis in all panels is household wealth (capital levels) in '000 USD ($). 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 capital holdings, using a kernel density estimator with the Epanechnikov kernel.

functions and simulate the distribution when households have exogenously incomplete information but believe 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 can quantify the effects of household incomplete information on inequality. Panel (c) in Figure 8 plots the difference between the two distributions.

The presence of incomplete information *per se* explains the lion's share of the widening of the wealth distribution. The dampened correlation between savings and returns—caused by the incompleteness of information—implies a widening of the wealth distribution that is qualitatively similar but stronger than that in our benchmark economy (Panel (a)).

For low-to-medium wealth households, the presence of incomplete information weakens the positive correlation that exists between returns and savings in response to fluctuations in capital (Figure 4). This, in turn, reduces their average wealth accumulation. In effect, these households make more "mistakes" with their savings choices—as they are unable to effectively exploit periods of high returns on capital—and end up poorer, as a result.

The increased randomness by which households make savings choices, by contrast, increases the share of wealthy households. These households' informed savings policies, as mentioned above, correlate negatively with fluctuations in the return on capital (Figure 4). The presence of incomplete information thus increases the average correlation between their savings and returns, increasing the average wealth of high-wealth households.

In essence, high-wealth households in our model—who also feature close to linear policy functions, as they are far away from the borrowing constraint—are akin to the households described in Piketty and Saez (2003): The authors show that the combination of exogenous random savings rates and linear policy functions can generate Pareto tails in the wealth distribution. In our case, however, the presence of 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.[36]

3. 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 effects that *heterogeneous information choices* have on the wealth distribution. We plot differences in Panel (d) in Figure 8.

All else equal, the more informed households in our benchmark economy are the poor, unemployed households and the rich households with substantial amounts of wealth (Figure 5). These households, who are better informed than their exogenous-information counterparts,

[36] That said, because the magnitude of the errors are comparatively small relative to what is needed, in the benchmark model this force is not sufficiently potent to generate a thick Pareto tail of wealth (cf. Section 5.7).

make fewer savings mistakes—they are, on average, better able to exploit differences in rates of return on capital. As a result, the existence of *heterogenous information* causes fewer wealth-poor ($0-50,000) as well as wealth-rich (above $750,000) households. The decline in "random savings" caused by the presence of additional information allows households near the bottom-end of the wealth distribution to better correlate their savings with its rate of return, allowing these households to better accumulate wealth. By contrast, the decline in "random savings" of the wealthy drives down their realized saving.

Summary and Quantification: 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 leading to "random savings behavior". On the other hand, the existence of heterogeneity in information choices, all else equal, dampens this increase by allowing the more exposed households—the poor and the wealthy—to acquire information to make better savings choices. This dampening effect is then further amplified by the weakening of the mean-reversion of capital further lessening any increase in inequality.[37] On balance, we find that the former effect dominates, leading to a modest overall increase in inequality (Table III). Table C.2 in the Appendix decomposes the overall increase in different summary measures of inequality (e.g., the Gini) and shows that *modest net increase* in inequality is comprised of *large, offsetting gross effects.* The 2 percent increase in the Gini is, for example, comprised of an 8 percent increase due to incomplete information counteracted by 4 and 2 percent decreases, respectively, due to general equilibrium dynamics and heterogeneity in information choices. The combined effects of heterogenous information are subtle—and, crucially, affect different parts of the distribution differentially.

5.6 Dynamics of the Wealth Distribution

The previous subsection studied the average consequences of heterogenous information for household inequality. The final columns of Table III, nevertheless, also shows that its introduction causes inequality to become substantially more cyclical. In fact, although the overall correlation is small, the correlation between the Gini coefficient and output almost doubles in our benchmark economy compared to that under full information (Table III). Busts (booms) become periods with relatively more (less) inequality. Indeed, as a result, allowing for the presence of heterogenous information brings the dynamics of household inequality closer to what is observed for the US economy. The correlation between the Gini coefficient and out-

[37]Notice that while heterogeneity in information acquisition by itself decreases inequality, our benchmark economy features more inequality than its exogenous information counterpart (Table III). This is due to the weaker general equilibrium dynamics in our benchmark model with heterogeneous information.

Table IV: Dynamics of the Wealth Distribution

	σ(Gini)	σ(90/50)	σ(99/50)	Cor(90th,10th)	Cor(99th,10th)
Benchmark Model	0.04	0.63	2.68	0.24	-0.24
Exogenous Information	0.03	0.42	1.93	0.40	-0.20
Full Information	0.01	0.15	0.57	0.72	0.26

Note: Based on a long simulation of the benchmark economy and the versions with exogenous and full information, the table shows the standard deviations (σ) over time of the Gini, the 90/50, and the 99/50 percentile ratios of the wealth distribution (k), as well as the correlations (Cor) of, respectively, the 90th and 99th percentiles with the 10th percentile of the wealth distribution.

put increases from -0.05 under full information to -0.09 in the benchmark economy. For comparison, the correlation between the two for the US economy is -0.12.[38]

Table IV identifies the forces that make inequality more countercyclical. It does so by comparing the dynamics of different parts of the wealth distribution in the benchmark economy to its full- and exogenous-information counterparts. The main feature that stands out are the substantially larger swings that occur in the benchmark economy. Relative to the case with full information, the standard deviations of both the Gini coefficient and, for example, the 99/50 percentile ratio are about 4–4.5 times higher in the benchmark economy—with inequality moments that are, further, substantially larger than in the exogenous-information case.

The key to understanding this increase in the volatility of wealth dispersion lies, once more, in the impact that incomple information has on savings choices across the wealth distribution. With full information, low-to-medium wealth households save less in booms, making their wealth accumulation *countercyclical*. The rich, by contrast, save more when capital is high and returns low, but their overall wealth accumulation remains *countercyclical*—driven by countercyclical returns to their stock of wealth. Taken together, the two imply that the whole wealth distribution moves up and down in tandem, limiting variations in inequality.

The presence of incomplete information, by contrast, makes savings, and thus wealth, at the bottom of the distribution *less* countercyclical, while *increasing* the countercyclicality at the top of the wealth distribution. Both arise due to an increase in savings errors that households make because of the presence of incomplete information. Combined, the two explain the increased volatility of inequality measures in Table IV, and the increased cyclicality

[38]The data counterpart is calculated based on annual data for the US Gini in household networth (obtained from the World Inequality Database as variable GHWEAL992J) and annual data on US Real GDP (obtained from FRED as variable GDPCA) for the years 1962-2023. Both annual series are detrended using an HP filter with smoothing parameter of 6.25 (Ravn and Uhlig (2002)).

of inequality. Finally, notice that—relative to the economies with *exogenous information*—the endogeneity of information in our benchmark economy dampens the cyclicality of savings by medium-wealth households (who have less-than-average information and own the bulk of the capital stock) further. This, in turn, causes the volatility of wealth inequality to be larger in our benchmark economy than in its exogenous-information counterpart.[39]

5.7 Discussions and Extensions

Before closing with two policy exercises that target the wealth-expectation nexus, we briefly turn to several extensions of our benchmark framework. The overall aim of these is to further elucidate the forces that drive our results. We recalibrate all alternative models to target the same capital-output ratio and (where appropriate) average errors as our benchmark model.

Angels and Demons: The frameworks considered so far with a given probability of information acquisition keep the costs of information acquisition from the benchmark model unchanged. Thus, while households have no choice about information acquisition in these alternative models, they still have to pay the resource and utility cost of information upon becoming informed. This focuses any comparison between the different models on the *endogeneity* and *incompleteness* of information itself—rather than the size of costs associated with it. We relax this assumption in Appendix C.4, where we instead make information freely (but randomly) available. We note that the exogenous-information economy without information costs is similar to the "Mankiw-Reis"-style economy studied in Auclert *et al.* (2020) and Carroll *et al.* (2020). Table C.5 and C.6 show that key moments change little irrespective of whether the given information is from a "Calvo-Angel" or a "Calvo-Demon".[40] Business cycle moments are next to unaffected: The maximum difference is less than a quarter of one percent (Table C.5). Inequality measures are slightly more affected, because the absence of a resource costs affects the disposable income of the poor (Table C.6). Nevertheless, the maximum difference in inequality measures is still below 1 percent. We conclude that assumptions about the information costs themselves matter little for the above comparisons between our benchmark economy and the alternative informational environments.

[39] Because unemployment is countercyclical, individual earnings volatility is generally higher in busts than in booms. As we mention above, this increases the value of information for unemployed households, causing their information acquisitions to be weakly countercyclical—more information is bought by unemployed households in a bust. Overall, however, information purchases are weakly procyclical, driven by the behavior of the employed. This modulates the above discussed effects.

[40] An anonymous referee provided the labeling for these alternative economies: When we make information freely available in the alternative frameworks with an (exogenously) pre-specified probability of information acquisition it is as-if households are touched by a "Calvo-Angel" upon becoming informed. By contrast, in the alternative frameworks that we studied above, where households still have to pay the costs of information upon becoming informed, it is as-if households are instead touched by a "Calvo-Demon".

Homotheticity of Information Costs: Section 5.3 shows that our benchmark model features a non-monotone relationship between the accuracy of unemployment expectations and wealth. Above, we have documented how households' complex incentives to acquire information, depending on their prior beliefs and wealth levels, lead to such a relationship. Appendix C.5 explores the consequences that the non-homotheticity of the two different information costs—the random utility costs κ and the real resource cost η—have for these results.

Notice that the utility cost in (3.1) affects a household's *level of utility*. Thus, any given κ shock can have different consequences for a household whose level of wealth—and hence utility—is low compared to high, because the value of information itself depends on the level of wealth. Figure C.1 compares the inverse u-shape from the benchmark model with one from a model in which the utility-cost shock is multiplicative in the utility value of cash-at-hand. The latter restores homotheticity. The relationship between the accuracy of unemployment expectations and wealth is nearly identical to that in the benchmark model—with only a small increase in the size of the inverse u-shape noticeable. Consistent with these findings, Table C.7 shows that the business-cycle and inequality statistics also do not depend much on the non-homotheticity of the utility cost. The maximum difference between our benchmark model and the moments of the model with the scaled utility cost is less than 2 percent.

By contrast, the presence of the small resource cost—recall that the calibrated resource cost is around 1/10 of a percent of average quarterly earnings—and the associated non-homotheticity matters more. Figure C.2 compares the relationship between the accuracy of unemployment forecasts and wealth in the benchmark model with a version of the model without resource costs of information. The absence of a resource cost causes a negative relationship between accuracy and wealth, unlike that in the data or in the benchmark model. Although the information policy function still shows non-monotinicities similar to those in benchmark model, the resource cost presents a barrier for wealth-poor households' information acquisition. Table C.8 further shows that the amplification and increase in inequality documented above is also somewhat dampened in the model without a resource cost of information. We conclude that the non-homotheticity embedded in the resource cost of information is an important feature for matching the survey data—and for the consequences of heterogenous information highlighted above—unlike the non-homotheticity implied by the utility cost.

The Cyclicality of Taxes: Following Krusell and Smith (1998) and others, our baseline framework assumes a balanced government budget in every period (Section 3.3). Because the labor endowment falls and government outlays for unemployment insurance rise in recessions, a balanced budget, however, implies *countercyclical tax rates*: Labor-income taxes rise in recessions and fall in booms. This makes household income, all else equal, more pro-cyclical, affecting savings and information-acquisition policies as a result.

To analyze the importance of countercyclical taxes for our results, Appendix C.3 studies an alternative version of our model in which the government perfectly smoothes taxes through an implicit insurance scheme with (unmodelled) financial intermediaries from the rest of the world. This follows the approach taken in Mitman and Rabinovich (2015), among others. In particular, the government charges a constant tax rate $\tau^{\star}$ such that the present value of labor-income taxes equals that of unemployment-benefit payments when appropriately discounted.[41] Table C.3 and C.4 compare key business-cycle and inequality moments across the familiar three versions of our benchmark model, using this specification.

As expected, aggregate volatility is lower with acyclical taxes. However, crucially, relative to the full-information case, the increase in aggregate volatility that arises due to heterogeneous information is similar to that in our benchmark economy, although somewhat smaller in relative magnitude. The standard deviation of output, for example, now increases by 8 percent compared to 11 percent previously. While the equilibrium level of inequality is similarly slightly smaller, the increase in inequality that arises due to heterogenous information is now somewhat larger—with a similar decomposition of the overall increase to that from before. The 99/1-ratio, for example, now increases by 7 percent versus 5 percent previously. We conclude that our main business-cycle and inequality implications are robust to changes in assumptions about the cyclicality of income taxes.

Matching the Wealth Distribution: A prominent issue with the Aiyagari-Bewley-Huggett-Imrohoroglu class of models that we depart from is that it is does not generate realistic wealth heterogeneity: The data display significantly more skewness in wealth than the models. Although the presence of stochastic death, in our case, helps match the lower-tail of the wealth distribution, the concentration of wealth among the richest is similarly too small in our benchmark framework. The Gini coefficient is, as a consequence, too low: On the basis of data from the 2022 SCF, the Gini coefficient on wealth in the US hovers around 0.80. In the benchmark model, by contrast, the equivalent number is only 0.51. The presence of heterogeneous information, as argued above, modestly increases wealth concentration among the rich—and thus the Gini coefficient—but does not do so by enough to match the data.

One of the main purposes of our line of research is to extend workhorse heterogenous-agent frameworks to simultaneously allow for heterogeneity in wealth and expectation formation. It therefore seems important to make sure that the forces we highlight extend to models that roughly match the observed wealth distribution. To analyze such an environment, Appendix

[41] The exact formula for the constant tax rate is $\tau = \mathbb{P}[z = z_l]\frac{\mu u_l}{u_l\mu+(1-u_l)} + \mathbb{P}[z = z_h]\frac{\mu u_h}{u_h\mu+(1-u_h)}$, where $\mathbb{P}[z = z_l]$and $\mathbb{P}[z = z_h]$ are the unconditional probabilities of a bust and a boom, respectively.

C.6 follows Bayer *et al.* (2024) and extends our baseline framework with a CES-structure for final production and a share of entrepreneurs in the population, who do not work but earn all pure rents in the economy. As a baseline, we set the share of entrepreneurs to 1/2 percent and the elasticity of substitution to target a 5 percent profit share. We follow Bayer *et al.* (2024) and assume a fixed transition-probability matrix allowing households to switch in and out of the entrepreneur state. The transition-probability matrix is set to match the estimated values in Guvenen *et al.* (2014), following the approach in Bayer *et al.* (2024). We recalibrate the extended model to target the same forecast error moments as our benchmark framework. Table C.9 shows that the extended model matches closely the Gini on wealth in the data.[42]

Figure C.3 compares the accuracy of unemployment forecasts across the wealth distribution with that in the benchmark model. Although the inverse u-shaped relationship is slightly shallower compared to that in the benchmark model, the overall strength of the relationship is similar to before. The main difference being that the addition of wealthy entrepreneurs makes the top percentiles of the wealth distribution more accurate than in our benchmark model. Consistent with this similarity, Table C.9 shows that the presence of heterogenous, incomplete information once more amplifies business-cycle fluctuations: The standard deviation of capital, for example, increases by 28 percent, although the effects are now clearly somewhat muted by the presence of a fixed share of wealthy-entrepreneurs households who inevitable acquire information. Finally, Table C.9 shows that, as in our benchmark framework, the presence of incomplete information once more widens the wealth distribution, especially the right tail, while heterogeneity in information by itself dampens this increase by allowing more exposed households to acquire information at higher rates. The presence of heterogenous, incomplete information, as such, causes e.g., the 99/1-ratio to increase by 6 percent.[43]

Overall, we conclude that the consequences of heterogenous, incomplete information that we highlighted above extend to a framework which better matches the wealth distribution. Clearly, the addition of another source of heterogeneity—whether households are workers or entrepreneurs—further complicates the already intricate consequences of heterogenous, incomplete information. Yet, the overall consequences that we find are similar and akin to those in our benchmark framework.

Across all extensions, a consistent pattern emerges: Heterogeneous information choices amplify business-cycle volatility and alter inequality dynamics. The specific magnitudes

[42] While the model matches the Gini coefficient on wealth, it is, however, not able to match the Pareto tail of wealth in the data, as the wealth distribution merely inherits the tail of the entrepreneurial earnings (Benhabib and Bisin, 2018).

[43] One noteworthy difference between our benchmark framework and the entrepreneur extension is that, while in the former case general-equilibrium dynamics dampened the overall increase inequality, in the latter case it contributes to the widening of the wealth distribution (Table C.9). This once more shows that the details of which households buy information matter for inequality.

vary—information costs matter more for matching empirical expectation patterns, while the entrepreneur extension demonstrates robustness under realistic inequality—but the core mechanisms persist. Whether households face cyclical or constant taxes, pay utility or resource costs for information, or operate in economies with concentrated or dispersed wealth, the fundamental two-way relationship between wealth and information systematically alters macroeconomic dynamics. This robustness strengthens confidence that our conclusions reflect genuine features of the wealth-expectation nexus rather than artifacts of particular modeling choices.

The results in this section further point to a broader conclusion. Because the presence of incomplete information alters households' economic choices, heterogeneity in information choice 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 and heterogeneity. The next section shows how the interplay between information choice, (pre-cautionary) savings, and the aggregate economy similarly modifies the predictions of simple economic policies that inadvertently target the expectation-wealth nexus.

6 Two Policy Experiments

The previous section highlighted the interaction between heterogenous information choices and the aggregate economy. We close the paper by illustrating how two policy reforms that directly modulate the expectation-wealth nexus can change the properties and dynamics of the economy in rich and unexpected ways. We do so for both the benchmark economy and the extended entrepreneur economy with higher wealth inequality. Our aim in this section is to illustrate that macroeconomic policies may have important additional consequences in environments with heterogenous information. By changing the distribution of households' information, and hence their expectations, macroeconomic policies fundamentally alter an economy's responsiveness to shocks, and hence the desirability of economic policies.

6.1 A Wealth Tax

For our first 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 US Congress in 2021 (e.g., Guvenen *et al.*, 2019; Saez and Zucman, 2022).[44] 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. The proceeds from the tax are rebated to households as lump-sum

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

Table V: Quantitative Effects of a Wealth Tax

	K	$\sigma(Y)$	Gini K	90/10	99/1	Info U.	Info E.
Benchmark Model	-5.56	6.35	-0.17	-2.63	-0.92	-6.83	-8.53
Exogenous Information	-5.35	3.00	-1.23	-3.21	-4.88	.	.
Full Information	-5.02	0.66	-3.06	-4.55	-10.26	.	.

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 in percent. K denotes the mean of the (log-)capital stock, while $\sigma(Y)$ denotes the variance of (log-)economy-wide output. The table computes the moments for the both calibrated model ("Benchmark Model") and the associated full-information and exogenous-information economies ("Full Information" and "Exogenous Information", respectively). 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.

transfers.[45] Household cash-at-hand m_i is therefore given by the expression:

$$m_i = (1 + r - \delta - \tau_k)k_i + (1 - \tau)\left[\epsilon_i w\bar{l} + (1 - \epsilon_i)\,\mu w\right] + T. \tag{6.1}$$

where the transfer T equals the average proceeds of the tax. We model the size of the wealth tax based on the wealth tax that was in effect in France from 2011 to 2017. The magnitude of the tax is also consistent with the recent proposal in the US Congress. In particular, we set $\tau_k = 0.0025$, corresponding to a one percent per annum wealth tax. Table V and Figure D.1 report the macroeconomic consequences of the wealth tax.[46]

The direct effect of the wealth tax is—unsurprisingly—to reduce aggregate savings, as seen by the 5 percent drop in the capital stock (Table V). For a given level of wealth and information, the wealth tax reduces households' incentive to save. However, for a given wealth level, information acquisition policies are approximately unaffected by the tax. Because the probability of information acquisition is heterogeneous across the wealth distribution, the wealth tax therefore also changes the level and incidence of information in the economy.

The lump-sum transfer, financed by the wealth tax, increases cash-at-hand for wealth-poor households, while wealth-rich see their resources reduced by the tax. Because both of these groups acquirer information at higher rates than average, the introduction of the wealth tax

[45] We have experimented with an alternative setup in which the proceeds are not rebated to households but instead treated as unvalued government spending. This alternative specification only captures part of the effects that we present in detail in this section—it only effects the right tail of the distribution since without the rebate there is no direct effect on the wealth-poor. The results from this specification (unreported) are in line with the detailed results we present here.

[46] 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.

strongly reduces information acquisition—by more than 8 percent, on average. This reduction, in turn, decreases the accuracy of household expectations, dampens the mean-reversion of capital, and thus leads to more volatility in output (by c. 6 percent).

The introduction of the wealth tax further changes the wealth distribution. Indeed, perhaps surprisingly, the Gini coefficient in our benchmark economy, summarizing wealth inequality, is almost unchanged by the introduction of the wealth tax. We plot the change in the wealth distribution following the wealth tax in Figure D.1 in the Appendix. There are three contrasting forces that explain the effects of the wealth tax on inequality.

First, the lump-sum transfer financed by the tax, all else equal, reduces income volatility, and thus households' precautionary wealth. Second, this left-ward shift in the wealth distribution is, in turn, reinforced by the "disincentive effect" from a lower after-tax return on capital, which further reduces savings for all households, including the rich. Combined, these two effects reduce wealth inequality. In our benchmark economy, however, there is also a third offsetting force that operates through heterogeneous information choices.

Because information increases with wealth for wealth-rich households, the tax-induced reduction in their wealth also decreases their information acquisition. This decrease in information makes 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 an increase in (the approximately) "random savings". Similar to the results in Piketty and Saez (2003), and for the same reasons as discussed in Section 5.4, this dampens the reduction of top-wealth holdings. The mass of households with more than \$1 million in wealth in Figure D.1 is, as a result, almost unchanged—despite the wealth tax. Thus, counterintuitively, introducing a wealth tax does not *substantially* decrease the share of extreme wealthy in the economy.

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 V, the direct impact of the tax on average wealth (i.e., the capital stock) is nearly identical; but that is where the similarities end. The volatility of output only increases marginally, in contrast to the stronger increase in the benchmark economy. The dampened mean-reversion of capital is absent in the full-information case, leading to little change in the endogenous volatility of output through capital accumulation.

Turning to the distributional implications of the tax, the differences 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 strongly. This is in contrast to the benchmark case, where the inequality between the top and bottom of the wealth distribution, as captured by the 99/1 percentile ratio is approximately unchanged.

Our results, thus, provide one example of how the general-equilibrium forces that arise from heterogeneous information choices can counteract the direct effects of a policy reform.

Finally, we study the introduction of a wealth tax in the extended entrepreneur model, which better matches the empirical wealth distribution. The results, presented in Appendix C.6, show that the effect of a wealth tax on aggregate volatility is smaller, and more alike across different informational environments, than before. However, crucially, the presence of incomplete information once more strongly mutes the equalizing effects of a wealth tax. Indeed, in the extended model, the Gini coefficient increases under heterogenous, incomplete information, albeit only by around 0.5 percent. The small share of rich entrepreneurs makes the right-hand tail of the wealth distribution less sensitive to economic policies, driving up the 99/1-percentile ratio, and hence the Gini coefficient.[47] Through its effects on "random savings", introducing a wealth tax thus modestly increases inequality in the extended model.[48]

6.2 An Increase in Unemployment Benefits

We close the paper with a second policy counterfactual that instead considers an increase in the replacement rate μ; the unemployed are the second category of households that we identified as acquiring information at higher rates than the average (Section 5). In particular, we consider an increase in μ by 10 percentage points, from 40 to 50 percent of the current wage. This increase, combined with the increase in labor-income taxes required to finance it, reduces the income risk from unemployment, and thus both the incentive to accumulate precautionary savings and to acquire information about the current state.

Table VI shows how the increase in unemployment benefits affects key moments of our benchmark economy. When the replacement rate is increased, the unemployed especially reduce both their savings and their rate of information acquisition, the latter by about 13 percent.[49] As in the case of the wealth tax, this reduction in information acquisition probabilities once more decreases the accuracy of expectations, dampens the mean-reversion of capital, and increases the volatility of output (by c. 4 percent).

Inequality is further affected. The savings and information behavior of the wealth-rich—for whom post-tax labor or replacement incomes account for a small share of total wealth—is

[47] Notice that the 90/10-ratio, however, falls, as in our benchmark model. The increase in the Gini coefficient is explained by its sensitivity to the extreme right-tail of the wealth distribution (Cowell and Flachaire, 2002).

[48] Although our model is not meant to capture all the dimensions of a wealth tax, we note that several studies, often building on data from Swiss Cantons (e.g., Marti *et al.*, 2023) or Norwegian regions (e.g., Iacono and Smedsvik, 2024), show that a *decrease* in wealth taxes can lead to an *increase* in *measured inequality*. However, as argued in Seim (2017), among others, the lion's share of the estimated effect often reflect changes in "evasion and avoidance rather than changes in actual savings behavior" (Seim, 2017).

[49] Since there are fewer unemployed than employed households, the average probability of information acquisition across all households falls by almost 3 percentage points.

Table VI: Quantitative Effects of Increased Unemployment Benefits

	K	$\sigma(Y)$	Gini K	90/10	99/1	Info U.	Info E.
Benchmark Model	-1.04	3.52	5.66	27.40	48.31	-12.63	-1.77
Exogenous Information	-0.84	1.06	4.64	25.08	41.49	.	.
Full Information	-0.61	0.36	3.12	22.73	33.54	.	.

Note: The table shows the effects of a ten percentage-point increase in the replacement rate on moments of the average cross-sectional capital distribution and the logarithm of economy-wide output in percent. K denotes the mean of the (log-)capital stock, while $\sigma(Y)$ denotes the variance of (log-)economy-wide output. The table computes the moments for the both benchmark economy ("Benchmark Model") and the associated full-information and exogenous-information economies ("Full Information" and "Exogenous Information", respectively). 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.

nearly unaffected by the direct effect of the policy. Yet, because the reduced precautionary savings by the wealth-poor increases returns to capital, it increases the average savings of the rich. As the reduction in information acquisition is concentrated among the wealth-poor unemployed, their reduction in savings is further re-inforced by their inability to save when future returns are high and wages low. This explains why the rise in inequality documented in Table VI and Figure D.2 is substantially more pronounced in the benchmark economy compared to the economies with full information and (to a lesser extent) with exogenous information. The same further holds true in the extended environment of Section 5.7, which more closely matches the wealth distribution in the data (Appendix C.6). An increase in unemployment benefits increases inequality by *substantially more* when one accounts for households' endogenous information choices. A policy designed to strengthen the safety net for the unemployed, in essence, inadvertently reduces their incentive to be informed, leading them to make larger savings errors that ultimately increase overall inequality.

While we in the above have abstained from making welfare statements about the desirability of the policies that we analyze, our positive findings indicate that policymakers should proceed cautiously when evaluating the consequences of tax and transfer policies. The effects on households' information choices may lead to implications which run counter to the stated objectives—for example, their effect on wealth inequality. For both the case of a wealth tax and an increase in unemployment benefits, inequality is larger after accounting for the endogeneity of household information choices. More generally, the above policy experiments illustrate 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 macroeconomics 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 have 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 have demonstrated how the wealth-expectation nexus increases the endogenous propagation of shocks and partially accounts for the lack of inequality in standard frameworks with incomplete markets. We have showed how the wealth-expectation nexus further fundamentally alters the predictions of government policies such as wealth taxes or unemployment benefits—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 experiments provide a "Lucas-style" criticism (Lucas, 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 consequences for macroeconomics dynamics and the cross-section.[50] 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 quantities. Does this mean policymakers should subsidize information or alter the manner in which they condition their policy instruments? Should such subsidies or policy instruments target a particular subset of the population? And how close do market-based or institutional solutions, which collect information on behalf of households and distribute it to them, come to the constrained efficient allocation? We leave these exciting questions for future research.

[50] In this sense, our results provide a Lucas-style criticism (Lucas, 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, 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.

Angeletos, G.-M., Huo, Z. and Sastry, K. A. (2021). Imperfect macroeconomic expectations: Evidence and theory. *NBER Macroeconomics Annual*, **35**, 1–86.

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*.

Baley, I. and Veldkamp, L. L. (2025). *The Data Economy: Tools and Applications*. Princeton University Press.

Bayer, C., Born, B. and Luetticke, R. (2024). Shocks, frictions, and inequality in us business cycles. *American Economic Review*, **114** (5), 1211–1247.

Benhabib, J. and Bisin, A. (2018). Skewed wealth distributions: Theory and empirics. *Journal of Economic Literature*, **56** (4), 1261–1291.

Bhandari, A., Borovička, J. and Ho, P. (2025). Survey data and subjective beliefs in business cycle models. *Review of Economic Studies*, **92** (3), 1375–1437.

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*.

Brunnermeier, M. K. and Parker, J. A. (2005). Optimal expectations. *American Economic Review*, **95** (4), 1092–1118.

Caplin, A. and Leahy, J. V. (2019). Wishful thinking. *National Bureau of Economic Research*.

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

— (2006). The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics letters*, **91** (3), 312–320.

—, 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). Dynamic stochastic general equilibrium models for monetary policy analysis. **3**, 285–367.

Coibion, O. and Gorodnichenko, Y. (2012). What can survey forecasts tell us about information rigidities? *Journal of Political Economy*, **120** (1), 116–159.

— and — (2015). Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review*, **105** (8), 2644–2678.

—, — 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**.

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.

—, Kaplan, G. and Song, J. (2014). How risky are recessions for top earners? *American Economic Review*, **104** (5), 148–153.

Harrington, M. (1997). *The other america*. Simon and Schuster.

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

Iacono, R. and Smedsvik, B. (2024). Behavioral responses to wealth taxation: evidence from a norwegian reform.

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.

Kacperczyk, M., Nosal, J. and Stevens, L. (2019). Investor sophistication and capital income inequality. *Journal of Monetary Economics*, **107**, 18–31.

Kaplan, G. and Violante, G. L. (2018). Microeconomic heterogeneity and macroeconomic shocks. *Journal of Economic Perspectives*, **32** (3), 167–194.

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).

Lorenzoni, G. (2009). A theory of demand shocks. *American economic review*, **99** (5), 2050–2084.

Lucas, R. E. (1972). Expectations and the neutrality of money. *Journal of economic theory*, **4** (2), 103–124.

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

— (1976). Econometric policy evaluation: a critique,. *The Phillips Curve and Labour Markets, Carnegie-Rochester Conference Series on Public Policy*, **1**.

Lusardi, A. and Mitchell, O. S. (2014). The economic importance of financial literacy: Theory and evidence. *American Economic Journal: Journal of Economic Literature*, **52** (1), 5–44.

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.*

— and Wiederholt, M. (2009). Optimal sticky prices under rational inattention. *American Economic Review*, **99** (3), 769–803.

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.

Marti, S., Martínez, I. Z. and Scheuer, F. (2023). Does a progressive wealth tax reduce top wealth inequality? evidence from switzerland. *Oxford Review of Economic Policy*, **39** (3), 513–529.

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

Mihet, R. (2022). Who benefits from innovations in financial technology? *University of Lausanne mimeo.*

Mitman, K. and Rabinovich, S. (2015). Optimal unemployment insurance in an equilibrium business-cycle model. *Journal of Monetary Economics*, **71**, 99–118.

Moll, B. (2024). The trouble with rational expectations in heterogeneous agent models: A challenge for macroeconomics. *London School of Economics, mimeo, available at https://benjaminmoll. com.*

Morduch, J. and Schneider, R. (2017). *The financial diaries: How American families cope in a world of uncertainty*. Princeton University Press.

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

Newman, K. S. (2009). *No shame in my game: The working poor in the inner city*. Vintage.

Peress, J. (2004). Wealth, information acquisition, and portfolio choice. *The Review of Financial Studies*, **17** (3), 879–914.

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

Ravn, M. O. and Uhlig, H. (2002). On adjusting the hodrick-prescott filter for the frequency of observations. *Review of economics and statistics*, **84** (2), 371–376.

Reiche, L. (2025). That's what she said: an empirical investigation on the gender gap in inflation expectations. *University of Oxford mimeo.*

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

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

Seim, D. (2017). Behavioral responses to wealth taxes: Evidence from sweden. *American Economic Journal: Economic Policy*, **9** (4), 395–421.

Shipler, D. K. (2005). *The working poor: Invisible in America*. Vintage.

Silvapulle, M. J. and Sen, P. K. (2005). *Constrained statistical inference: Inequality, order and shape restrictions*. John Wiley & Sons.

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*.

Woodford, M. (2002). *Imperfect common knowledge and the effects of monetary policy*. Tech. rep., National Bureau of Economic Research.

Online Appendix for
"Expectation and Wealth Heterogeneity in the Macroeconomy"

Tobias Broer Alexandre Kohlhas Kurt Mitman Kathrin Schlafmann

September 2025

Contents

A Motivating Evidence

A.1 Additional Estimates

Table A.7: 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)	–	–	–
Male	–	-0.004	-0.004
		(0.008)	(0.008)
Education	–	-0.072***	-0.072***
		(0.010)	(0.010)
Non-participation	–	0.019*	0.020*
		(0.011)	(0.011)
Age	–	0.013***	0.013***
		(0.002)	(0.002)
Age^2	–	-0.0001***	-0.0001***
		(0.00002)	(0.00002)
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 (college or not), labor market status, and sex of the respondent, as well as time fixed effects. Column (3) shows estimates "pre-covid"; that is, before January 2020. Robust standard errors in parentheses. Sample: 2013M10-2020M3. $*$ p<.1, $**$ p<.05, $***$ p<.01

Table A.8: Test for the Monotonicity of Regression Coefficients

	Unemployment Forecasts	
	Left-hand side	Right-hand side
Figure 1 (Panel a and b)	0.01	0.32
Figure 1 (Panel c and d)	$<$0.01	$<$0.01
	Inflation Forecasts	
	Panel a (lhs)	Panel a (rhs)
Figure 2	0.89	0.01
	House Price Forecasts	
	Panel a (lhs)	Panel a (rhs)
Figure 2	0.09	0.03

Note: The table reports p-values from the Likelihood Ratio test by Silvapulle and Sen (2005) of monotonically declining regression coefficients (H_0 : monotonic decline) against the alternative (H_A : non-monotinicty) using the SCE data discussed in Section 2. Results are reported using 10,000 draws from a semi-nonparametric Bollen-Stine bootstrap procedure. The figure and panel labels correspond to those used in the main text.

Table A.9: Unemployment Expectations and Wealth Percentiles

	Dependent variable:
	Absolute Error
Wealth Percentile	0.113*
	(0.059)
Wealth Percentile Squared	$-$0.162**
	(0.067)
Constant	1.237***
	(0.015)
Controls	✓
Observations	40,998
F Statistic	6.876*** (df = 2; 40995)

Note: The table shows estimates from a regression of the absolute value of individual unemployment errors on the wealth percentile that the individual respondent belongs to, in addition to controls: the age, education (college or not), labor market status, and sex of the respondent, as well as time fixed effects. Robust standard errors in parentheses. Sample: 2013M10-2020M3. $*$ p<.1, $**$ p<.05, $***$ p<.01

A.2 Data Construction

The SCE is a monthly internet survey of c. 1,300 "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.[51]

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 difference between individual best estimates and the actual (12-month-ahead) outcomes of US consumer price index inflation and inflation in 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 instead 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 $P_t(u_{t+12} > u_t)$.

[51] We match wealth observations with the household's monthly expectations using the household's id variable.

Consistent with ample evidence that professional forecasters provide more accurate predictions than even those from modern statistical and economic models (Stark *et al.*, 2010; Faust and Wright, 2013; and Bhandari *et al.*, 2025), we proxy the true probability by the consensus forecast from the SPF, which we denote $P_{SPF,t}\left(u_{t+12} > u_t\right)$. 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:

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

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.[52] 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.10 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 substan-

[52] 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.10: 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).

tially larger for households than for professional forecasters—equal to 1.6 and 0.7 percentage points (pp), respectively. Furthermore, Table A.10 demonstrates that household expectations are also 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, respectively. Finally, Table A.10 shows that household expectations are also substantially *more heterogeneous* than SPF forecasts. Specifically, household unemployment expectations and point forecasts for CPI inflation have a much higher cross-sectional standard deviation than the forecasts of professionals. The standard deviation of forecast errors for CPI inflation across households is, for example, about three times larger than across professional forecasters.

A.3 Forecasting VAR

We use a standard quarterly forecasting VAR to compute 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 change 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 an 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), and 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). This did not materially affect our results.

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	0.10
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 (b)	0.987
Death probability ($1-\rho$)	0.0045
Relative risk aversion (γ)	5.00
Resource cost of information (η)	0.0028
Scale parameter of utility cost of information (α^{κ})	$1/15e^{-4}$

B.2 Business Cycle Moments: Data Comparison

Table B.2: Comparison of Business Cycle Moments

Variable (x)	σ_x	σ_x/σ_y	Corr(x_t, x_{t-1})	Corr(x_t, y_t)	Corr(x_t, y_{t-1})
Panel (a): U.S. Data (1947Q1-2024Q4)					
Output (y)	1.64	1.00	0.79	1.00	0.79
Investment	7.26	4.42	0.78	0.82	0.61
Consumption.	1.39	0.85	0.72	0.79	0.54
Panel (b): Benchmark Model					
Variable (x)	σ_x	σ_x/σ_y	Corr(x_t, x_{t-1})	Corr(x_t, y_t)	Corr(x_t, y_{t-1})
Output (y)	3.38	1,00	0.83	1.00	0.83
Investment	11.93	3.53	0.76	0.97	0.75
Consumption	1.25	0.37	0.99	0.70	0.72
Panel (c): Full Information					
Variable (x)	σ_x	σ_x/σ_y	Corr(x_t, x_{t-1})	Corr(x_t, y_t)	Corr(x_t, y_{t-1})
Output (y)	3.04	1.00	0.78	1.00	0.78
Investment	10.53	3.46	0.72	0.97	0.70
Consumption	1.19	0.39	0.98	0.72	0.72

Note: The table reports the comparison of business cycle moments in different versions of our baseline model to that in the US data. The data on output, consumption, and investment come from FRED (code: `A939RX0Q048SBEA`, `A939RX0Q048SBEA`, $NFIRSAXDCUSQ$). We cumulate capital from investment using a 2.5% depreciation rate. All data variables are expressed in per capita terms and in logs. The data are at a quarterly frequency, and all cyclical components are extracted using an HP filter with $\lambda = 1{,}600$.

B.3 Time-series for Capital and Priors

Incomplete information makes individual 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 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 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 B.1 depicts the time series of the actual

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

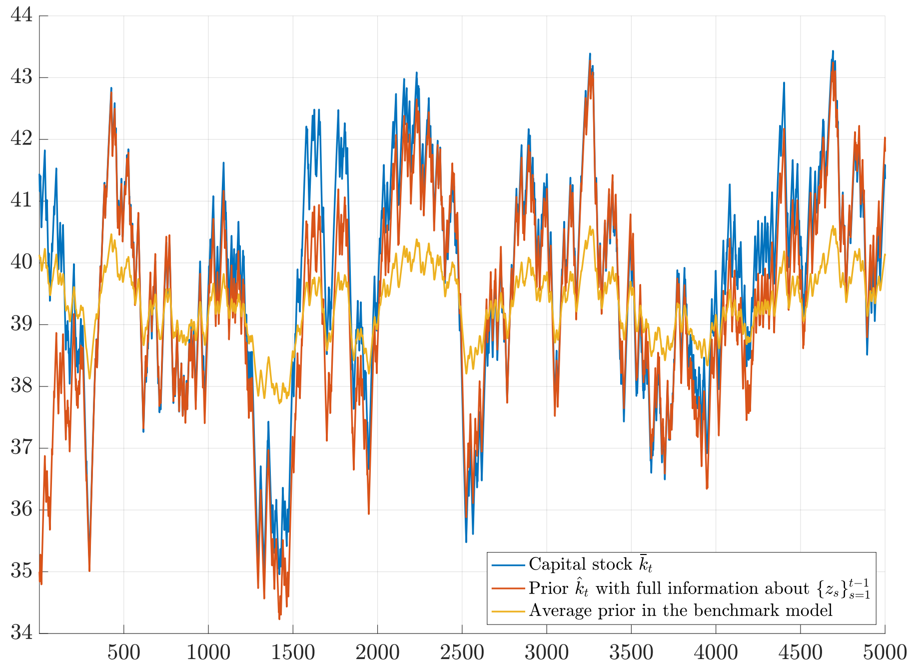

Note: Based on a simulation of the calibrated benchmark model, the figure shows time series of the mean (aggregate) capital stock $K_t = \bar{k}_t$ (blue line), the prior expectation of current aggregate capital $\hat{K}_t = \hat{k}_t$ of households who acquire information about the current productivity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

capital stock (blue line) and the prior expectation of an individual that has an arbitrary belief about capital in period 0 but then acquires information in every period (red line). We also, for comparison, plot the average prior expectation in our benchmark economy (yellow line). An individual that always acquires information would have prior expectations that closely track the realized value (with a correlation above 0.95).[53]

[53] In the figure, we start the prior 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.

C Extensions and Robustness

C.1 Mean-biased Capital Expectations

Table C.1: Mean-biased Capital Expectations—% Relative to Full Information

	Panel (a): Business Cycle Moments				
	$\sigma(K)$	$\sigma(Y)$	$\sigma(I)$	$\sigma(C)$	Cor(C, Y)
Endog. Info (P(learn K)=1.0)	5.49	1.03	6.40	-0.34	-15.15
Endog. Info (P(learn K)=0.1)	11.47	1.89	8.43	-0.90	-17.90
Benchmark model	64.24	11.03	13.61	5.08	-3.08
	Panel (b): Inequality Moments				
	Gini(G)	90/10	99/1	Cor(K, Y)	Cor(G, Y)
Endog. Info (P(learn K)=1.0)	-0.60	0.82	-2.21	3.16	-14.25
Endog. Info (P(learn K)=0.1)	-0.53	0.92	-2.46	5.25	-14.79
Benchmark model	1.60	2.66	4.98	24.66	86.15

Note: The table shows model moments in the calibrated model ("Benchmark Model"), as well as in models in which households exogenously obtain information about the level of the aggregate capital stock with a fixed probability $p \in [0, 1]$ ("Endogenous Information (P(learn K)=p)"). The table reports the percentage difference relative to the full-information version for the same moments as those computed in Table II and Table III.

C.2 Decomposition of Changes in the Wealth Distribution

Table C.2: Decomposition of Wealth Distribution Changes—% Relative to Full Information

	Overall	GE	Incomplete Info.	Heterogenous Info.	Interaction
Gini Statistic	1.50	-4.04	7.54	-1.65	-0.35
90/10−ratio	2.64	-2.99	7.31	-1.40	-0.28
99/1−ratio	4.67	-13.92	20.67	-3.73	1.65

Note: This table decomposes the overall change in the wealth distribution (relative to the full-information case) into the forces highlighted in Section 5.5. We focus, for concreteness, on the Gini, the 90/10-ratio, and the 99/1-ratio. Similar results hold for other summary inequality measures.

C.3 Acyclical Income Taxes

Table C.3: Business Cycle Moments—Acyclical Income Taxes

	Panel (a): Level of Moments					
	$\sigma(K)$	$\sigma(Y)$	$\sigma(I)$	$\sigma(C)$	Cor(C,Y)	Cor(I,Y)
Benchmark Model	3.98	3.15	9.33	0.98	0.64	0.99
Exogenous Information	3.30	3.05	9.00	0.94	0.61	0.99
Full Information	2.40	2.93	8.14	0.93	0.69	0.98
	Panel (b): Percent Difference w.r.t. Full Information					
	$\sigma(K)$	$\sigma(Y)$	$\sigma(I)$	$\sigma(C)$	Cor(C,Y)	Cor(I,Y)
Benchmark Model	66.08	7.72	14.68	4.86	-8.20	0.55
Exogenous Information	37.88	4.13	10.59	0.10	-11.79	0.29

Note: This 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 depicts these moments for the benchmark model ("Benchmark Model") as well as the two comparison models with full and exogenous information ("Full Information" and "Exogenous Information", respectively).

Table C.4: Inequality Moments—Acyclical Income Taxes

	Panel (a): Level of Moments				
	Gini G	90/10	99/1	Cor(K,Y)	Cor(G,Y)
Benchmark Model	0.51	14.48	337.3	0.54	-0.08
Exogenous Information	0.51	14.45	331.8	0.49	-0.07
Full Information	0.50	14.04	315.8	0.43	-0.05
	Panel (b): Percent Difference w.r.t. Full Information				
	Gini G	90/10	99/1	Cor(K,Y)	Cor(G,Y)
Benchmark Model	2.33	3.19	6.82	25.11	76.44
Exogenous Information	1.75	2.97	5.06	14.42	50.50

Note: The table shows 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 depicts these moments for the benchmark model ("Benchmark Model") as well as the two comparison models with full and exogenous information ("Full Information" and "Exogenous Information", respectively).

C.4 Costs of Information and Alternative Models

Table C.5: Business Cycle Moments—No Costs of Information

	$\sigma(K)$	$\sigma(Y)$	$\sigma(I)$	$\sigma(C)$	$\mathrm{Cor}(C, Y)$
Benchmark Model	5.05	3.38	11.98	1.25	0.7
	Panel (a): With Costs of Information				
Exogenous Information	4.19	3.22	11.56	1.19	0.67
Full Information	3.07	3.04	10.54	1.19	0.72
	Panel (b): Without Costs of Information				
Exogenous Information	4.19	3.22	11.55	1.19	0.67
Full Information	3.07	3.04	10.54	1.19	0.72
	Panel (c): Percentage Difference due to Costs				
Exogenous Information	-0.04	-0.01	-0.04	-0.08	-0.03
Full Information	-0.02	0,00	0,00	-0.21	-0.03

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 and output, respectively (e.g., $\mathrm{Cor}(C, Y)$). The table shows these moments for the benchmark model ("Benchmark Model"), the two comparison models with full and exogenously limited information ("Full Information" and "Exogenous Information", respectively), as well as versions of the comparison models in which all costs of information are set equal to zero.

Table C.6: Inequality Moments—No Costs of Information

	$\mathrm{Gini}(K)$	90/10	99/1	$\mathrm{Cor}(K, Y)$	$\mathrm{Cor}(\mathrm{G}, Y)$
Benchmark Model	0.51	14.53	345.3	0.62	-0.09
	Panel (a): With Costs of Information				
Exogenous Information	0.51	14.51	339.4	0.56	-0.07
Full Information	0.50	14.16	328.9	0.49	-0.05
	Panel (b): Without Costs of Information				
Exogenous Information	0.51	14.5	337.6	0.56	-0.07
Full Information	0.50	14.13	326	0.49	-0.05
	Panel (c): Percentage Difference due to Costs				
Exogenous Information	-0.20	-0.11	-0.53	-0.02	-0.16
Full Information	-0.31	-0.18	-0.89	-0.01	-0.55

Note: The table shows the mean of the logarithm of capital (K), the Gini coefficient of the capital distribution (Gini/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., $\mathrm{Corr}(\mathrm{G}, Y)$). The table shows these moments for the benchmark model ("Benchmark Model"), the two comparison models with full and exogenously-limited information ("Full Information" and "Exogenous Information", respectively), as well as versions of the comparison models in which all costs of information are set equal to zero.

C.5 Homotheticity of Information Costs

Figure C.1: Accuracy Across the Wealth Distribution—Scaled Utility Cost

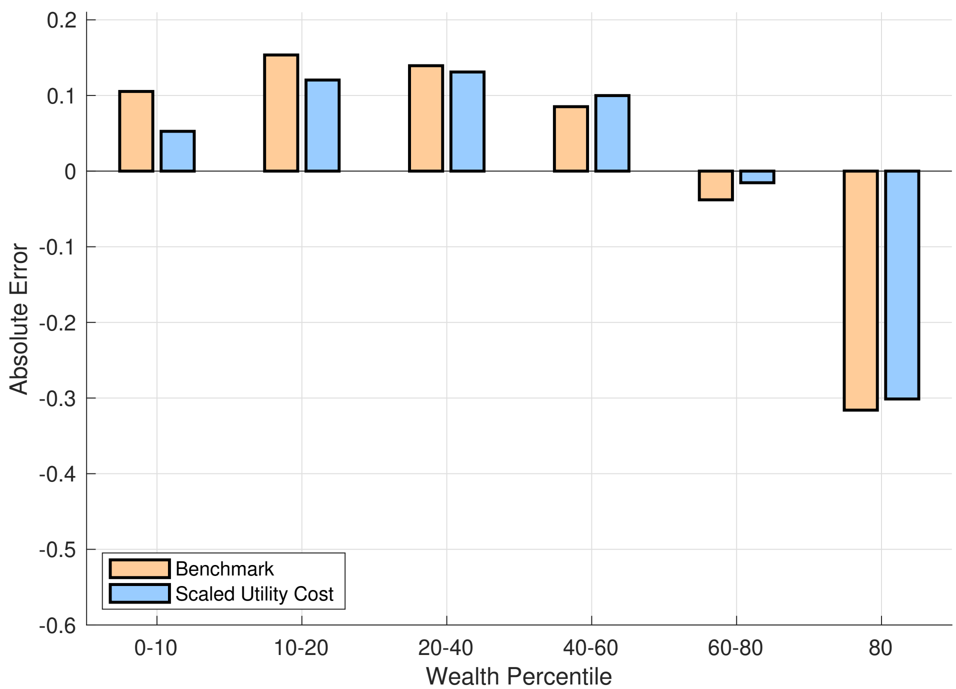

Note: The figure plots the difference between the average one-year-ahead accuracy of unemployment forecasts within wealth deciles/quintiles and the overall average taken across all wealth levels. Accuracy is measured by the absolute value of unemployment errors (Online Appendix A.2). We plot these accuracies both for the calibrated benchmark model and for the calibrated model with the scaled utility cost.

Table C.7: Model Moments—Scaled Utility Cost

	Panel (a): Business Cycle Moments				
	$\sigma(K)$	$\sigma(Y)$	$\sigma(I)$	$\sigma(C)$	$\mathrm{Cor}(C,Y)$
Benchmark	5.05	3.38	11.98	1.25	0.70
Scaled Utility Cost	4.96	3.36	11.92	1.24	0.70
Difference %	-1.77	-0.52	-0.43	-0.67	-0.37
	Panel (b): Inequality Moments				
	$\mathrm{Gini}(K)$	90/10	99/1	$\mathrm{Cor}(K,Y)$	$\mathrm{Cor}(G,Y)$
Benchmark	0.51	14.53	345.30	0.62	-0.09
Scaled Utility Cost	0.51	14.49	343.10	0.61	-0.09
Difference %	-0.19	-0.33	-0.64	-0.85	-2.11

Note: The table shows the same model moments as in Section 5 for the benchmark (endogenous information) model ("Benchmark") with and without the scaled utility cost of information κ set equal to zero.

Figure C.2: Accuracy Across the Wealth Distribution—No Resource Cost

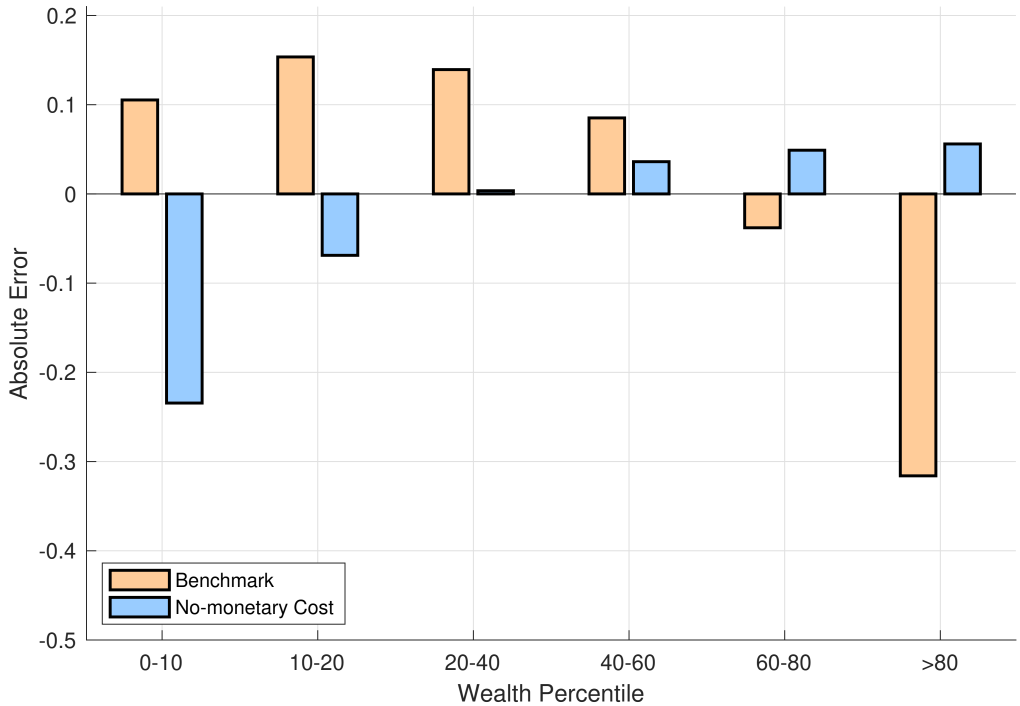

Note: The figure plots the difference between the average one-year-ahead accuracy of unemployment forecasts within wealth deciles/quintiles and the overall average taken across all wealth levels. Accuracy is measured by the absolute value of unemployment errors (Online Appendix A.2). We plot these accuracies both for the calibrated benchmark model and for the calibrated model with no-resource cost.

Table C.8: Model Moments—No Resource Costs

	Panel (a): Business Cycle Moments				
	$\sigma(K)$	$\sigma(Y)$	$\sigma(I)$	$\sigma(C)$	Cor(C, Y)
With Resource Cost	5.05	3.38	11.98	1.25	0.70
Without Resource Cost	4.11	3.20	11.48	1.19	0.68
Difference %	-18.49	-5.10	-4.11	-5.01	-3.82
	Panel (b): Inequality Moments				
	Gini(K)	90/10	99/1	Cor(K, Y)	Cor(G, Y)
With Resource Cost	0.51	14.53	345.3	0.62	-0.09
Without Resource Cost	0.50	14.11	327.96	0.56	-0.07
Difference %	-1.50	-2.88	-5.02	-9.08	-17.58

Note: The table shows the same model moments as in Section 5 for the benchmark (endogenous information) model ("Benchmark Model") with and without the resource cost of information η set equal to zero.

C.6 Matching the Wealth Distribution

C.6.1 Extended Model Framework

We assume a representative final-goods producer bundles varieties $j \in \mathcal{J}$ of differentiated goods according to the Dixit-Stiglitz aggregator:

$$Y_t = \left(\int y_{jt}^{\frac{\omega-1}{\omega}} d_j \right)^{\frac{\omega}{\omega-1}} \tag{A1}$$

with elasticity of substitution ω. Each of the differentiated goods is sold at price p_j, so that the ideal price level equals $P_t = \left(\int p_{jt}^{1-\omega} dj \right)^{\frac{1}{1-\omega}}$, which we normalize to one. Consistent with this structure, the demand for each variety is

$$y_{jt} = p_{jt}^{-\omega} Y_t.$$

Intermediate goods firms produce using the production technology in (3.3) and are characterized by the same assumptions as firms in our benchmark framework.

There are two types of households in the economy: A mass $m \in (0, 1)$ of worker-households and a mass $1 - m$ of entrepreneur-households. Worker-households are identical to households in our benchmark model. Entrepreneurs, by contrast, receive all pure rents (i.e., profit income) in the economy but no labor income. Entrepreneurs are across all other dimensions identical to worker-households. We assume the existence of fixed transition probability matrix, Π_T, modeled as in Bayer *et al.* (2024), so that worker-household transition to become entrepreneur-households, and vice versa. All other model details are as described in Section 3.

C.6.2 Alternative Calibration and Quantification

We follow an identical calibration strategy to that of our benchmark model. We set the additional parameters—the mass of entrepreneurs and the CES elasticity—to 0.5 percent and 20, respectively, consistent with Bayer *et al.* (2024). We set the transition matrix, Π_T, to the values estimated in Guvenen *et al.* (2014), consistent with the approach in Bayer *et al.* (2024). Figure C.3 compares the accuracy of unemployment forecasts across the wealth distribution with that in the benchmark model; Table C.9 shows the business-cycle and inequality moments, as well as the decomposition of the change in the summary measures of inequality.

Figure C.3: Accuracy Across the Wealth Distribution—Extended and Benchmark Model

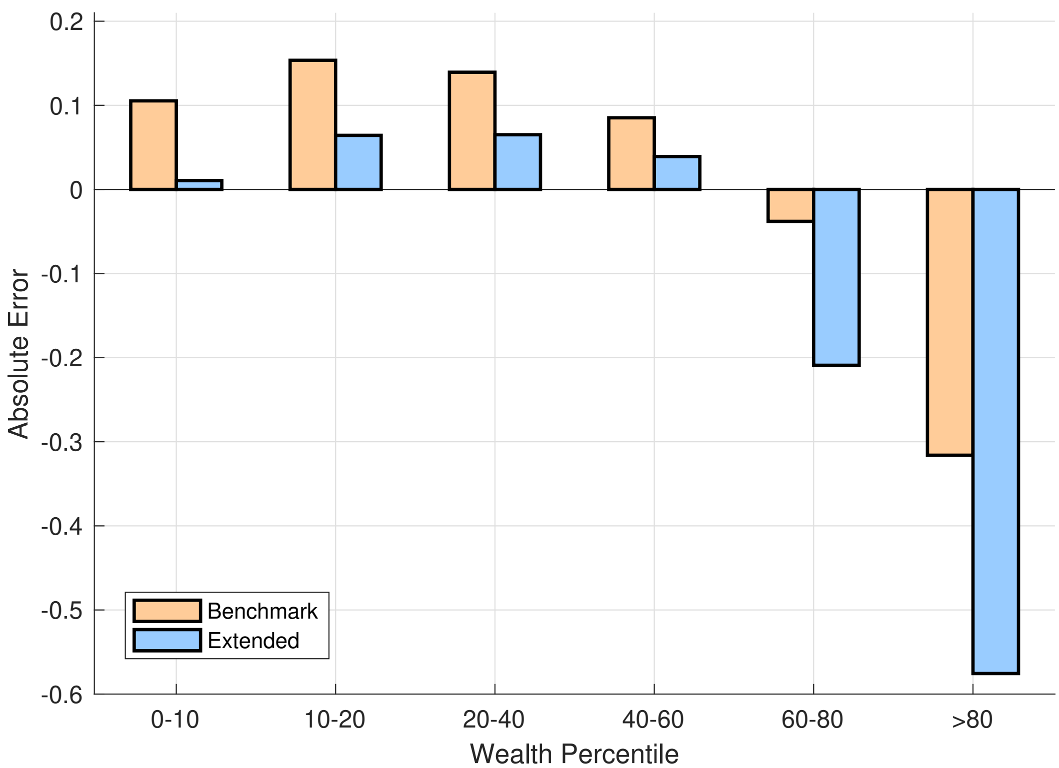

Note: The figure plots the difference between the average one-year ahead accuracy of unemployment forecasts within wealth deciles/quintiles and the overall average taken across all wealth levels. Accuracy is measured by the absolute value of unemployment errors (Online Appendix A.2). We plot these accuracies both for the calibrated benchmark model and for the calibrated extended model that matches the wealth distribution.

Table C.9: Extended Model Moments: Business Cycle and Inequality

	Panel (a): Business Cycle Moments				
	$\sigma(K)$	$\sigma(Y)$	$\sigma(I)$	$\sigma(C)$	Cor(C, Y)
Endogenous Information (A)	2.45	2.91	11.01	1.00	0.61
Exogenous Information (B)	2.23	2.90	10.78	1.08	0.61
Full Information (C)	1.91	2.86	9.57	1.17	0.79
Difference % (A vs C)	28.11	1.61	15.07	-15.22	-23.14

	Panel (b): Inequality Moments				
	Gini(K)	90/10	99/1	Cor(K, Y)	Cor(G, Y)
Endogeous Information (A)	0.79	3.68	1239.86	0.41	-0.14
Exogenous Information (B)	0.79	3.80	1268.09	0.41	-0.13
Full Information (C)	0.79	4.01	1317.89	0.38	-0.12
Difference % (A vs C)	0.12	-8.25	5.92	5.69	22.34

	Panel (c): Inequality Moments—Decomposition (%pp)		
	Gini(K)	90/10	99/1
General Equilibrium	1.74	13.69	15.60
Incomplete Information	0.98	-0.64	15.84
Heterogenous Information	-2.55	-18.78	-25.20
Overall Difference	0.12	-8.25	5.92

Note: Panel (a) and (b) show the same model moments as in Table C.5 and Table C.6 but for the extended benchmark model (endogenous information) as well as its full-information and exogenous-information counterparts. Panel (c) provides a decomposition of the change in the inequality statistics as in Table C.2.

C.7 Internalized Uncertainty about K

The log capital stock evolves according to the posited law of motion:

$$\log K_{t+1} = a_{z_t} + b_{z_t} \log K_t, \qquad z_t \in \{0,1\}. \tag{A2}$$

Since in practice $b_1 > b_0$ and $a_1 > a_0$, the time series for $\log K_t$ is bounded between:

$$\underline{\log K} = \frac{a_0}{1-b_0}, \qquad \overline{\log K} \ = \frac{a_1}{1-b_1}.$$

Let $\phi_N(\cdot)$ and $\Phi_N(\cdot)$ denote the standard normal pdf. and cdf., respectively. At every t, we assume that the posterior of $\log K_t$ is *approximated* by a truncated normal $\mathcal{TN}\Big(m_t, v_t, [L, U]\Big)$, i.e.:

$$\log K_t \mid \Omega_{it} \overset{\text{approx.}}{\sim} \mathcal{TN}\Big(m_{it}, v_{it}, [\underline{\log K}, \overline{\log K}]\Big).$$

We let μ_{it} and ς^2_{it} denote the household's mean and variance of $\log K_t$ *before* accounting for truncation. The household then updates its beliefs in accordance with:

Informed Households: Let $\alpha_{it} \equiv \varsigma_{it}^{-1}\left(\underline{\log K} - \mu_{it}\right)$ and $\beta_{it} \equiv \varsigma_{it}^{-1}\left(\overline{\log K} - \mu_{it}\right)$. Then, consistent with Equation (A2), for given (m_{it}, v_{it}) and realized z:

$$\mu_{it} \ = \ a_z + b_z\, m_{it}, \qquad \varsigma^2_{it} \ = \ b_z^2 v_{it} \tag{A3}$$

so that

$$m_{it+1} = \mu_{it} + \varsigma_{it} \frac{\phi_N(\alpha_{it}) - \phi_N(\beta_{it})}{\Phi_N(\beta_{it}) - \Phi_N(\alpha_{it})}, \tag{A4}$$

$$v_{it+1} \ = \varsigma^2_{it}\left[1 + \frac{\alpha_{it}\phi_N(\alpha_{it}) - \beta_{it}\phi_N(\beta_{it})}{\Phi_N(\beta_{it}) - \Phi_N(\alpha_{it})} - \left(\frac{\phi_N(\alpha_{it}) - \phi_N(\beta_{it})}{\Phi_N(\beta_{it}) - \Phi_N(\alpha_{it})}\right)^2\right]. \tag{A5}$$

Uninformed Households: Let $p_{it,z} = \mathbb{E}(z_t = 1 \mid \Omega_{it})$ denote an uninformed household's expectation of a boom. Given (m_{it}, v_{it}), it follows that:

$$\mu_{it} = (1 - p_{it,z})\Big(a_0 + b_0\, m_{it}\Big) + p_{it,z}\Big(a_1 + b_1\, m_{it}\Big) \tag{A6}$$

$$\varsigma^2_{it} = (1 - p_{it,z})\, b_0^2 v_{it} + p_{it,z}\, b_1^2 v_{it} + p_{it,z}(1 - p_{it,z})\, \Delta^2_{it}, \tag{A7}$$

where $\Delta_{it} \equiv (a_1 - a_0) + (b_1 - b_0)\, m_{it}$. Truncation is identical to that in (A4) and (A5).

Implementation: The solution algorithm is updated as follows. Instead of assuming households have a point expectation about tomorrow's capital stock, when solving their consumption-savings and information acquisition problem in *Stage 1* and *Stage 2*, households now internalize their uncertainty about tomorrow's capital stock, and that their beliefs evolve in accordance with Equations (A3) to (A7). Beyond how beliefs are updated, all other steps of the solution algorithm are identical to before.

Quantitative Results: We calibrate the model in the same manner as all extensions considered in Section 5. Consistent with the greater uncertainty faced by households, to match the data requires a larger resource and utility cost of information. Specifically, we set $\eta = 0.006$ and the mean of $\kappa = 0.0001$ to match the data. The model-implied errors imply a similar, although slightly milder, inverse-u shape to that in the benchmark model. Table C.10 shows our main results, which on balance confirm the insights from the benchmark specification.

Table C.10: Model Moments: Business Cycle and Inequality

	Panel (a): Business Cycle Moments				
	$\sigma(K)$	$\sigma(Y)$	$\sigma(I)$	$\sigma(C)$	Corr(C, Y)
Endogenous Information (A)	4.12	3.20	11.47	1.17	0.67
Exogenous Information (B)	3.64	3.13	11.16	1.17	0.66
Full Information (C)	3.20	3.05	10.55	1.17	0.71
Difference % (A vs C)	28.75	4.92	8.72	0.05	-5.97

	Panel (b): Inequality Moments				
	Gini(K)	90/10	99/1	Cor(K, Y)	Cor(G, Y)
Endogenous Information (A)	0.49	13.57	318.71	0.56	-0.07
Exogenous Information (B)	0.48	13.25	306.44	0.54	-0.06
Full Information (C)	0.47	12.84	279.27	0.50	-0.04
Difference % (A vs C)	4.26	5.69	14.12	12.00	75.01

Note: Panel (a) and (b) show the model moments as in Table C.5 and Table C.6 but for the extended benchmark model (endogenous information) as well as its full-information and exogenous-information counterparts.

D Counterfactual Policy Experiments

D.1 A Wealth Tax

Figure D.1: Wealth Taxes and Changes in the Wealth Distribution

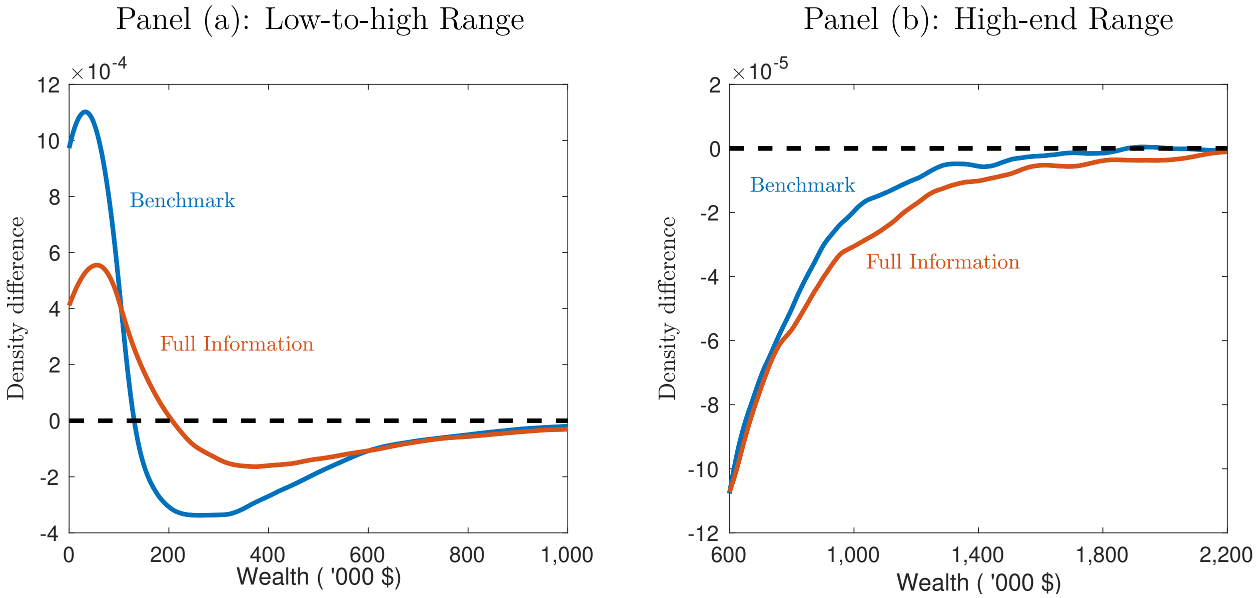

Note: The figure illustrates changes in the average wealth distribution relative to the zero-wealth tax case. We illustrate these changes for both our calibrated model ("Benchmark Model") and the associated full-information economy ("Full Information"). We use 2020 values of US household income to convert 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.

D.2 Increased Unemployment Benefits

Figure D.2: Unemployment Benefit Increases and Changes in the Wealth Distribution

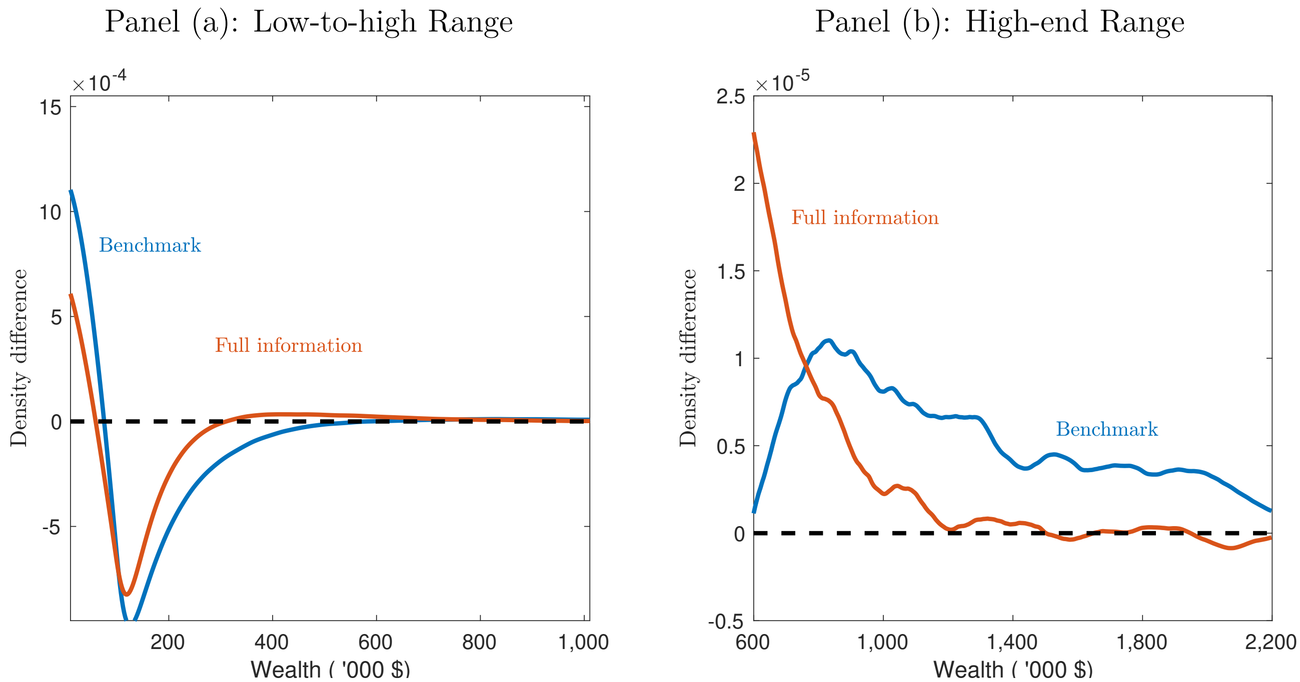

Note: The figure illustrates changes in the average wealth distribution relative to the zero-wealth tax case. We illustrate these changes for both our calibrated model ("Benchmark Model") and the associated full-information economy ("Full Information"). We use 2020 values of US household income to convert 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.

D.3 Policy Experiments in the Extended Model Framework

Table D.1: Extended Model: Effects of a Wealth Tax

Entrepreneur Model	K	$\sigma(Y)$	Gini(K)	90/10	99/1	Info U.	Info E.
Endogenous Information	-6.93	0.44	0.52	-1.29	3.30	1.32	-9.74
Exogenous Information	-6.86	0.46	0.36	-2.64	2.79	0.00	0.00
Full Information	-7.58	0.44	-0.13	-7.66	-4.93	0.00	0.00

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 in percent. K denotes the mean of the (log-)capital stock, while $\sigma(Y)$ denotes the variance of (log-)economy-wide output. The table computes the moments for the both calibrated model ("Benchmark Model") and the associated full-information and exogenous-information economies ("Full Information" and "Exogenous Information", respectively). 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.

Table D.2: Extended Model: Effects of Increased Unemployment Benefits

Entrepreneur Model	K	$\sigma(Y)$	Gini(K)	90/10	99/1	Info U.	Info E.
Endogenous Information	-3.35	0.32	5.62	17.47	157.25	-20.29	4.27
Exogenous Information	-3.27	0.26	5.34	15.56	146.81	0.00	0.00
Full Information	-4.40	0.40	4.91	9.65	48.41	0.00	0.00

Note: The table shows the effects of a ten percentage-point increase in the replacement rate on moments of the average cross-sectional capital distribution and the logarithm of economy-wide output in percent. K denotes the mean of the (log-)capital stock, while $\sigma(Y)$ denotes the variance of (log-)economy-wide output. The table computes the moments for the both benchmark economy ("Benchmark Model") and the associated full-information and exogenous-information economies ("Full Information" and "Exogenous Information", respectively). 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.