Learning to Be Rich: How Expectations Amplify Wealth Inequality*

Pau Belda
Janko Heineken
Adrian Ifrim

November 2026
Latest version here

Households persistently disagree about future stock market returns. Survey data reveal that this disagreement is mainly driven by stable individual differences, which are largely unexplained by observable household characteristics, even though average optimism rises sharply with wealth. We develop a general-equilibrium heterogeneous-agent model to study the origins of this disagreement and quantify its consequences for the distribution of wealth. In the model, households learn about stock market returns from their own portfolio performance, so idiosyncratic investment experiences generate persistent disagreement. Analytically, greater disagreement can thicken the upper tail of the wealth distribution by creating rare winners among moderately optimistic households, even as it lowers average long-run wealth growth. Quantitatively, learning from personal experience increases the top-decile wealth share by 12 percentage points. Methodologically, departing from full-information rational expectations not only helps reconcile the model with survey evidence but also makes heterogeneous-agent models with aggregate risk substantially more tractable.

Keywords: Heterogeneous Expectations, Wealth Inequality, Learning, Internal Rationality.

*Preliminary draft. First version: May 2023. Previous versions have circulated under the title *"Heterogeneous Expectations and Wealth Inequality"*. Belda: Bank of England (Pau.Belda@bankofengland.co.uk), Heineken: European Commission (Janko.Heineken@ec.europa.eu), Ifrim: European Commission (Adrian.Ifrim@ec.europa.eu). We have benefited from comments from Albert Marcet, Juan Rubio-Ramírez, Wouter Den Haan, Olivier Coibion, Inês Xavier, and seminar participants at the NBER - Summer Institute, BSE - Summer Forum, SED meeting, Federal Reserve Board, CREi, Bellaterra Macroeconomics Workshop, ERMAS, Bank of England. The views expressed here are solely of the authors and do not represent, in any way, those of the Bank of England, the European Commission or any of its committees. This research received funding from the European Research Council under the European Union's Horizon2020, research and innovation program GA project number 788547 (APMPAL-HET).

1. INTRODUCTION

Theories of wealth inequality have increasingly shifted their emphasis from differences in what labor earns to differences in what wealth earns. A long tradition traced the concentration of wealth to heterogeneity in earnings, which modern incomplete-markets models formalize as idiosyncratic, uninsurable labor-income risk.[1] Yet earnings heterogeneity alone struggles to reproduce the extreme concentration of wealth observed in the data,[2] and a growing literature has turned to heterogeneity in the returns that households earn on their wealth.[3] Within this literature, a central open question is why returns rise with wealth. We propose that differences in households' expectations about stock market returns help explain this pattern by shaping the portfolios they hold and the returns they earn.

In particular, we study how disagreement about stock market returns arises and how it shapes the distribution of wealth. We develop a heterogeneous-agent model with aggregate risk in which households do not know how equilibrium asset prices are determined and learn about them from their own investment experiences. We discipline the model with household-level and aggregate data and show that it jointly accounts for the persistence and wealth gradient of beliefs, the distribution of income, wealth, and portfolios, and salient asset-pricing facts, including a large equity premium. Our main finding is that heterogeneous beliefs make an important contribution to the concentration of wealth at the top.

The theory is motivated by two facts about household return expectations.[4] First, consistent with Giglio et al. (2021), households tend to remain optimistic or pessimistic over time, and these persistent differences in beliefs are largely unexplained by observable characteristics such as wealth. Second, we document that, beneath this individual dispersion, beliefs are strongly ordered by wealth. Averaging within wealth percentiles, expected stock returns rise by roughly 4 percentage points from the bottom to the top of the wealth distribution, and forecast errors decline, so that over our sample period, in which realized returns were predominantly positive, wealthier households' expectations were also closer to realized returns.

[1] A large literature studies how earnings heterogeneity, incomplete markets, and saving behavior shape the wealth distribution. See, among others, Huggett (1993) and Aiyagari (1994), and the reviews in De Nardi and Fella (2017) and Benhabib and Bisin (2018).

[2] Stachurski and Toda (2019) prove that in canonical incomplete-markets models with a common, constant return on wealth and a common discount factor, the stationary wealth distribution cannot have a heavier tail than the earnings distribution. Quantitatively, Castañeda, Díaz-Giménez, and Ríos-Rull (2003) match top wealth shares only by assuming very large and transitory earnings realizations at the top, and De Nardi, Fella, and Paz Pardo (2016), Benhabib and Bisin (2018), and Hubmer, Krusell, and Smith (2021) find that empirically grounded earnings processes leave top wealth shares substantially understated.

[3] For empirical evidence on heterogeneity in returns to wealth, see Fagereng et al. (2020), Bach, Calvet, and Sodini (2020), and Xavier (2021). For models studying its implications for wealth concentration, see Gabaix et al. (2016), Benhabib, Bisin, and Luo (2019), Xavier (2021), and Hubmer, Krusell, and Smith (2021).

[4] We document both facts in the RAND American Life Panel, which follows households over 2008–2017 and elicits subjective probabilities about one-year stock market returns. We use these probabilities to construct expected returns. Wealth in the RAND data is measured from housing and retirement-account holdings; Section 2.1 describes the construction and data in detail.

Wealth therefore explains little of why a given household is persistently optimistic or pessimistic, yet richer households are substantially more optimistic on average. A theory of belief heterogeneity must deliver both a low household-level R^2 and a steep wealth–belief schedule.

Learning from one's own investment experience can account for both facts. In the model, households invest through intermediaries who report portfolio returns that combine market performance with idiosyncratic shocks. Unable to disentangle these components, households interpret their own portfolio returns as informative about future market performance. Bayesian updating therefore turns personal return histories into heterogeneous beliefs, even among households exposed to the same aggregate market. The process is self-referential: beliefs determine portfolio exposure, personal returns jointly affect beliefs and wealth, and subsequent portfolio decisions transmit these differences into future wealth accumulation. When learning is slow, independent personal shocks cumulate into persistent disagreement about the same market. At the same time, favorable histories raise wealth, optimism, and risky exposure together, which can generate the empirical ordering of average beliefs by wealth even though wealth is not the primitive source of disagreement. We call this feedback *learning to be rich*. The UBS/Gallup investor survey provides supporting evidence for its first link. As in Vissing-Jørgensen (2003), investors who have experienced higher returns on their own portfolios report higher expected stock market returns; conditional on survey-wave fixed effects, which absorb common market conditions, the elasticity of expected market returns with respect to own past portfolio returns is positive and significant.[5]

In an analytic version of the model, greater belief disagreement can thicken the upper tail of the wealth distribution even as it lowers average long-run wealth growth.[6] Accurate beliefs maximize expected log wealth growth, while both pessimism and optimism reduce it by inducing too little or too much risk-taking. Hence, a symmetric mean-preserving spread in beliefs lowers average growth. Yet its effects on extreme wealth are different. Within each belief group, idiosyncratic return shocks generate a Pareto upper tail whose thickness depends on wealth growth and risk exposure. Moderate optimism increases exposure to these shocks to first order, while reducing growth only to second order. Consequently, the heaviest tail belongs to moderately optimistic households, rather than those with accurate or excessively optimistic beliefs. Because the population's upper tail is dominated by the group with the heaviest tail, greater disagreement can increase wealth concentration through rare, exceptionally wealthy optimists, even as average wealth growth declines.[7]

[5] Jiang et al. (2025) provide complementary evidence from Chinese retail investors: among investors exposed to the same market movements, those who recall higher returns on their own portfolios report higher expected market returns and subsequently buy more stock, and recalled own returns explain beliefs better than actual returns do.

[6] In the analytic version of the model, beliefs are frozen at their persistent component, asset prices are held fixed, and aggregate return risk is shut down.

[7] The growth result is closely related to Martin and Papadimitriou (2022), who show in a log-utility environment that wealth flows toward investors whose beliefs prove most accurate. The tail result shows that selection toward accuracy in typical outcomes coexists with selection toward optimism in extreme ones.

General equilibrium adds a feedback from beliefs to the realized returns that shape subsequent choices. Greater optimism raises the demand for stocks and hence their price, so optimistic households partly generate the capital gains that enter their personal investment experiences. The feedback is not unambiguously amplifying, however. Higher prices lower subsequent stock returns, and weaker demand for safe assets can raise the equilibrium safe rate, both of which work against optimists. General equilibrium therefore changes both the persistence of disagreement and its consequences for wealth accumulation, and the net effect is a quantitative question.

We discipline the quantitative model with a rich set of household-level and aggregate moments. Because no single data set measures beliefs and complete balance sheets together, we use the characteristics observed in both the RAND panel and the Survey of Consumer Finances to impute expected returns into the latter, which yields a joint distribution of beliefs, portfolios, income, and wealth. We calibrate the labor-income process to the U.S. distribution of earnings, discipline aggregate risk with U.S. dividend and asset-return data, and allow for persistent heterogeneity in patience and risk aversion, so that beliefs are not asked to account for all observed heterogeneity in portfolios. All households start from the same prior. Conditional on these inputs, beliefs, portfolio choices, asset prices, capital income, and the wealth distribution are determined endogenously.

The quantitative model closely matches the observed concentration of financial wealth: the top 10 percent hold 73.5 percent, compared with 75.7 percent in the data, while the top 1 percent hold 39.4 percent, compared with 37.6 percent. The model also generates a positive relationship between wealth and expected stock returns, which rise from 4.9 percent among the bottom half to 7.9 percent among the top 1 percent. Finally, the model reproduces the strong persistence of beliefs observed in the RAND survey: households tend to remain optimistic or pessimistic over time, with these persistent individual differences explaining 63.5 percent of the variation in expectations, compared with 57.3 percent in the data.

Learning from personal investment experience substantially amplifies wealth concentration at the top. To quantify this effect, we compare the baseline with an economy in which households learn from the common market return rather than their own portfolio returns. Both economies start from the same prior, share identical structural parameters and learning gains, and determine beliefs, portfolio allocations, stock prices, and safe returns endogenously. Removing personal-experience learning reduces the top-decile wealth share from 73.5 to 61.8 percent, a decline of 11.7 percentage points, and the top-1-percent share from 39.4 to 31.2 percent.

We then ask whether the same mechanism can account for the historical rise in U.S. wealth concentration. Whereas the stationary analysis determines asset prices endogenously, this exercise conditions on the observed U.S. stock-price and dividend series over 1960–2024 and studies transitional dynamics. Starting from the observed initial wealth distribution, with initial beliefs chosen to reproduce the wealth gradient measured in RAND, portfolio choices, wealth accumulation, and beliefs evolve endogenously. Without targeting the path of wealth shares, the economy with heterogeneous learning generates a sustained increase in top wealth

shares that captures much of the historical rise in U.S. concentration, whereas a common-belief economy exposed to the same market history generates much flatter dynamics. Heterogeneous beliefs thus help account not only for the level of wealth inequality but also for its evolution.

Methodologically, departing from full-information rational expectations permits a coherent model of individual return learning while retaining a substantially simpler solution architecture. We adopt Internal Rationality (Adam and Marcet, 2011): households maximize expected utility under a well-specified subjective probability measure over the payoffs and information they face, but need not know the equilibrium mapping from the cross-sectional distribution of households to prices. Households observe individual risky returns and current bond prices, and they forecast returns directly from their subjective models. This allows realized portfolio experiences to inform beliefs about future market performance while households remain fully rational conditional on their information. The approach also sidesteps a central computational difficulty of heterogeneous-agent models with aggregate risk. Under rational expectations, equilibrium prices generally depend on the evolving cross-sectional distribution, requiring households to forecast its law of motion in order to forecast future prices. Under Internal Rationality, households instead forecast the returns on their investments directly from subjective models, so the distribution does not enter their individual state space. This separation property underlies the Two-Stage Deep Learning (2SDL) algorithm we develop: we first solve the household problem under subjective beliefs, approximating policy functions with a neural network, and only then aggregate the resulting asset demands and impose market clearing to recover equilibrium prices.

Our paper contributes to the literature on heterogeneous beliefs and wealth inequality by providing a unified account of the persistent disagreement observed in household microdata and the concentration of wealth. The closest quantitative studies—Acedański (2017), Chiu (2024), Broer et al. (2022), and Lei (2026)—combine different subsets of the ingredients in our framework: heterogeneous beliefs, heterogeneous returns and portfolios, and general-equilibrium asset prices. Lei (2026), for instance, lets cohorts learn from the aggregate return histories they have lived through, generating disagreement across cohorts, while Broer et al. (2022) generates heterogeneous expectations through endogenous information acquisition within a rational-expectations framework. Our contribution is to bring these ingredients together in a model in which pervasive, individual-level disagreement about stock returns emerges endogenously from a common prior once households are allowed to learn about asset prices from their own investment experiences.

We also build on two literatures on learning. In adaptive-learning models of asset prices, investors learn about the price process itself, and this learning has been used mainly to account for boom–bust dynamics (e.g., Timmermann, 1996; Carceles-Poveda and Giannitsarou, 2008; Adam, Marcet, and Nicolini, 2016; Adam, Marcet, and Beutel, 2017; Winkler, 2020), with some extensions introducing heterogeneous investors and endogenous trade (Adam et al., 2015; Belda, 2026). Our focus is instead the cross-section: the joint distribution of beliefs, portfolios, and wealth across the

population. Experience-based learning (Malmendier and Nagel, 2011; Malmendier and Nagel, 2016) shows that households overweight the macroeconomic outcomes they have lived through, so that cohorts exposed to different aggregate histories hold different beliefs. We shift attention from aggregate histories to individual ones: households extrapolate from the returns on their own portfolios, so even households with identical aggregate histories come to disagree. This connects to evidence that personal and local experiences shape beliefs about aggregate conditions (Kuchler and Zafar, 2019; D'Acunto et al., 2021; Fidrmuc, Hainz, and Hölzl, 2024). We bring this individual-to-aggregate extrapolation into portfolio choice, where it becomes a source of persistent disagreement, heterogeneous returns, and wealth inequality.

We further provide a microfoundation for persistent heterogeneity in returns to wealth that links measured beliefs to portfolio choice and wealth concentration. Household data from several countries reveal large and persistent differences in returns across households, arising both from portfolio composition and from return differences within asset classes (Fagereng et al., 2020; Bach, Calvet, and Sodini, 2020; Xavier, 2021), and quantitative models show that such return heterogeneity can be central to top wealth concentration (Gabaix et al., 2016; Benhabib, Bisin, and Luo, 2019; Xavier, 2021; Hubmer, Krusell, and Smith, 2021). In most of this work, however, the return process or the relation between portfolio returns and wealth is taken as given. Closest to our equilibrium environment, Fernández-Villaverde and Levintal (2024) and Gomez (2025) determine portfolio exposure, asset prices, and the wealth distribution jointly. We add household beliefs to this interaction and discipline them directly with expectations data, which allows the model to account for the distributions of beliefs, portfolios, returns, and wealth together.

Finally, we contribute to the computation of heterogeneous-agent models with aggregate risk. Recent work also uses low-dimensional market observables to reduce the dimensionality of these models (Yang et al., 2025; Cioffi, Hurtado, and Nuño, 2026), motivated primarily by computation: price-based forecasting rules approximate equilibrium expectations without tracking the full wealth distribution. Our motivation is economic. We depart from rational expectations because the pervasive disagreement in household surveys calls for a theory of subjective expectations, and Internal Rationality supplies one in which households optimize under subjective return forecasts that the survey data discipline directly. Tractability is a by-product of the theory rather than its purpose.

The rest of the paper is organized as follows. Section 2 documents the facts about expected returns and wealth. Section 3 presents the model and the mechanism of learning from own portfolio experience. Section 4 derives the analytical results on disagreement, wealth growth, and wealth tails. Section 5 describes the solution algorithm. Section 6 presents the calibration, the quantitative results, and the historical transition exercise. Section 7 concludes.

2. Facts about Expected Returns and Wealth

This section documents the empirical facts about household return beliefs that discipline the model. We use the RAND American Life Panel (ALP), which follows

the same households over time and jointly records stock-market expectations and household characteristics. Two facts are central. First, as in Giglio et al. (2021), disagreement is predominantly persistent, yet standard observables have very low explanatory power for the persistent component of beliefs. In RAND, individual fixed effects account for more than half of the panel variation in expected returns, while wealth alone explains less than 10 percent of the cross-sectional variation in these individual effects. Second, low explanatory power at the household level does not imply that beliefs are unstructured. Once we average idiosyncratic disagreement within wealth percentiles, a remarkably tight gradient emerges: expected stock returns rise systematically with wealth and forecast errors fall. Thus, the empirical object is a combination of substantial within-wealth disagreement and a highly structured conditional mean of beliefs across the wealth distribution.

Section 2.1 describes the survey and our construction of expected returns. Section 2.2 studies the persistence of disagreement following Giglio et al. (2021) and asks how much of the persistent component is explained by wealth and other observables. Section 2.3 characterizes the wealth gradient in beliefs directly, both using our scalar expected-return measure and using the underlying subjective probabilities reported by households. Appendix A provides additional details on the RAND questions and sample construction. Appendix B reports the forecast-error predictability tests, which are useful for benchmarking RAND against the broader expectations literature but are not central to the long-run mechanism studied in the paper.

2.1. Data and Measurement of Return Expectations

RAND American Life Panel. The RAND ALP is a U.S. household panel that contains 61 survey waves between 2008 and 2017. The expectations module forms an unbalanced panel of roughly 3,000 respondents. Its main advantage for our purposes is that repeated observations of stock-market beliefs can be linked to household wealth and other household characteristics. The panel dimension allows us to distinguish persistent differences across households from common time variation in expectations, while the wealth information allows us to study how this persistent disagreement is distributed across households.

The survey does not ask respondents for a single point forecast. Instead, it elicits probabilities for several events concerning the value of a mutual fund invested in blue-chip stocks one year ahead. Let p_{it}^{+} denote household i's reported probability in wave t that the fund increases in value, p_{it}^{20} the probability that it increases by more than 20 percent, and p_{it}^{-20} the probability that it falls by more than 20 percent. These answers identify probabilities for four mutually exclusive regions of

the one-year-ahead return distribution:

$$p_{1,it} \equiv \Pr_{it}(R_{t+1} < 0.8) = p_{it}^{-20}, \tag{1}$$

$$p_{2,it} \equiv \Pr_{it}(0.8 \le R_{t+1} < 1) = 1 - p_{it}^{+} - p_{it}^{-20}, \tag{2}$$

$$p_{3,it} \equiv \Pr_{it}(1 \le R_{t+1} < 1.2) = p_{it}^{+} - p_{it}^{20}, \tag{3}$$

$$p_{4,it} \equiv \Pr_{it}(R_{t+1} \ge 1.2) = p_{it}^{20}. \tag{4}$$

Here R_{t+1} denotes the gross one-year return, so the four regions correspond to a return below -20 percent, a return between -20 and zero, a return between zero and 20 percent, and a return above 20 percent. We drop observations for which the answers imply a negative probability for any of the four regions. For example, the response $(p^{+}, p^{20}, p^{-20}) = (0.65, 0.15, 0.10)$ implies probabilities $(0.10, 0.25, 0.50, 0.15)$ across the four bins. Appendix A reports the wording of the survey questions and further details on this mapping.

From subjective probabilities to expected returns. Several of our empirical exercises require a scalar measure of expected returns. We construct this measure from the four probabilities in equation (4). Following von Gaudecker and Wogrolly (2022), let x_j denote the representative gross return associated with bin j. We define household i's subjective expected return in wave t as

$$m_{it} \equiv \mathbb{E}_{it}[R_{t+1}] = \sum_{j=1}^{4} p_{j,it} x_j. \tag{5}$$

The representative values x_j are chosen using the historical distribution of stock-market returns, as described in Appendix A. This conversion gives us a common scalar belief measure while retaining the information contained in the probability questions.

The conversion in equation (5) is needed only for exercises that require expected returns or forecast errors. Our fractional-logit analysis in Section 2.3 instead uses the four reported probabilities directly. That exercise is therefore informative about the wealth–belief relationship without relying on the particular mapping from probability bins into a scalar expected return.

The estimation samples differ across exercises because the fixed-effect, wealth, and portfolio regressions require different combinations of nonmissing and internally consistent responses. We therefore report the number of observations and individuals for each specification rather than imposing a common sample throughout. As an additional panel-quality restriction, our persistence exercises retain respondents with sufficiently repeated observations; Appendix A shows how the number of households varies with alternative minimum-observation requirements.

Wealth and portfolio variables. Throughout the RAND analysis, household wealth is the sum of housing wealth and retirement-account holdings, and the risky

portfolio share is the share of retirement-account wealth invested in stocks. Both are reported directly by respondents. RAND does not record directly held stocks and bonds, which is one reason we complement it with the Survey of Consumer Finances in Section 6.1.

2.2. Persistent Disagreement

We begin by asking whether disagreement primarily reflects common fluctuations in beliefs or persistent differences across households. This distinction matters for the mechanism developed below. Common movements in expectations are naturally relevant for aggregate fluctuations. Persistent differences in expected returns, in contrast, can generate persistent differences in portfolio exposure and wealth accumulation even when households face the same aggregate market history.

Fixed-effect decomposition. Following Giglio et al. (2021), we decompose expected returns into an individual component and a common time component. Specifically, we estimate variants of

$$m_{it} = \alpha_i + \mu_t + \varepsilon_{it}, \tag{6}$$

where α_i is an individual fixed effect and μ_t is a survey-wave fixed effect. We estimate equation (6) with time effects only, individual effects only, and both sets of effects. The comparison of the resulting R^2 values provides a simple measure of the relative importance of common time variation and persistent individual heterogeneity.

Table 1 Panel A reports the results. Time fixed effects explain 12.8 percent of the variation in expected returns. Individual fixed effects explain 57.3 percent. Including both sets of fixed effects raises the R^2 only slightly further, to 57.8 percent. The estimation sample contains 19,956 household-wave observations from 1,120 individuals.

The comparison is quantitatively stark. Expected returns do move over time, but most of the variation in the panel is associated with stable differences across households. A household that is relatively optimistic in one survey wave therefore tends to remain relatively optimistic in subsequent waves. This finding closely mirrors the persistent individual heterogeneity documented by Giglio et al. (2021) among Vanguard investors, despite the substantial differences between the two samples.

What explains the individual component? The dominance of individual fixed effects raises a second question: are these persistent differences simply a reflection of persistent observable household characteristics? To examine this, we recover $\widehat{\alpha}_i$ from the specification with individual and time fixed effects and estimate the cross-sectional regression

$$\widehat{\alpha}_i = a + \beta \log \overline{W}_i + \Gamma' X_i + u_i, \tag{7}$$

where $\overline{W}_i$ is household i's average wealth over its RAND observations and X_i contains the available household controls. The controls include income, age, homeownership, retirement-account status, and housing wealth.

We use average rather than contemporaneous wealth in equation (7) because the dependent variable is itself a household-level, time-invariant object. Pairing the same estimated fixed effect with repeated observations of contemporaneous wealth would mechanically create multiple observations for a single outcome and would give transitory wealth fluctuations a role in explaining a persistent belief component. Averaging wealth over the same panel instead provides a household-level counterpart to $\widehat{\alpha}_i$ and reduces noise from short-run wealth movements. This choice is purely descriptive: because beliefs can themselves affect portfolios and wealth accumulation, β should not be interpreted causally.

Panel B of Table 1 reports the explanatory power of these cross-sectional regressions. Wealth alone explains 8.9 percent of the variation in $\widehat{\alpha}_i$. The remaining observables explain 6.6 percent. Adding wealth to these controls increases the R^2 by only 2.7 percentage points. As a benchmark, the corresponding cross-sectional regressions using the expected-return measure itself deliver almost identical numbers: wealth alone explains 9.0 percent, the controls explain 6.4 percent, and adding wealth raises explanatory power by 2.6 percentage points.

These results closely reproduce the message of Giglio et al. (2021). Persistent disagreement is not well summarized by observable household types: wealth leaves more than 90 percent of the cross-sectional variation in the estimated individual component unexplained. This is the relevant sense in which the R^2 values are small. It does not follow, however, that the joint distribution of wealth and beliefs lacks structure. A variable can explain little of household-level variation when there is substantial dispersion around the conditional mean, while the conditional mean itself varies strongly and systematically with that variable. Section 2.3 shows that RAND exhibits exactly this pattern.

2.3. Wealth, Portfolio Exposure, and the Distribution of Beliefs

We next characterize how disagreement is organized across the wealth distribution. The distinction between this exercise and the regressions in Table 1 is important. The R^2 asks how much of the variation in individual beliefs can be explained by wealth. The binscatters instead trace the conditional mean of beliefs across the wealth distribution, averaging over much of the residual household-level heterogeneity. The first object is small; the second is highly structured. In RAND, once households are grouped by wealth percentiles, the correlation between percentile-average wealth and percentile-average expected returns is extremely high, and the corresponding relationship with forecast errors is strongly negative. Thus, the low individual-level R^2 and the steep binscatter relationship are two sides of the same empirical fact: there is substantial disagreement within wealth groups, but the center of the belief distribution shifts systematically with wealth.

TABLE 1. Persistent Heterogeneity in RAND Expected Returns

Panel A. Fixed-effect decomposition			
Dependent variable: m_{it}	Time FE	Individual FE	Time + Individual FE
R^2	0.128	0.573	0.578
Observations	19,956	19,956	19,956
Individuals	1,120	1,120	1,120
Panel B. Cross-sectional explanatory power			
Dependent variable	Wealth only	Controls only	Incremental R^2 from wealth
$\widehat{\alpha}_i$	0.089	0.066	0.027
Expected return	0.090	0.064	0.026
Observations	745	745	745

Notes. Panel A reports the R^2 from three versions of equation (6): survey-wave fixed effects only, individual fixed effects only, and both. Panel B reports the R^2 from cross-sectional regressions on wealth and household controls. For the first row, the dependent variable is the estimated individual effect from the specification with both individual and time fixed effects. Wealth is average household wealth over the respondent's RAND observations. Controls include income, age, homeownership, retirement-account status, and housing wealth. The final column reports the increase in R^2 from adding wealth to the controls-only specification rather than the total R^2 of the combined specification.

Expected returns and forecast errors across wealth. Figure 1 Panel A plots the scalar expected-return measure in equation (5) against household wealth, defined as in Section 2.1, using RAND observations only. Each point averages observations within a wealth percentile. Expected stock returns rise systematically with wealth, and the percentile means lie close to a monotone schedule. Across the wealth distribution, the difference is economically meaningful—roughly four percentage points between poorer and richer households. The tight relationship among percentile means should not be interpreted as implying that wealth predicts individual beliefs precisely. Rather, averaging within percentiles filters out the large residual dispersion documented in Section 2.2 and reveals the systematic component of the wealth–belief relationship.

Panel B examines whether this optimism is associated with larger or smaller forecast errors. For each observation, we match the survey forecast to the realized one-year market return over the corresponding forecast horizon and define

$$FE_{it} = R_{t+1} - m_{it}. \tag{8}$$

With this sign convention, a positive forecast error means that the household was too pessimistic and a negative forecast error means that it was too optimistic. Forecast errors decline with wealth. Poorer households therefore tend to be more pessimistic relative to the market outcomes realized over our sample, whereas the expectations of wealthier households are closer, on average, to realized returns.

This accuracy statement is sample-specific. The realized one-year returns corresponding to the RAND observations average 14.88 percent, with a minimum of 1.38

percent and a maximum of 32.33 percent in the sample used for the probability-distribution exercise below. Because realized returns were predominantly positive over these forecast horizons, shifting probability mass toward positive outcomes improves average forecast accuracy. The result should therefore not be read as a general claim that optimism is intrinsically more accurate.

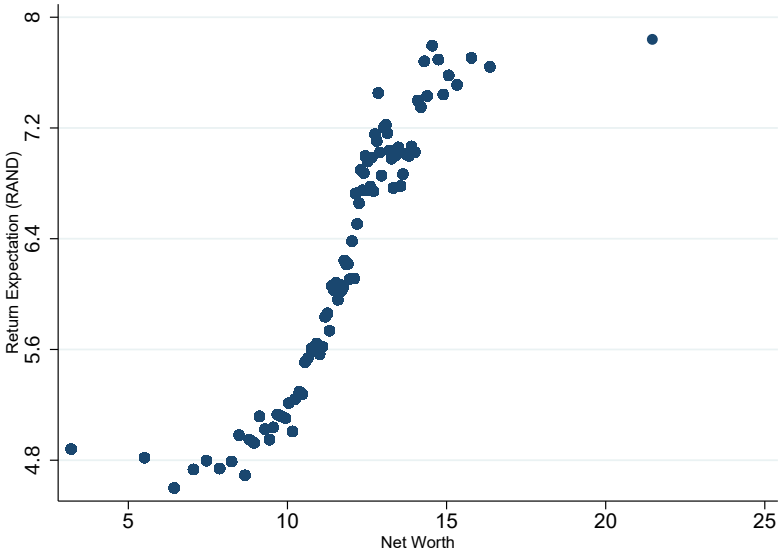

(A) Expected returns and wealth.

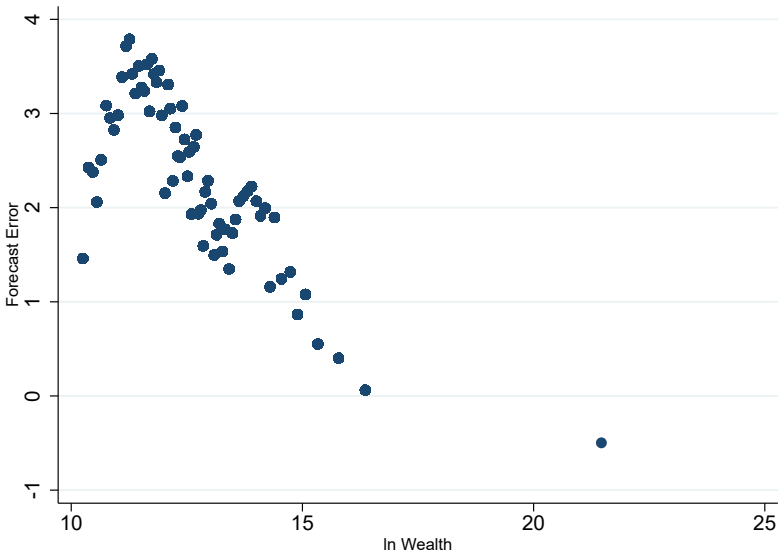

(B) Forecast errors and wealth.

FIGURE 1. **Beliefs and wealth in RAND.** Panel A plots the binned relationship between household wealth and the scalar expected-return measure constructed from the RAND probability questions. Panel B plots the corresponding relationship for the one-year forecast error, defined as realized minus expected return. In both panels, the horizontal axis is the logarithm of household wealth, defined in Section 2.1 as the sum of housing and retirement-account wealth, and each point averages observations within a wealth percentile. Both panels use RAND observations only.

The underlying probability distribution. The scalar measure m_{it} is convenient for the fixed-effect decomposition and for computing forecast errors, but it necessarily summarizes a coarse subjective distribution with a single number. We therefore return to the original probability responses and ask how wealth changes the entire reported distribution. This exercise has two advantages. First, it shows which part of the distribution generates the positive wealth gradient in expected returns. Second, it does not depend on the representative bin values x_j used in equation (5).

For each of the four return regions j, we estimate a separate fractional-logit specification,

$$p_{j,it} = \Lambda\Big(\alpha_j + \beta_{W,j} \log W_{it} + \beta_{V,j} V_{it} + \Gamma_j' X_{it}\Big), \tag{9}$$

where $\Lambda(z) = \exp(z)/(1+\exp(z))$, W_{it} is household wealth, V_{it} is the household's risky portfolio share, X_{it} contains homeownership, retirement-account status, housing wealth, income, and age. We estimate the four equations separately because the dependent variables are probabilities bounded between zero and one and because the economic question concerns how wealth and portfolio exposure reallocate probability mass across distinct regions of the subjective return distribution.

The coefficients of a fractional logit are not directly interpretable as changes in probabilities. We therefore report average changes in predicted probabilities under two counterfactuals. For each household-wave observation, we first compute the predicted probability at the observed covariates. We then change one variable while leaving all other covariates unchanged and compute the prediction again. For counterfactual c,

$$\Delta \widehat{p}^{\,c}_{j,it} = \widehat{p}_j(X^c_{it}) - \widehat{p}_j(X_{it}), \tag{10}$$

and Table 2 reports

$$\overline{\Delta \widehat{p}}^{\,c}_{j} = \frac{1}{N} \sum_{i,t} \Delta \widehat{p}^{\,c}_{j,it}. \tag{11}$$

The first counterfactual replaces W_{it} with $2W_{it}$ for each observation. Thus, "wealth $\times 2$" is the average change in the predicted probability when the same household is assigned twice its observed wealth, holding the other covariates fixed. The second counterfactual raises the risky portfolio share by 10 percentage points—for example, from 30 to 40 percent—again leaving the remaining covariates unchanged.

Table 2 reports the results. Doubling wealth has essentially no effect on the probability assigned to a loss below −20 percent. Instead, it shifts probability mass from moderate losses to moderate gains: the predicted probability of a return between −20 and zero falls by 1.37 percentage points, while the probability of a return between zero and 20 percent rises by 1.40 percentage points. The probability assigned to returns above 20 percent increases by a much smaller 0.15 percentage points. Thus, the positive association between wealth and expected returns does not arise mainly because richer households place substantially more probability on extreme booms. It reflects a shift in the center of the subjective distribution, from moderately negative to moderately positive outcomes.

The risky-share counterfactual produces a closely related pattern. A 10 percentage point increase in risky exposure reduces the predicted probability of a moderate

TABLE 2. Wealth, Portfolio Exposure, and Subjective Return Probabilities

	$R < -20\%$	$-20\% < R < 0\%$	$0\% < R < 20\%$	$R > 20\%$
Wealth ×2	0.01	-1.37^{***}	1.40^{***}	0.15^{**}
Risky share +10 pp	0.08	-1.66^{*}	1.35^{***}	0.33^{***}

Notes. The table reports average changes in predicted probabilities, in percentage points, from four separate fractional-logit regressions of the form in equation (9). "Wealth ×2" replaces each household's observed wealth with twice that value while holding the remaining covariates fixed. "Risky share +10 pp" raises the risky portfolio share by 10 percentage points. Controls are homeownership, retirement-account status, housing wealth, income, and age. The sample contains 12,025 observations. Realized one-year returns over the corresponding sample have mean 14.88 percent, minimum 1.38 percent, and maximum 32.33 percent. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

loss by 1.66 percentage points and raises the probability of a moderate gain by 1.35 percentage points. It also raises the probability assigned to returns above 20 percent by 0.33 percentage points. Households with greater risky exposure are therefore systematically more optimistic in much the same way as wealthier households.

The results in this subsection are conditional associations, not causal estimates. In particular, the relationship between beliefs and risky portfolio exposure can run in both directions: beliefs can affect portfolio choice, while portfolio outcomes can subsequently affect beliefs. This endogeneity is central to the mechanism developed in the model rather than something the regressions are intended to resolve.

The empirical evidence therefore imposes two simultaneous requirements on the theory. First, it must reproduce the large persistent dispersion in beliefs emphasized by Giglio et al. (2021), including the fact that wealth and standard observables explain only a small share of that dispersion. Second, it must reproduce the strong structure in the conditional mean: after averaging within wealth percentiles, optimism rises sharply with wealth, forecast errors fall, and the underlying subjective distribution shifts from moderate losses toward moderate gains. A successful model must therefore generate both a low micro-level R^2 and a steep wealth–belief schedule. The learning mechanism introduced in the next section is designed to account for these two features jointly.

3. MODEL

We develop a heterogeneous-agent model that connects the facts in Section 2 to the distribution of wealth. The environment is deliberately standard except for beliefs. Households face idiosyncratic labor-income risk, choose how much to save, and allocate their financial wealth between a safe asset and risky assets. The risky return contains both an aggregate stock-market component and an idiosyncratic component, consistent with the substantial within-asset-class return heterogeneity documented in household data (Fagereng et al. (2020); Hubmer, Krusell, and Smith (2021)). Aggregate asset prices are determined in general equilibrium.

We depart from the standard model in two related dimensions. First, households have imperfect knowledge of the equilibrium process governing asset prices.

Following the Internal Rationality approach of Adam and Marcet (2011), they nevertheless make dynamically optimal choices under a well-defined subjective probability measure. Second, rather than assigning households permanently optimistic or pessimistic beliefs, all households start from the same prior and learn about future stock returns from their own investment experience. Because households hold different portfolios and receive different idiosyncratic risky returns, they experience the same aggregate market through different personal return histories. These experiences feed back into beliefs and subsequent portfolio choices.

The distinction between these two ingredients is important. Internal Rationality specifies what households know and how they make choices under imperfect market knowledge. Learning from experience specifies how the belief state entering that subjective probability measure evolves over time.

3.1. Model structure

Households. Time is discrete and infinite. There is a unit mass of households indexed by i. Household i chooses consumption c_t^i, end-of-period financial wealth A_t^i, and the share v_t^i of financial wealth allocated to risky assets. Its problem is

$$\max_{\{c_t^i, A_t^i, v_t^i\}_{t=0}^{\infty}} \mathbb{E}_0^{\mathscr{P}^i} \sum_{t=0}^{\infty} (\beta^i)^t \frac{(c_t^i)^{1-\gamma^i}}{1-\gamma^i}, \tag{12}$$

subject to

$$c_t^i + A_t^i = W_t^i + R_t^{A,i}(v_{t-1}^i) A_{t-1}^i, \tag{13}$$

and

$$A_t^i \geq \underline{a} D_t > 0, \qquad \underline{v} \leq v_t^i \leq \bar{v} \tag{14}$$

Here W_t^i denotes labor income and $\mathscr{P}^i$ is household i's subjective probability measure, described in Section 3.2.

The small positive minimum account balance $\underline{a} D_t$ and minimum equity share $\underline{v}$ ensure that every household receives a personal risky-return observation. They formalize participation in the intermediated investment arrangement below; neither bound is intended to drive the wealth distribution. The restrictions also rule out negative net worth and short positions.

In the quantitative model, (β^i, γ^i) vary over a finite support of household types that also differ in labor income (Section 6.2). Types are highly persistent but not permanent; because transitions between them are rare, we suppress the time dependence of preferences in the notation. This heterogeneity provides variation in saving and risk taking that is independent of beliefs. Importantly, beliefs are not assigned by type. Households of different types start from the same prior about stock returns, and their subsequent beliefs evolve separately through experience.

Financial wealth consists of risky and safe positions,

$$A_t^i = P_t S_t^i + B_t^i, \tag{15}$$

where P_t is the ex-dividend stock-market price, S_t^i denotes risky asset holdings, and

B_t^i is the market value of the position in the one-period safe asset. The risky share is therefore

$$v_t^i \equiv \frac{P_t S_t^i}{A_t^i}, \tag{16}$$

whenever $A_t^i > 0$.

This representation separates the two margins through which beliefs affect wealth accumulation. A_t^i determines how much the household carries forward, while v_t^i determines the exposure of those savings to risky returns. Households place dollar investment orders through an intermediary: they choose (A_t^i, v_t^i) but do not directly choose the quantities of underlying stock shares S_t^i. The intermediary executes these orders at the market-clearing price P_t and keeps the underlying quantities and price from the household's information set. The intermediary's conversion from dollars into asset quantities is specified below.

Aggregate and idiosyncratic risky returns. The aggregate risky asset is a claim to a dividend stream D_t. Dividend growth is

$$g_{t+1} \equiv \frac{D_{t+1}}{D_t} = \bar{g}_D \exp\left(-\frac{1}{2}\sigma_D^2 + \sigma_D \varepsilon_{t+1}^D\right), \qquad \varepsilon_{t+1}^D \sim \mathcal{N}(0,1). \tag{17}$$

The gross return on the aggregate stock market is

$$R_{t+1}^S = \frac{P_{t+1} + D_{t+1}}{P_t}. \tag{18}$$

Individual households need not earn exactly this return on the risky component of their portfolios. We allow for household-specific investment performance,

$$R_{t+1}^{S,i} = R_{t+1}^S \eta_{t+1}^i, \tag{19}$$

where

$$\log \eta_{t+1}^i \overset{i.i.d.}{\sim} \mathcal{N}\left(-\frac{1}{2}\sigma_\eta^2, \sigma_\eta^2\right). \tag{20}$$

Thus

$$\mathbb{E}[\eta_{t+1}^i] = 1.$$

The shock η_{t+1}^i is independent across households and over time and independent of aggregate dividend and labor-income innovations.

This component captures the empirical fact that households earn substantially different returns even within narrow asset classes. Such differences can reflect individual asset selection, concentration, diversification, timing, investment skill, or access to heterogeneous risky opportunities. We model their consequences without taking a stand on which of these underlying forces is responsible for the dispersion.

The distinction between R_{t+1}^S and $R_{t+1}^{S,i}$ will also be important for the learning mechanism. R_{t+1}^S is the common market return, which enters the objective payoff but is not separately reported to households. $R_{t+1}^{S,i}$ is the risky return actually experienced

and reported by the intermediary to household i. Hence two households exposed to the same aggregate market can experience different investment outcomes. The individual performance component is implemented within intermediaries; with mean-one independent shocks and a continuum of households, its aggregate payoff wedge vanishes, while idiosyncratic net investment returns remain heterogeneous.

The mean-one restriction ensures that idiosyncratic return risk does not introduce a separate aggregate growth factor. With a continuum of households and innovations independent of predetermined portfolio positions, the idiosyncratic component washes out in the aggregate by a cross-sectional law of large numbers. It can nevertheless have large distributional consequences because it generates independent multiplicative wealth histories.

Let Q_t denote the price of a claim to one unit of consumption next period. The safe gross return is

$$R^B_{t+1} = \frac{1}{Q_t}. \tag{21}$$

The gross return on household i's financial wealth is therefore

$$R^{A,i}_{t+1}(v^i_t) = v^i_t R^{S,i}_{t+1} + (1 - v^i_t) R^B_{t+1}. \tag{22}$$

Combining equations (19) and (22),

$$R^{A,i}_{t+1}(v^i_t) = v^i_t R^S_{t+1} \eta^i_{t+1} + (1 - v^i_t) R^B_{t+1}. \tag{23}$$

The decomposition makes clear that heterogeneous portfolio returns have two sources. First, households choose different risky shares v^i_t. Second, conditional on the same risky share, households receive different idiosyncratic risky returns η^i_{t+1}. Both sources generate heterogeneous investment experiences and therefore matter for the learning mechanism below.

Labor income. Households also face uninsurable idiosyncratic labor-income risk. In the quantitative model,

$$W^i_t = \bar{w} D_t \varepsilon(\tau^i_t), \qquad \tau^i_t \in \{1, \ldots, 6\}, \tag{24}$$

with

$$\tau^i_{t+1} \sim \Pi(\tau^i_t, \cdot). \tag{25}$$

Labor income is proportional to the aggregate dividend, so that wages and dividends share a common stochastic trend and the economy is stationary in dividend-normalized units. The household type τ^i_t, which also indexes preferences, follows a Markov chain that is common to all households (Section 6.2). Types are highly persistent, but there is no permanent household-specific wage type: every household can move between types.

This separation is useful for identification of the mechanism. Differences in expected stock returns are not mechanically attached to differences in labor earnings or preferences: no type carries a belief. Instead, persistent disagreement must emerge

from households' investment histories. Labor-income risk retains the standard incomplete-markets motive for precautionary saving and also produces variation in wealth and portfolio choices that can affect subsequent investment experience.

Asset markets and equilibrium prices. There is one unit of the aggregate risky asset. In addition to the risky Lucas tree, the economy features a productive one-period safe asset. At date t, households can purchase $\bar{B}_t = \bar{b} D_t > 0$ units of safe trees at price Q_t per unit. Each safe tree produces one unit of the consumption good with certainty at date $t+1$ and then disappears, implying a gross safe return

$$R^B_{t+1} = \frac{1}{Q_t}.$$

The quantity of safe trees available is exogenous and scales with aggregate dividends, while Q_t adjusts in equilibrium so that households' desired holdings equal their available supply.[8]

Asset-market clearing requires

$$\int S^i_t \, di = \int \frac{v^i_t A^i_t}{P_t} \, di = 1, \tag{26}$$

$$\int B^i_t \, di = \int (1 - v^i_t) A^i_t \, di = Q_t \bar{B}_t. \tag{27}$$

The two prices (P_t, Q_t) adjust so that these conditions hold.

Intermediated trading and information. The intermediary accepts household instructions for $v^i_t A^i_t$ units of expenditure on risky assets and $(1-v^i_t)A^i_t$ on safe assets. At the market-clearing prices, the corresponding orders in *physical* units are

$$S^i_t = \frac{v^i_t A^i_t}{P_t}, \qquad N^{B,i}_t = \frac{(1-v^i_t)A^i_t}{Q_t}, \qquad B^i_t = Q_t N^{B,i}_t. \tag{28}$$

Households receive statements of their own realized risky returns $R^{S,i}_t$ and safe returns R^B_t, and know their aggregate account balances and portfolio weights. They do not observe underlying stock prices, market-index returns, individual market-share holdings, or the idiosyncratic shock η^i_t separately. The intermediary does not report both an underlying share quantity and its dollar valuation, which would reveal P_t. Households have perfect recall of the information actually delivered to them; this is an information restriction, not forgetfulness of previously observed stock prices.

[8]The positive supply of the safe asset represents productive capital rather than borrowing by another household, firm, or government. Expenditure $Q_t \bar{B}_t$ on newly planted safe trees is a use of current resources, while trees planted in the previous period generate $\bar{B}_{t-1}$ units of consumption at t. Hence, no government budget constraint, debt rollover, or transfers are required. A positive supply is also necessary for the aggregate household portfolio to contain safe assets, as it does in the data (Table 8); with zero net supply, risky assets would mechanically constitute all aggregate financial wealth.

At date t, the intermediary quotes the price Q_t of a new one-period safe claim, equivalently its *guaranteed gross return* $R^B_{t+1} = 1/Q_t$. Households can therefore recover and record the history of safe-asset prices. The return on inherited safe positions is $R^B_t = 1/Q_{t-1}$, whereas $1/Q_t$ is the known return on claims purchased at t. Unlike the safe-asset price, the underlying *stock* price and market return are withheld.

Hence households are price takers, but asset prices are fully endogenous. They depend on the joint cross-sectional distribution of wealth, income, portfolio positions, preferences, and beliefs. The informational assumption introduced next concerns households' knowledge of this equilibrium mapping; it does not remove these general-equilibrium feedbacks.

3.2. Internal Rationality

The full-information rational-expectations benchmark combines two logically distinct assumptions. First, households make optimal intertemporal choices given their information. Second, they know the correct probability law governing all payoff-relevant variables, including equilibrium asset prices. We retain the first requirement and relax the second.

Following Adam and Marcet (2011), households are *internally rational*: they maximize expected utility under a coherent subjective probability measure without necessarily knowing the objective equilibrium law of asset returns. We distinguish the subjective probability law, the information available to the household, and the decisions made conditional on that information.

Subjective probability law. Formally, let $(\Omega, \mathscr{B}, \mathscr{P}^i)$ denote household i's subjective probability space. The sample space Ω contains possible histories of primitive uncertainty: the latent expected-log-return state, aggregate and idiosyncratic return innovations, dividend and safe-asset price innovations, and household-type transitions. The sigma-algebra $\mathscr{B}$ specifies measurable events, while $\mathscr{P}^i$ assigns subjective probabilities to them. These payoff-relevant processes are taken as stochastic and outside the control of an individual price-taking household, even though aggregate asset prices are determined in general equilibrium. Consumption, saving, and portfolio allocations are instead choices contingent on the information received.

Conditional on a common Gaussian initial prior $Z^i_{-1} \sim \mathcal{N}(\zeta_0, P_0^-)$, the subjective aggregate-return and latent-state processes are

$$\log R^S_t = Z^i_{t-1} + \varepsilon^{S,i}_t, \tag{29}$$

$$Z^i_t = (1-\rho)\bar{Z} + \rho Z^i_{t-1} + \varepsilon^{Z,i}_t, \qquad 0 < \rho < 1, \tag{30}$$

where $\varepsilon^{S,i}_t \sim \mathcal{N}(0, \sigma_S^2)$ and $\varepsilon^{Z,i}_t \sim \mathcal{N}(0, \sigma_Z^2)$. The individual risky return remains

$$\log R^{S,i}_t = \log R^S_t + \log \eta^i_t, \qquad \log \eta^i_t \sim \mathcal{N}(-\tfrac{1}{2}\sigma^2_\eta, \sigma^2_\eta), \tag{31}$$

with the same marginal idiosyncratic-return law as under the objective measure.

Under $\mathscr{P}^i$, the shocks $\varepsilon_t^{S,i}$, $\varepsilon_t^{Z,i}$, and the centered $\log\eta_t^i + \sigma_\eta^2/2$ are mutually independent, independent over time, and independent of the initial latent state. Thus $\mathrm{Cov}^{\mathscr{P}^i}(\varepsilon_t^{S,i}, \log\eta_t^i) = 0$, $\mathrm{Cov}^{\mathscr{P}^i}(\varepsilon_t^{Z,i}, \log\eta_t^i) = 0$, and $\mathrm{Cov}^{\mathscr{P}^i}(\varepsilon_t^{Z,i}, \varepsilon_t^{S,i}) = 0$. The market-return innovation is common within the subjective aggregate process; the household-specific η^i innovations are independent across investors. Different investors condition their subjective laws on different personal histories, so their posterior forecasts can differ even though the aggregate market return is a single objective variable.

Households retain the objectively specified lognormal dividend-growth law (17) and the type transition law (25); W_t^i follows (24). To close the subjective model of the safe asset, households specify a Markov law for the observed safe-asset price Q_t, equivalently for the gross return $1/Q_t$ it promises:

$$\log\left(\frac{1}{Q_{t+1}}\right) = \bar{r}_B + \varphi_B\left(\log\left(\frac{1}{Q_t}\right) - \bar{r}_B\right) + \sigma_B\varepsilon_{t+1}^{B,i}, \qquad |\varphi_B| < 1, \quad \sigma_B > 0. \tag{32}$$

Conditional on the observed Q_t, the return $1/Q_t$ on a safe claim purchased today is known; only prices, and hence returns, offered on later purchases are uncertain. The innovations ε_t^D and $\varepsilon_t^{B,i}$ are independent standard normals and are independent of the latent return state, the aggregate return and idiosyncratic shocks, and the Markov type transitions. The initial dividend, safe-asset price and type are independent of Z_{-1}^i. In particular, the subjective covariances between dividend or safe-price news and the latent market-return state are zero; observing D_t or Q_t does not supply an additional measurement of Z_{t-1}^i. These independence assumptions are maintained for the analytical filter below, not asserted as properties of objective general equilibrium, where stock prices, dividends and safe rates may comove.

Equations (29)–(32), together with the wage and dividend laws, initial distributions and independence restrictions, specify the household's joint subjective law for future payoff-relevant uncertainty. The law for Q_t preserves the contractual identity $R_{t+1}^B = 1/Q_t$. Households need not know or forecast the equilibrium mapping from the cross-sectional distribution to (P_t, Q_t); their subjective laws for personal risky returns and future safe-asset prices need not coincide with the objective laws generated by market clearing.

Information filtration. Let $\mathscr{I}_t^i \subseteq \mathscr{B}$ denote the information available to household i immediately before its date-t choices. It is the filtration generated by the observed histories and the household's past decisions:

$$\mathscr{I}_t^i = \sigma\Big(A_{-1}^i, v_{-1}^i, \{W_s^i, D_s, R_s^{S,i}, Q_s, \tau_s^i\}_{s\le t}, \{c_s^i, A_s^i, v_s^i\}_{s<t}\Big). \tag{33}$$

The initial asset positions and the household's current type (which indexes preferences and income) are known to the household. Observations dated t arrive before the period-t consumption and portfolio decision. In particular, the household observes Q_t and therefore knows the promised return $R_{t+1}^B = 1/Q_t$ on new safe claims; the realized return on inherited claims is $R_t^B = 1/Q_{t-1}$. In contrast, the intermediary does not provide the underlying stock-price history, share quantities, or aggregate

market returns. The household observes $R_t^{S,i}$ but cannot separately identify R_t^S and η_t^i; under its subjective law, dividends, safe-asset prices, and account statements do not reveal the missing aggregate return. The restriction concerns the entire history of observations, not the household's recall of previously supplied prices.

Optimal decisions. Given $\mathscr{I}_t^i$, household i chooses (c_t^i, A_t^i, v_t^i) to solve (12) subject to (13)–(14). These choices are $\mathscr{I}_t^i$-measurable; unlike the primitive shocks, they are determined by the household's policy. Expectations of future outcomes are conditional expectations under $\mathscr{P}^i$ given $\mathscr{I}_t^i$. Optimality uses the *joint* subjective law, including the correlation between next period's wealth and the personal-return signal that updates beliefs. In this sense, households are rational conditional on their information and subjective model, without knowing the objective general-equilibrium price process.

The corresponding correct-beliefs benchmark is rational expectations *conditional on this restricted information set*. A full-information rational-expectations economy, in which agents also observe stock prices, has a different information structure and need not be nested literally. Section 5 explains why neither subjective learning nor the two-stage computation requires households to forecast the law of motion of the cross-sectional wealth distribution.

3.3. Learning from own portfolio experience

Internal Rationality describes how households make choices under imperfect market knowledge, but it leaves open a central economic question: why do their beliefs differ in the first place? One possibility would be to endow households with permanently different expected returns. We use precisely that reduced-form representation in Section 4 to isolate analytically the consequences of persistent disagreement. As a description of the origin of beliefs, however, it is incomplete. Taken literally, some households would simply be born optimistic and subsequently become rich partly because of that optimism.

Our quantitative model instead generates disagreement endogenously from investment experience. All households begin from the same prior. They do not observe a complete record of aggregate stock-market returns; instead, the intermediary reports their own risky investment performance. Because personal risky returns depend on the common market return and on idiosyncratic performance, households exposed to the same aggregate market accumulate different histories of investment outcomes. A household can rationally interpret its own return as noisy information about the market because it does not separately observe the two components.

This idea is related to the literature on experience-based learning. Malmendier and Nagel (2011), for example, show that investors put disproportionate weight on stock-market returns experienced during their lifetimes. In those models, heterogeneity largely reflects differences in the aggregate histories experienced by different cohorts. Our mechanism operates within a cohort and even within the same calendar period: households exposed to exactly the same aggregate market receive different private signals because η_t^i differs across investors.

Recent evidence also points more directly toward personal investment experience as an input into expectations. Jiang et al. (2025) show that investors' recalled own investment outcomes are strongly related to beliefs about future market returns and to subsequent trading behavior. Our model does not attempt to model memory retrieval itself. Rather, its restricted information set captures the economically relevant fact that personal investment outcomes, rather than an observed aggregate benchmark, are used to infer future market performance.

Personal performance as a Bayesian measurement. For each participant with $A^i_{t-1} > 0$ and $v^i_{t-1} > 0$, the intermediary reports the realized risky return $R^{S,i}_t$. Equivalently, the household could recover this return from its total account performance and the known safe return:

$$R^{S,i}_t = \frac{R^{A,i}_t - (1 - v^i_{t-1}) R^B_t}{v^i_{t-1}}. \tag{34}$$

Although portfolio exposure determines the magnitude of gains and losses, this exact recovery removes v^i_{t-1} from the precision of the risky return observation conditional on positive participation. We do not interpret lower equity exposure as a noisier observation of the same individual risky return. The distinction matters for rational filtering.

From the independent lognormal idiosyncratic return in (31), the correctly centered personal experience statistic is

$$\mathscr{S}^i_t \equiv \log R^{S,i}_t + \frac{\sigma^2_\eta}{2} = \log\left[\frac{R^{A,i}_t - (1 - v^i_{t-1}) R^B_t}{v^i_{t-1}}\right] + \frac{\sigma^2_\eta}{2}. \tag{35}$$

By (29), this gives the exact linear-Gaussian measurement equation

$$\mathscr{S}^i_t = Z^i_{t-1} + \nu^i_t, \quad \nu^i_t = \varepsilon^{S,i}_t + \log \eta^i_t + \frac{\sigma^2_\eta}{2}, \quad H \equiv \mathrm{Var}^{\mathscr{P}^i}(\nu^i_t) = \sigma^2_S + \sigma^2_\eta. \tag{36}$$

The private term is genuine noise, but it is also the source of cross-sectional disagreement. No aggregate-market-return observation is being dropped from a larger information set: the intermediary has never provided it. Nor does observing dividends, the safe-rate quote, or a personal account statement reconstruct the latent aggregate return under the specified subjective information structure.

Belief updating. Households update their beliefs using a steady-state Kalman filter. Let

$$\zeta^i_t \equiv \mathbb{E}^{\mathscr{P}^i}[Z^i_{t-1} \mid \mathscr{I}^i_{t-1}], \qquad P^- \equiv \mathrm{Var}^{\mathscr{P}^i}(Z^i_{t-1} \mid \mathscr{I}^i_{t-1}), \tag{37}$$

denote the prior mean and variance carried into date t. We initialize the prior variance at its stationary Kalman-filter value, so P^- is constant across households

and over time. All households start from the same prior,

$$\zeta_0^i = \zeta_0, \qquad P_0^{i,-} = P^- \quad \text{for all } i. \tag{38}$$

Upon observing the personal return signal $\mathscr{S}_t^i$, Bayesian updating gives

$$\widehat{Z}_{t-1|t}^i = (1-g_L)\zeta_t^i + g_L \mathscr{S}_t^i, \qquad g_L \equiv \frac{P^-}{P^- + H} \in (0,1), \tag{39}$$

where g_L is the constant Kalman gain. Combining this posterior with the subjective state transition (30) yields

$$\zeta_{t+1}^i = (1-\rho)\bar{Z} + \rho(1-g_L)\zeta_t^i + \rho g_L \mathscr{S}_t^i. \tag{40}$$

Thus $\zeta_{t+1}^i = \mathbb{E}^{\mathscr{P}^i}[Z_t^i \mid \mathscr{I}_t^i]$ is the household's forecast of the expected log market return over t to $t+1$. The realized personal return informs the previous latent state Z_{t-1}^i, which predicts Z_t^i through its persistence ρ.

The constant gain is consistent with exact Bayesian inference. In steady state, the posterior and prior variances satisfy the Riccati equations

$$P^+ = (1-g_L)P^-, \qquad P^- = \rho^2 P^+ + \sigma_Z^2. \tag{41}$$

Equivalently, for a chosen gain g_L, the subjective state innovation variance must satisfy

$$P^- = \frac{g_L H}{1-g_L}, \qquad \sigma_Z^2 = \frac{g_L H}{1-g_L}\left[1-\rho^2(1-g_L)\right]. \tag{42}$$

Since all households receive a signal with the same variance H each period, the filter remains at its stationary variance. Constant-gain learning is therefore an exact implication of the subjective linear-Gaussian model. The coefficient on past beliefs is $\rho(1-g_L)$, while the response to personal investment experience is ρg_L.

The state ζ_{t+1}^i forecasts a log return, not a gross return. Under the stationary filter, the conditional variance of the next log market return is $P^- + \sigma_S^2$. Hence the expected gross return is

$$m_t^i \equiv \mathbb{E}^{\mathscr{P}^i}[R_{t+1}^S \mid \mathscr{I}_t^i] = \exp\left(\zeta_{t+1}^i + \tfrac{1}{2}(P^- + \sigma_S^2)\right). \tag{43}$$

The individual risky-return distribution additionally includes σ_η^2. Since the subjective idiosyncratic multiplicative shock has mean one and is independent of aggregate returns, $\mathbb{E}^{\mathscr{P}^i}[R_{t+1}^{S,i} \mid \mathscr{I}_t^i] = m_t^i$. Idiosyncratic risk therefore increases the variance of personal payoffs without changing their conditional expected gross return.

Persistent individual disagreement. The mechanism generates persistent beliefs from identical priors without assigning permanent optimism types. To make this precise, condition on an aggregate stock-return history and let $\bar{\zeta}_t$ be the cross-sectional mean of prior beliefs. With an atomless population and independent mean-one idiosyncratic return shocks, the centered cross-sectional deviation $\tilde{\zeta}_t^i \equiv \zeta_t^i - \bar{\zeta}_t$

satisfies

$$\widetilde{\zeta}^i_{t+1} = \rho(1-g_L)\widetilde{\zeta}^i_t + \rho g_L\left(\log\eta^i_t + \tfrac{1}{2}\sigma^2_\eta\right). \tag{44}$$

Thus the common market history disappears from the cross-section, while personal investment luck generates persistent differences. Provided $0 < \rho(1-g_L) < 1$, the stationary conditional variance and the k-period autocorrelation are

$$\mathrm{Var}(\widetilde{\zeta}^i \mid \text{aggregate history}) = \frac{(\rho g_L)^2\sigma^2_\eta}{1-\rho^2(1-g_L)^2}, \qquad \mathrm{Corr}(\widetilde{\zeta}^i_{t+k}, \widetilde{\zeta}^i_t) = [\rho(1-g_L)]^k. \tag{45}$$

These expressions explain how a small gain and persistent expected-return state can generate the large individual components emphasized in Section 2.

Own returns and expected market returns. At fixed inherited wealth, portfolio weight, safe return, and prior belief, (34) and (40) imply

$$\frac{\partial\zeta^i_{t+1}}{\partial R^{A,i}_t} = \frac{\rho g_L}{R^{A,i}_t - (1-v^i_{t-1})R^B_t} > 0, \qquad A^i_{t-1} > 0, \quad v^i_{t-1} > 0. \tag{46}$$

Since the conditional predictive variance is common, the corresponding semi-elasticity in the gross expectation is positive as well:

$$\frac{\partial\log m^i_t}{\partial\log R^{A,i}_t} = \rho g_L \frac{R^{A,i}_t}{R^{A,i}_t - (1-v^i_{t-1})R^B_t} > 0. \tag{47}$$

Hence, the model predicts that good own portfolio returns therefore raise the forecast of aggregate stock returns even after conditioning on common market circumstances.

Personal luck raises both resources and optimism. Let $X^i_t = W^i_t + R^{A,i}_t A^i_{t-1}$ denote resources available before the date-t saving decision. Conditional on predetermined wealth, equity exposure and labor income, and the realized common market and safe returns, (23) and (40) imply

$$\mathrm{Cov}\left(X^i_t, \zeta^i_{t+1} \mid A^i_{t-1}, v^i_{t-1}, W^i_t, R^S_t, R^B_t, \mathscr{I}^i_{t-1}\right) = \rho g_L A^i_{t-1} v^i_{t-1} R^S_t \sigma^2_\eta > 0. \tag{48}$$

The identity follows from $\mathbb{E}[\eta^i_t] = 1$ and $\mathrm{Cov}(\eta^i_t, \log\eta^i_t) = \sigma^2_\eta$. The same favorable idiosyncratic investment outcome raises both contemporaneous resources and optimism. Greater equity exposure does not increase the *precision* of a perfectly reported personal risky return, but it magnifies the wealth consequence of the very shock that changes beliefs. This creates a channel from past beliefs, through risky exposure and realized wealth, to subsequent portfolio decisions and wealth accumulation.

Persistent pessimism. The mean-reverting subjective state provides a separate discipline for average optimism. Let $m_0 \equiv \mathbb{E}_0[\log R^S_t]$ denote the objective stationary average *log* market return. The signal is correctly centered with respect to the idiosyncratic shock, so under the objective law $\mathbb{E}_0[\mathscr{I}^i_t] = m_0$. Taking objective

expectations of (40) in a stationary equilibrium yields

$$\mathbb{E}_0[\zeta^i] = \frac{(1-\rho)\bar{Z} + \rho g_L m_0}{1-\rho(1-g_L)}. \tag{49}$$

Thus a perceived long-run anchor $\bar{Z} < m_0$ produces an objective mean posterior log-return forecast below m_0, even with an arbitrarily long investment history. The restriction is a subjective perception of long-run profitability, not a biased signal or a failure of Bayesian updating. There is nevertheless a direct result for gross forecasts. Holding the objective return process fixed and assuming the relevant exponential moments exist, the stationary recursion is affine in $\bar{Z}$, with coefficient $(1-\rho)/[1-\rho(1-g_L)]$. Consequently,

$$\frac{\partial \log \mathbb{E}_0[m_t^i]}{\partial \bar{Z}} = \frac{1-\rho}{1-\rho(1-g_L)} > 0. \tag{50}$$

A sufficiently pessimistic long-run anchor can therefore lower average *gross* expected returns below their realized objective mean, not merely lower the mean forecast of log returns.

The feedback from experience to wealth. Equations (23) and (40) generate a feedback between beliefs, portfolios, and wealth. Favorable personal return histories raise both wealth and optimism. The latter increases desired risky exposure, which amplifies the effects of subsequent idiosyncratic shocks on wealth. The filter is Bayesian, but the econometric origin of disagreement remains household-specific luck rather than a permanent optimism type. Greater exposure strengthens the *wealth–belief covariance* in (48), not the gain on an observed risky return.

Evidence on learning from own returns. The UBS–Gallup investor survey analyzed by Vissing-Jørgensen (2003) provides supportive evidence on the key empirical link in (46). Unlike RAND, the survey contains measures of investors' realized own-portfolio returns together with their expectations of future stock-market returns.

Past own-portfolio performance is strongly positively associated with expected market returns. In the raw micro data, the correlation between past own-portfolio returns and expected market returns is approximately 0.44; across bins of past portfolio performance, the correlation is approximately 0.88. Figure x represents this relationship.

To examine this more directly, we estimate

$$\log \mathbb{E}_{it}\left(R_{t+1}^m\right) = \alpha_t + \gamma_m \log R_{t-1}^{A,i} + X_{it}'\lambda_m + u_{it}, \tag{51}$$

where $R_{t-1}^{A,i}$ is respondent i's past own-portfolio return, R_{t+1}^m is the future market return, X_{it} contains income, savings, and age, and α_t denotes survey-wave fixed effects. Returns entering logarithms are expressed as gross returns. The mechanism predicts $\gamma_m > 0$: conditional on the common information available at a given date,

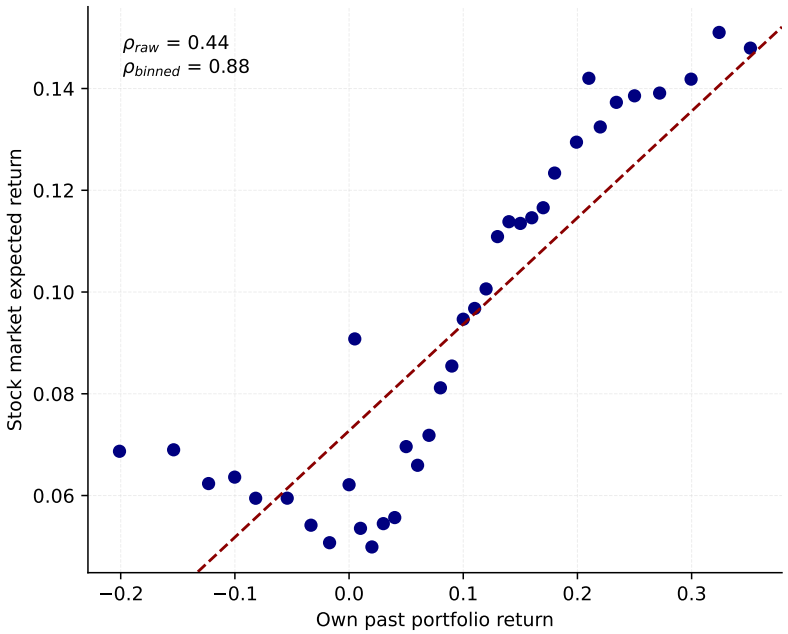

FIGURE 2. **Own past portfolio returns and expected stock market returns 1y ahead in UBS/Gallup Survey data.** The figure plots the binned relationship between the reported own portfolio return in the previous year and the scalar expected return measure for the stock market one year ahead.

investors with better personal portfolio returns should report higher expected market returns.

Column (3) is the most informative specification for our mechanism because survey-wave fixed effects absorb market conditions common to investors interviewed at the same time. The coefficient on past own-portfolio returns falls relative to the unconditional specification but remains strongly positive:

$$\widehat{\gamma}_m = 0.352, \qquad \mathrm{se}(\widehat{\gamma}_m) = 0.078.$$

Thus, within the same survey wave, investors who experienced better portfolio performance report substantially higher expectations of future market returns. We interpret this result as evidence for the direction of the experience channel rather than as a causal estimate of g_L or ρ. Realized portfolio returns are endogenous to previous portfolio choices and therefore to earlier beliefs.

3.4. Competitive equilibrium under Internal Rationality

We use the standard competitive-equilibrium concept, modified only to allow households to evaluate future outcomes under their subjective probability measures. Given preferences, the objective laws of motion for dividends, labor income, and idiosyncratic risky returns, initial asset positions, and the collection of subjective measures $\{\mathscr{P}^i\}_i$, a *competitive equilibrium under Internal Rationality* is a collection

TABLE 3. Own Portfolio Experience and Expected Market Returns

	(1)	(2)	(3)
$\log R^{A,i}_{t-1}$	0.7854***	0.7438***	0.3516***
	(0.0923)	(0.0882)	(0.0778)
Controls	No	Yes	Yes
Survey-wave fixed effects	No	No	Yes
R^2	0.067	0.083	0.168
Observations	26,607	26,607	26,607

Notes. The dependent variable is the log expected stock-market return. Controls are income, savings, and age. Standard errors, in parentheses, are clustered by survey wave. *** $p < 0.01$.

of household allocations and portfolio choices

$$\left\{c^i_t, A^i_t, v^i_t, S^i_t, B^i_t\right\}_{i,t\geq 0},$$

asset prices $\{P_t, Q_t\}_{t\geq 0}$, and belief states $\{\zeta^i_t\}_{i,t\geq 0}$ such that:

1. *Household optimality.* Taking the intermediary's quoted safe return and its own observed investment performance as given, and evaluating future uncertainty under $\mathscr{P}^i$, each household's choices solve (12) subject to (13)–(14) and the asset decomposition in (15).
2. *Belief updating.* Households update the personal-performance observation (36) using the Bayesian recursion (40), from the common prior in (38).
3. *Market clearing.* At every date and history, stock and bond markets clear according to (26) and (27).
4. *Return consistency.* Realized aggregate and individual asset returns satisfy (18)–(21), and the exogenous shocks follow their objective laws of motion.

This is a competitive equilibrium in allocations and prices with a restricted-information, internally rational household problem. Households submit dollar orders to intermediaries; their individually optimal choices aggregate to stock and bond demands that clear at objective market prices. The departure from full-information rational expectations is twofold: households observe personal returns rather than the market price history, and their subjective return laws need not coincide with the objective law generated by market clearing. Section 5 recasts this equilibrium recursively and makes precise the distinction between the individual information state and the objective state needed for price determination.

4. Analytical Results: Beliefs, Growth, and Wealth Tails

This section isolates the mechanism through which subjective stock return beliefs shape long-run wealth accumulation and the upper tail of the wealth distribution. We abstract from aggregate return risk to study the stationary cross-sectional wealth distribution directly. With aggregate risk, this distribution is itself random, which introduces analytical complications that are not essential to isolating the effect of

disagreement on cross-sectional wealth inequality. Aggregate risk is restored in the quantitative model. The analysis distinguishes between *typical* and *extreme* wealth outcomes. Accurate beliefs maximize objective expected log wealth growth. They do not, however, generate the thickest upper tail. Moderate optimism raises exposure to multiplicative idiosyncratic returns at only a second-order cost in typical growth, increasing the probability of exceptionally successful wealth histories. This distinction has a direct implication for disagreement: the symmetric, expanding mean-preserving spreads studied below can lower average long-run wealth growth while thickening the upper tail, up to the belief that generates the thickest tail.

4.1. Analytical benchmark

We simplify the environment in Section 3 in six dimensions. First, we freeze beliefs at their persistent component. This reduction is motivated by the evidence in Section 2: individual fixed effects account for most of the panel variation in expected stock returns. The analytical exercise therefore asks how persistent differences in subjective expected returns map into persistent differences in portfolio choice and wealth accumulation. Second, households share a common discount factor β and coefficient of relative risk aversion $\gamma > 1$. Thus, the only permanent source of household heterogeneity in the benchmark is beliefs. Third, stock returns are i.i.d. and safe rates are constant. Fourth, labor income is constant and non-pledgeable, $\bar{W} > 0$. Fifth, portfolio choices are unconstrained, so that the optimal risky share is given by the interior Merton solution. Sixth, there is no aggregate return risk: households face only idiosyncratic risky-return innovations.

The analytical benchmark retains the idiosyncratic component of realized risky returns present in the baseline model. This feature is motivated by the substantial within-asset-class return heterogeneity documented in household data (Fagereng et al. (2020); Hubmer, Krusell, and Smith (2021)). It is also important conceptually: idiosyncratic shocks generate independent multiplicative return histories within a cross-section, even in the absence of aggregate return fluctuations.

Subjective beliefs. Under household i's subjective probability measure, its risky return satisfies

$$R_t^{S,i} = \zeta^i \varepsilon_t^{P,i}, \qquad \log \varepsilon_t^{P,i} \overset{i.i.d.}{\sim} \mathcal{N}\left(-\frac{1}{2}\sigma^2, \sigma^2\right). \tag{52}$$

Because $\mathbb{E}^{\mathcal{P}^i}[\varepsilon_t^{P,i}] = 1$, ζ^i is household i's subjective expected gross stock return. Households agree about return risk and disagree only about the mean.

Let $R^B = e^{r_B}$ denote the gross safe return and define subjective expected excess returns by

$$\Delta^i \equiv \log \zeta^i - r_B. \tag{53}$$

The corresponding objective quantity is

$$\Delta^0 \equiv \log \zeta^0 - r_B, \qquad \Delta^0 > 0. \tag{54}$$

Under the objective measure, the risky return is

$$R_t^{S,i} = \zeta^0 \eta_t^i, \tag{55}$$

where

$$\log \eta_t^i \overset{i.i.d.}{\sim} \mathcal{N}\left(-\frac{1}{2}\sigma_\eta^2, \sigma_\eta^2\right), \tag{56}$$

and the innovations are independent across households and over time. We set

$$\sigma^2 \equiv \sigma_\eta^2. \tag{57}$$

Thus all risky-return uncertainty in the analytical benchmark is idiosyncratic, and belief disagreement concerns expected returns, not perceived risk.

High-wealth policy functions. Labor income is constant and non-pledgeable, $\bar{W} > 0$. Because labor income enters cash-on-hand additively, the household problem is not globally homothetic in financial wealth. We therefore use the Merton–Samuelson solution as a high-wealth approximation. In the upper tail,

$$\frac{\bar{W}}{R_t^{A,i} A_{t-1}^i + \bar{W}} \simeq 0, \tag{58}$$

so portfolio and consumption-saving policies approach their proportional Merton–Samuelson counterparts. This approximation is restrictive for the bottom of the wealth distribution, where labor income and borrowing constraints are important, but is well suited to the object of interest here: the asymptotic upper tail.

The risky portfolio share is

$$v^i \simeq \frac{\log\left(\mathbb{E}_t^{\mathscr{P}^i}[R_{t+1}^S]\right) - \log R^B}{\gamma\sigma^2} = \frac{\Delta^i}{\gamma\sigma^2}. \tag{59}$$

Households with higher expected excess returns therefore hold more risky assets. The Samuelson saving propensity is

$$\kappa^i \simeq \left[\beta \mathbb{E}^{\mathscr{P}^i}\left[(R_{t+1}^{A,i})^{1-\gamma}\right]\right]^{1/\gamma}. \tag{60}$$

Under the Merton–Samuelson approximation, the log return on household i's portfolio is

$$\log R_{t+1}^{A,i} \simeq r_B + v^i \Delta^i - \frac{1}{2}(v^i)^2\sigma^2 + v^i \sigma \varepsilon_{t+1}, \qquad \varepsilon_{t+1} \sim \mathcal{N}(0,1). \tag{61}$$

Hence, using the lognormal moment formula,

$$\log \mathbb{E}^{\mathscr{P}^i}\left[(R_{t+1}^{A,i})^{1-\gamma}\right] = (1-\gamma)\left[r_B + v^i\Delta^i - \frac{\gamma}{2}(v^i)^2\sigma^2\right] \tag{62}$$

$$= (1-\gamma)\left[r_B + \frac{(\Delta^i)^2}{2\gamma\sigma^2}\right], \tag{63}$$

where the second equality uses $v^i = \Delta^i/(\gamma\sigma^2)$. Therefore,

$$\log\kappa^i \simeq \frac{1}{\gamma}\left[\log\beta + (1-\gamma)\left(r_B + \frac{(\Delta^i)^2}{2\gamma\sigma^2}\right)\right]. \tag{64}$$

For $\gamma > 1$,

$$\frac{\partial \log\kappa^i}{\partial\Delta^i} = -\frac{\gamma-1}{\gamma^2\sigma^2}\Delta^i < 0. \tag{65}$$

A higher perceived expected return therefore raises risky exposure but lowers the fraction of resources carried into the next period. The latter result is the familiar income-effect response when the elasticity of intertemporal substitution, $1/\gamma$, is below one: the wealth effect of a higher expected return dominates the substitution effect.[9]

4.2. Subjective beliefs and typical wealth growth

Under the proportional policy rules, financial wealth evolves according to

$$A_t^i = \kappa^i R_t^{A,i} A_{t-1}^i + \kappa^i \bar{W}. \tag{66}$$

The positive wage term is quantitatively negligible relative to wealth in the upper tail, but is qualitatively important: it provides the additive component that prevents wealth from being absorbed at zero.

Under the objective measure, household i's portfolio return is $R_t^{A,i} = (1-v^i)R^B + v^i R_t^{S,i}$. Using a second-order lognormal approximation,

$$\begin{aligned}\log R_t^{A,i} \simeq\ & r_B + v^i\Delta^0 - \frac{1}{2}(v^i)^2\sigma^2 \\ & + v^i\sigma_\eta z_t^i,\end{aligned} \tag{67}$$

where z_t^i are independent standard-normal idiosyncratic innovations, with $\log\eta_t^i = -\sigma_\eta^2/2 + \sigma_\eta z_t^i$. At high wealth, the additive wage component is negligible for growth. Objective expected log wealth growth is therefore

$$\begin{aligned}\lambda^i &\equiv \mathbb{E}^0\left[\log\left(\kappa^i R_t^{A,i}\right)\right] \\ &= \log\kappa^i + r_B + v^i\Delta^0 - \frac{1}{2}(v^i)^2\sigma^2.\end{aligned} \tag{68}$$

[9]This comparative static holds perceived return variance fixed. A change in perceived risk would generate a separate precautionary-saving response.

Substituting (59) and (64) gives

$$\lambda^i = \frac{1}{\gamma}\left[\log\beta + r_B + \frac{(\Delta^0)^2}{2\sigma^2} - \frac{(\Delta^i - \Delta^0)^2}{2\sigma^2}\right]. \tag{69}$$

Belief errors therefore reduce objective expected log wealth growth quadratically:

$$\lambda^i - \lambda^0 = -\frac{(\Delta^i - \Delta^0)^2}{2\gamma\sigma^2}, \qquad \lambda^0 \equiv \lambda(\Delta^0). \tag{70}$$

The marginal effect of subjective expected returns makes the underlying forces transparent:

$$\begin{aligned}\frac{\partial\lambda^i}{\partial\Delta^i} &= \underbrace{\Delta^0\frac{\partial v^i}{\partial\Delta^i}}_{\substack{\text{return gain}\\(+)}} + \underbrace{\left(-\sigma^2 v^i\frac{\partial v^i}{\partial\Delta^i}\right)}_{\substack{\text{volatility drag}\\(-)}} + \underbrace{\frac{\partial\log\kappa^i}{\partial\Delta^i}}_{\substack{\text{lower saving}\\(-),\ \gamma>1}} \\ &= \frac{\Delta^0 - \Delta^i}{\gamma\sigma^2}.\end{aligned} \tag{71}$$

The first term is the expected-return benefit of increasing equity exposure. The second is the standard geometric-growth cost of portfolio volatility: larger risky positions increase the dispersion of realized returns, and bad realizations have disproportionately large effects on compounded wealth. The third term reflects the consumption-saving response in (65): for $\gamma > 1$, greater perceived investment opportunities induce households to consume more of current resources.

It follows immediately that

$$\frac{\partial\lambda(\Delta)}{\partial\Delta}\begin{cases} > 0, & \Delta < \Delta^0, \\ = 0, & \Delta = \Delta^0, \\ < 0, & \Delta > \Delta^0. \end{cases} \tag{72}$$

Hence objective expected log wealth growth is uniquely maximized by the accurately informed household. Optimism and pessimism of equal magnitude generate the same loss in typical long-run growth. This selection toward accurate beliefs is closely related to the mechanism in Martin and Papadimitriou (2022), where investors who are correct in hindsight become relatively wealthy.

Stationarity. Equation (66) is a stochastic affine recursion. Under the usual regularity conditions, a stationary wealth distribution exists if

$$\lambda^i = \mathbb{E}^0\left[\log\left(\kappa^i R_t^{A,i}\right)\right] < 0. \tag{73}$$

Equivalently,

$$\underbrace{-\log\kappa^i}_{\text{consumption drag}} + \underbrace{\frac{1}{2}(v^i)^2\sigma^2}_{\text{volatility drag}} > \underbrace{r_B + v^i\Delta^0}_{\text{return component}}. \tag{74}$$

This condition changes the long-run wealth dynamics relative to a purely multiplicative Merton environment. Without the additive income term, negative expected log growth eventually drives financial wealth toward zero, while positive expected log growth generates unbounded accumulation. Here the wage flow repeatedly replenishes financial wealth, while (73) prevents the multiplicative component from exploding. The result is a non-degenerate stationary random variable.

Moreover, because accurate beliefs maximize $\lambda(\Delta)$, stationarity of the accurately informed household,

$$\lambda(\Delta^0) < 0, \tag{75}$$

implies stationarity for every other belief:

$$\lambda(\Delta) \leq \lambda(\Delta^0) < 0. \tag{76}$$

Thus more optimistic beliefs do not mechanically imply faster long-run wealth growth or eventual dominance by the most optimistic households.

4.3. Subjective beliefs and the wealth tail

We now turn from typical wealth growth to extreme wealth outcomes. With aggregate risk shut down, households within any given belief class face independent idiosyncratic return histories. The stationary wealth distribution is therefore a fixed cross-sectional object rather than a distribution that moves with the aggregate history.

For a household with belief Δ, let $F(a \mid \Delta)$ denote the stationary cross-sectional wealth distribution. Under the lognormal high-wealth approximation, its upper tail satisfies

$$1 - F(a \mid \Delta) \sim C(\Delta) a^{-\xi(\Delta)}, \qquad a \to \infty, \tag{77}$$

where $C(\Delta) \in (0, \infty)$ is a constant. The Kesten–Goldie result applies to the approximating affine wealth recursion under the regularity conditions verified in Appendix D.

Define the multiplicative component of wealth accumulation as

$$M_t^i \equiv \kappa^i R_t^{A,i}. \tag{78}$$

From (67),

$$\log M_t^i = \lambda(\Delta^i) + v^i \sigma_\eta z_t^i, \qquad z_t^i \overset{i.i.d.}{\sim} \mathcal{N}(0,1). \tag{79}$$

For any fixed $\Delta > 0$, the positive Pareto exponent solves the standard Kesten moment condition

$$\mathbb{E}_\eta \left[M_t^\xi \right] = 1, \tag{80}$$

which, using the lognormal moment formula, becomes

$$\xi \lambda(\Delta) + \frac{1}{2} \xi^2 v(\Delta)^2 \sigma_\eta^2 = 0. \tag{81}$$

The positive solution is

$$\xi(\Delta) = -\frac{2\lambda(\Delta)}{v(\Delta)^2\sigma_\eta^2}. \tag{82}$$

Equation (82) gives the Pareto exponent of the stationary cross-sectional wealth distribution within a belief class. Beliefs affect this exponent through two channels: they determine both the drift of wealth accumulation, $\lambda(\Delta)$, and exposure to idiosyncratic multiplicative risk, $v(\Delta)^2\sigma_\eta^2$.

Substituting the Merton share (59) and the growth rate (69) into (82) gives the Pareto coefficient as a closed-form function of beliefs. Define

$$K \equiv -2\sigma^2\left(\log\beta + r_B\right). \tag{83}$$

Proposition 1 (Beliefs and the wealth tail)**.** *Suppose* $\Delta^0 > 0$*,* $\sigma_\eta > 0$*, and* $\lambda(\Delta^0) < 0$*. Then:*

1. *For every* $\Delta > 0$*, within the lognormal affine approximation, the Pareto coefficient of the stationary cross-sectional wealth distribution of households with belief* Δ *is*

$$\xi(\Delta) = \gamma\frac{\Delta^2 - 2\Delta^0\Delta + K}{\Delta^2}. \tag{84}$$

2. $\xi(\Delta)$ *is strictly decreasing on* $(0, \Delta^\star)$ *and strictly increasing on* $(\Delta^\star, \infty)$*, where*

$$\Delta^\star \equiv \frac{K}{\Delta^0}. \tag{85}$$

3. *Stationarity of the accurately informed household implies* $K > (\Delta^0)^2$ *and hence*

$$\Delta^\star > \Delta^0. \tag{86}$$

The thickest cross-sectional wealth tail is therefore generated by households that are optimistic relative to the objective expected return, with

$$\xi(\Delta^\star) = \gamma\left(1 - \frac{(\Delta^0)^2}{K}\right) \qquad \textit{and} \qquad \lim_{\Delta\to\infty}\xi(\Delta) = \gamma. \tag{87}$$

Proof. See Appendix D. ■

Finite mean wealth. Stationarity alone does not ensure that the stationary distribution has a finite mean. For the Pareto tail in (77), finite mean requires $\xi(\Delta) > 1$. To discuss aggregate wealth and top wealth shares, we impose the sufficient parameter restriction

$$K > \frac{\gamma}{\gamma - 1}(\Delta^0)^2, \tag{88}$$

which implies $\xi(\Delta^\star) > 1$ and hence $\xi(\Delta) > 1$ for every $\Delta > 0$. For the belief distributions considered below, this ensures finite mean cross-sectional wealth. The restriction is stronger than the stationarity condition (75).

Unlike typical wealth growth, the Pareto coefficient is thus not minimized at

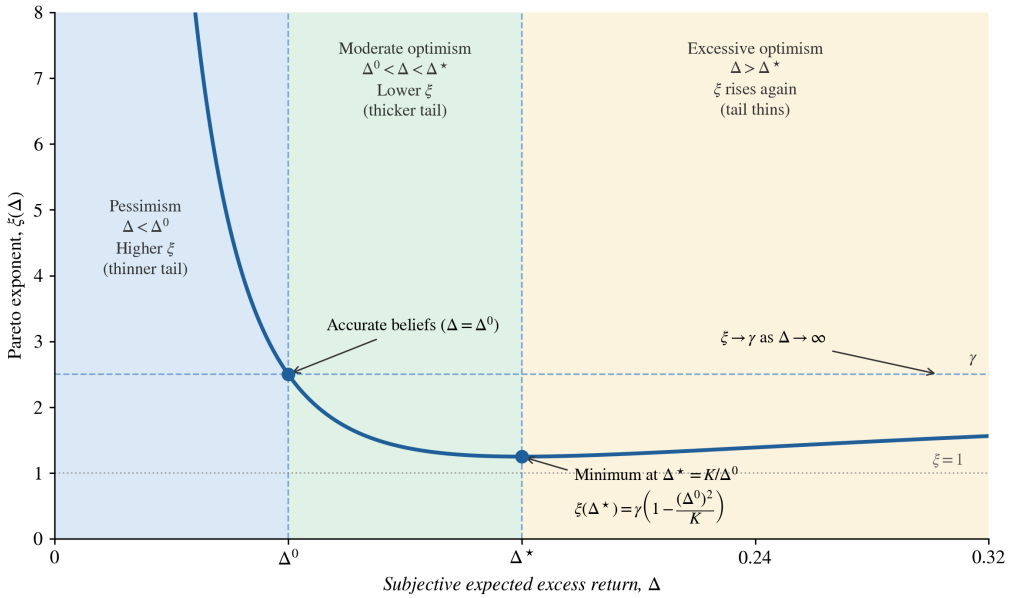

FIGURE 3. **Subjective expected excess return and the Pareto exponent.** The figure plots $\xi(\Delta)$ from Proposition 1 for an illustrative annual parameterization: $\gamma = 2.5$, $\sigma_\eta = 0.20$, $\Delta^0 = 0.08$, $r_B = \log(1.0075)$, and β chosen so that $K = 2(\Delta^0)^2$ ($\beta \simeq 0.846$). With these values $\xi(\Delta^0) = \gamma = 2.5$, $\Delta^\star = 2\Delta^0$, and $\xi(\Delta^\star) = 1.25 > 1$, ensuring finite mean wealth.

the true belief: the belief that maximizes typical growth, Δ^0, lies on the decreasing branch of $\xi(\Delta)$.

Economic mechanism. The difference between the beliefs that maximize typical performance and extreme wealth has a simple local interpretation. Consider a small optimistic deviation from the truth, $\Delta = \Delta^0 + h$, with $h > 0$. Risky exposure increases to first order:

$$v(\Delta^0 + h) - v(\Delta^0) = \frac{h}{\gamma\sigma^2}. \tag{89}$$

By contrast, the loss in expected log wealth growth is only second order:

$$\lambda(\Delta^0 + h) - \lambda(\Delta^0) = -\frac{h^2}{2\gamma\sigma^2}. \tag{90}$$

A small amount of optimism therefore buys a first-order increase in exposure to multiplicative idiosyncratic returns at only a second-order cost in typical wealth growth. Among a large population of such households, some experience unusually favorable sequences of idiosyncratic returns and become extremely wealthy.

This force does not continue indefinitely. As beliefs become increasingly optimistic, the deterioration in objective geometric wealth growth eventually dominates the additional exposure to idiosyncratic risk. Beyond $\Delta^\star$, the Pareto coefficient therefore rises again. The thickest cross-sectional wealth tail is generated neither by the accurately informed household nor by the most optimistic household, but by an interior degree of optimism. Figure 3 illustrates this non-monotonic relationship.

We next aggregate across belief classes. Let subjective expected excess returns be distributed in the population according to G_Δ. The cross-sectional wealth distribution is the mixture

$$F^{XS}(a) = \int F(a \mid \Delta)\, dG_\Delta(\Delta). \tag{91}$$

Each belief class contributes a Pareto tail with coefficient $\xi(\Delta)$. In the far upper tail, classes with larger Pareto coefficients decay faster and become asymptotically negligible relative to the class, or set of classes, with the smallest coefficient. Therefore, under the usual regularity conditions for mixtures of Pareto tails (e.g., finite belief support or appropriate uniform tail bounds),

$$\xi^{XS} = \inf_{\Delta \in \mathrm{supp}(G_\Delta)} \xi(\Delta). \tag{92}$$

The extreme upper tail of the social wealth distribution is therefore governed not by the average household or the average belief, but by the belief types that generate the heaviest conditional tails. In particular, if

$$\Delta^\star \in \mathrm{supp}(G_\Delta), \tag{93}$$

then

$$\xi^{XS} = \xi(\Delta^\star). \tag{94}$$

4.4. Belief disagreement and wealth inequality

The previous results characterize wealth accumulation conditional on a household's belief. We now use them to ask a population-level question: what happens to wealth inequality when households disagree more, holding the average belief fixed?

To isolate disagreement from average optimism or pessimism, consider a family of belief distributions G_d, indexed by $d \geq 0$, obtained from symmetric mean-preserving spreads around the objectively correct belief:

$$\Delta = \Delta^0 + dz, \qquad z \sim H, \qquad \mathrm{supp}(H) = [-1, 1], \tag{95}$$

where H is symmetric around zero and $\mathbb{E}[z] = 0$. Hence

$$\mathbb{E}_{G_d}[\Delta] = \Delta^0, \qquad \mathrm{supp}(G_d) = [\Delta^0 - d, \Delta^0 + d]. \tag{96}$$

We restrict attention to $0 \leq d < \Delta^0$ so that subjective expected excess returns remain positive for every household.

Proposition 2 (Disagreement and wealth inequality)**.** *Suppose* $\Delta^0 > 0$*,* $\sigma_\eta > 0$*, and the finite-mean restriction* (88) *holds. Consider the family of symmetric belief distributions in* (95)*. Then:*

1. Average objective expected log wealth growth satisfies

$$\mathbb{E}_{G_d}[\lambda(\Delta)] = \lambda(\Delta^0) - \frac{\mathrm{Var}_{G_d}(\Delta)}{2\gamma\sigma^2} = \lambda(\Delta^0) - \frac{d^2\,\mathrm{Var}(z)}{2\gamma\sigma^2}. \tag{97}$$

Thus greater disagreement lowers average long-run wealth growth even though beliefs remain correct on average.

2. *For any* $h \in (0, \Delta^0)$,

$$\xi(\Delta^0 - h) - \xi(\Delta^0 + h) = 4\gamma \frac{\Delta^0 h \left[K - (\Delta^0)^2 + h^2\right]}{\left[(\Delta^0)^2 - h^2\right]^2} > 0. \tag{98}$$

Hence, among two households whose beliefs are equally inaccurate in opposite directions, the optimistic household has the thicker stationary wealth tail.

3. *Define*

$$d^\star \equiv \Delta^\star - \Delta^0 = \frac{K}{\Delta^0} - \Delta^0 > 0. \tag{99}$$

The social Pareto coefficient is

$$\xi^{XS}(d) = \xi\left(\min\{\Delta^0 + d, \Delta^\star\}\right). \tag{100}$$

Therefore $\xi^{XS}(d)$ *is strictly decreasing in disagreement as long as* $\Delta^0 + d < \Delta^\star$. *If* $d^\star < \Delta^0$, *then*

$$\xi^{XS}(d) = \begin{cases} \xi(\Delta^0 + d), & 0 \le d < d^\star, \\ \xi(\Delta^\star), & d^\star \le d < \Delta^0. \end{cases} \tag{101}$$

Proof. See Appendix D. ■

The proposition isolates a distributional effect of disagreement that is distinct from average belief bias. Increasing d introduces pessimists and optimists symmetrically, leaving the population-average belief equal to the truth. Because the loss in typical wealth growth is quadratic in the belief error, both sides of the distribution perform worse on average and aggregate expected log wealth growth falls.

The consequences for extreme wealth are asymmetric. Optimists increase their exposure to multiplicative idiosyncratic returns, whereas pessimists reduce it. Thus, among households that are equally wrong in opposite directions, the optimistic type generates the thicker wealth tail. The social upper tail is consequently selected by the optimistic side of an otherwise unbiased belief distribution.

Within this particular family of symmetric, expanding belief supports, greater disagreement can therefore simultaneously reduce average long-run wealth growth and thicken the upper tail. This conclusion does not extend to arbitrary mean-preserving spreads, which can remove tail-dominating beliefs from the support. Once the support of beliefs contains $\Delta^\star$, further disagreement does not change the asymptotic Pareto coefficient. The amount of probability mass around the tail-dominating beliefs can nevertheless continue to affect how quickly the asymptotic tail becomes quantitatively important and, hence, finite-tail objects such as top wealth shares. Tail-exponent comparisons do not by themselves establish a monotone response of finite top wealth shares.

Relative optimism. Propositions 1 and 2 are stated relative to the objective expected return, but their content is a comparative static in beliefs. For any two

households on the decreasing branch of $\xi(\Delta)$, the one with the higher expected return holds more risky assets and generates the thicker tail, whether or not either belief exceeds Δ^0. In the quantitative model of Section 6, every wealth group's expected stock return, between 5.1 and 8.6 percent, lies below the model's mean stock return of 12.8 percent. Thus the groups are relatively pessimistic and less pessimistic households take more risk, consistently with the analytical portfolio-choice mechanism. The result about the decreasing branch of $\xi(\Delta)$, however, is established only for the simplified benchmark with fixed beliefs, common preferences, and no aggregate risk; it is not a proved property of the quantitative economy. Throughout the paper, "optimism" refers to this relative sense unless stated otherwise.

5. Solving the Model: Two-Stage Deep Learning

Section 3 defined the competitive equilibrium in sequence form. We now write the model recursively and use that representation to solve the quantitative economy. The key computational implication of Internal Rationality is that the state needed for household optimization differs from the state needed to determine equilibrium prices. In particular, the intermediated household observes its own risky return and the quoted safe return, but not the stock price. This separation allows us to solve the household problem first and impose market clearing in a second stage.

We implement this idea with a numerical procedure that we call *Two-Stage Deep Learning* (2SDL). Section 5.1 presents the recursive formulation and defines a recursive competitive equilibrium under Internal Rationality. Sections 5.2 and 5.3 then describe the two stages of the algorithm.

5.1. Recursive formulation and equilibrium

We begin by distinguishing the objective state of the economy from the information available to an individual household. Let Γ_t denote the cross-sectional distribution of predetermined household states, including inherited financial wealth and portfolio shares, labor-income and preference types, and prior belief states ζ_t^i. Aggregate dividends D_t and the previous period's stock and bond prices (P_{t-1}, Q_{t-1}) are also predetermined from the perspective of the period-t price-clearing problem. Given household policy functions, market clearing determines two current prices, so the objective pricing function can be written

$$(P_t, Q_t) = p(\Gamma_t, D_t, P_{t-1}, Q_{t-1}). \tag{102}$$

We retain the lagged prices explicitly because the current stock return is $(P_t+D_t)/P_{t-1}$ and the realized return on inherited safe positions is $1/Q_{t-1}$. Stock prices are not directly observed by the household; safe-asset prices are observed through their quoted returns.

The individual information state. A household's decision state contains the variables proposed in Section 3: wage income, inherited savings and equity share, dividends, its prior belief, and the personal risky return reported by the intermediary.

It also contains Q_{t-1} and Q_t, both recoverable from the safe-asset quotes. The former determines the realized return $R_t^B = 1/Q_{t-1}$ on inherited safe savings; the latter determines the known return $R_{t+1}^B = 1/Q_t$ offered on new savings. Thus

$$(\boldsymbol{x}_t^i)^{IR} = \left(W_t^i, A_{t-1}^i, v_{t-1}^i, D_t, Q_{t-1}, Q_t, \zeta_t^i, R_t^{S,i}\right). \tag{103}$$

Preference type is suppressed in this notation. Unlike the stock price, Q_{t-1} and Q_t belong to the household's observed information, and both are needed when safe returns vary over time.

The household's beginning-of-period resources satisfy

$$X_t^i \equiv c_t^i + A_t^i = W_t^i + \left[v_{t-1}^i R_t^{S,i} + (1 - v_{t-1}^i) R_t^B\right] A_{t-1}^i. \tag{104}$$

At date t the household first observes $R_t^{S,i}$, forms the experience statistic in (35), updates its forecast from ζ_t^i to ζ_{t+1}^i according to (40), and then chooses (c_t^i, A_t^i, v_t^i). The inclusion of both the incoming prior and the current personal return in (103) makes this timing explicit. Equivalently, one may use cash-on-hand, the updated return forecast, dividends and the current safe-asset price as decision states; we retain the expanded state because it makes the dependence on inherited portfolio positions transparent.

For the purpose of comparison, under full-information rational expectations households know the equilibrium pricing map p and the law of motion of the cross-sectional distribution. The relevant information then includes

$$(\boldsymbol{x}_t^i)^{RE} = \left(W_t^i, A_{t-1}^i, v_{t-1}^i, D_t, P_{t-1}, Q_{t-1}, \Gamma_t\right), \tag{105}$$

with current asset prices recovered from the pricing function. The household problem therefore requires forecasting the high-dimensional object Γ_t. By contrast, an internally rational household receives its personal return and safe-rate quote as observables, and simulates their *subjective* joint distribution directly, without forecasting the aggregate distribution or its associated price map.

Dollar policies and demands for asset units. It is important to distinguish the objects produced by household optimization from the demands entering market clearing. Define

$$\left(c_t^i, A_t^i, v_t^i\right) = \left(h_c^i, h_A^i, h_v^i\right)\left((\boldsymbol{x}_t^i)^{IR}\right). \tag{106}$$

These are policies for consumption, the *dollar amount* saved, and the fraction of savings allocated to risky assets; they are not policies for the numbers of stock shares or safe trees. At market-clearing prices, intermediaries convert dollar allocations into asset quantities:

$$S_t^i = \frac{h_v^i(\boldsymbol{x}_t^i) h_A^i(\boldsymbol{x}_t^i)}{P_t}, \qquad N_t^{B,i} = \frac{[1 - h_v^i(\boldsymbol{x}_t^i)] h_A^i(\boldsymbol{x}_t^i)}{Q_t}. \tag{107}$$

Stock and safe-tree supplies are 1 and $\bar{B}_t$, respectively. Thus individual dollar policies are well defined without knowledge of the stock price, while their implied *unit* demands require prices. This is precisely the distinction exploited by market clearing.

Recursive competitive equilibrium. Given subjective measures $\{\mathscr{P}^i\}_i$, objective exogenous shock laws, and the initial distribution and prices, a recursive competitive equilibrium under Internal Rationality consists of the individual policies (106), a price mapping

$$(P_t, Q_t) = p(\Gamma_t, D_t, P_{t-1}, Q_{t-1}), \tag{108}$$

and an induced law for Γ_t, such that:

1. *Household optimality.* For each observed household state, the policies solve the subjective optimization problem of Section 3 using the Bayesian posterior and the quoted return on the new safe asset.
2. *Market clearing.* The prices (108) induce observed personal returns and safe quotes and ensure that the unit demands (107) sum to the asset supplies.
3. *Distributional consistency.* The realized prices, aggregate dividends, idiosyncratic returns and policies jointly determine the subsequent distribution of financial positions, labor-income states and prior beliefs according to (40).

The consistency requirement is standard on the *objective* side of the economy: market-clearing prices determine realized returns, households respond to their own returns, and their choices determine subsequent aggregate states. Internal Rationality removes the additional requirement that individual subjective forecasts reproduce the objective price mapping p.

Two-stage decomposition. Conditional on the specified subjective measures, household policies can first be computed without imposing market clearing:

$$\mathscr{P}^i \quad \Longrightarrow \quad h^i\big((\boldsymbol{x}_t^i)^{IR}\big). \tag{109}$$

Given these policies, the objective distribution and previous prices, Stage 2 constructs return and quote observations at each candidate price, converts household dollar demands into asset units, and imposes market clearing:

$$\big\{h^i\big((\boldsymbol{x}_t^i(P,Q))^{IR}\big)\big\}_i,\ \Gamma_t, D_t, P_{t-1}, Q_{t-1} \quad \Longrightarrow \quad (P_t, Q_t). \tag{110}$$

Although the *realized* individual state depends on equilibrium prices, the *subjective policy function* does not use the mapping from prices to the cross-sectional distribution. This is the basis of the 2SDL algorithm.

5.2. Stage 1: household policies under subjective beliefs

In the first stage, we solve the household problem conditional on the subjective joint process for personal investment returns, dividends, labor-income types, and future safe-rate quotes. It is convenient to parameterize the policy in terms of the

saving rate κ_t^i and the risky portfolio share v_t^i. We approximate both policies jointly with a neural network,

$$(\kappa_t^i, v_t^i) = \mathcal{N}_\theta((\boldsymbol{x}_t^i)^{IR}), \tag{111}$$

where θ collects network parameters. Given resources X_t^i, the saving policy determines $A_t^i = \kappa_t^i X_t^i$ and $c_t^i = (1-\kappa_t^i)X_t^i$, subject to the minimum participation balance. The portfolio policy divides A_t^i into risky and safe *dollar* investment. The intermediary executes these allocations after the household has chosen them.

Economic theory provides a useful starting point for the approximation. We initialize the network close to the analytical Merton–Samuelson solution. In the benchmark with constant investment opportunities, the optimal risky share balances subjective excess returns against risk, and the saving rate has an analytical consumption–saving counterpart. Training learns the nonlinear corrections generated by income risk, changing posterior expected returns, variation in quoted safe rates, and occasionally binding constraints.

Training states are generated by simulating household histories under their subjective measures $\mathscr{P}^i$. In particular, Stage 1 draws the *joint* Gaussian market-return and personal-return innovations, latent return states, dividend and safe-quote shocks, and labor-income transitions specified in Section 3. It propagates personal returns and the Bayesian filter jointly with individual financial resources. No objective stock-price path, equilibrium pricing function or cross-sectional distribution is simulated or forecast at this stage. This is where Internal Rationality delivers the computational separation.

The fact that personal returns are signals does not eliminate their role as payoffs. Households take into account that a high realized $R_{t+1}^{S,i}$ simultaneously raises next-period cash-on-hand and next-period optimism. Thus Euler equations must use the correct *joint* distribution of payoffs and posterior beliefs, rather than drawing these two innovations independently. Conditional on positive participation, the observation $\mathscr{S}_{t+1}^i = \log R_{t+1}^{S,i} + \sigma_\eta^2/2$ is independent of the *magnitude* of v_t^i; accordingly the interior portfolio Euler condition contains the standard marginal excess payoff, without an additional intensive-margin value-of-information term:

$$\mathbb{E}_t^{\mathscr{P}^i}\left[u_i'(c_{t+1}^i)\left(R_{t+1}^{S,i} - 1/Q_t\right)\right] = 0, \qquad \underline{v} < v_t^i < \bar{v}. \tag{112}$$

Here next-period consumption already depends on the personal-return signal through the optimally updated belief. Equation (112) is an interior condition and must be replaced by KKT inequalities when a portfolio boundary binds; the small strictly positive equity participation bound eliminates the separate discrete experiment of whether to observe any stock return.

Simulation alone concentrates training observations in the region visited under current candidate policies. We therefore retain the original combination of simulated histories and off-path collocation states, with additional mass near the borrowing/saving and portfolio constraints.

The network is trained on Euler–KKT conditions. Let $\mathscr{E}_A(\boldsymbol{x};\theta)$ denote the residual for the saving choice and $\mathscr{E}_v(\boldsymbol{x};\theta)$ the portfolio residual. Because assets and risky

shares are subject to

$$A_t^i \geq \underline{a} D_t > 0, \qquad \underline{v} \leq v_t^i \leq \bar{v}, \tag{113}$$

first-order conditions include complementary slackness at the boundaries. We impose these conditions using the Fischer–Burmeister function

$$\phi(a,b) = \sqrt{a^2 + b^2} - a - b, \tag{114}$$

which equals zero if and only if $a \geq 0$, $b \geq 0$, and $ab = 0$. Applying the transformation to the Euler gaps and constraint slacks gives differentiable residuals. Schematically, the training objective is

$$\mathscr{L}(\theta) = \mathbb{E}_{\boldsymbol{x}\sim\mathscr{D}}\left[\Phi_A(\boldsymbol{x};\theta)^2 + \Phi_{v,\underline{v}}(\boldsymbol{x};\theta)^2 + \Phi_{v,\bar{v}}(\boldsymbol{x};\theta)^2 + \lambda\Pi(\boldsymbol{x};\theta)\right], \tag{115}$$

where $\mathscr{D}$ combines subjective simulated states and collocation states, and Π penalizes infeasible present or future resources. The residuals are evaluated using the updated conditional predictive *distributions*, not only the posterior mean. If the filter is initialized at its stationary Riccati solution, posterior variance is common and deterministic, so it does not add a heterogeneous state to the neural network. A finite numerical collocation domain is only an approximation domain: mechanically clipping posterior beliefs at its boundaries would invalidate the exact Gaussian Bayesian recursion and must not be treated as part of the theoretical updating rule.

At the end of Stage 1, we obtain household policy functions that are optimal under their subjective measures. Neither an objective equilibrium price map nor a perceived law of motion for the cross-sectional distribution is used in their computation.

5.3. Stage 2: general equilibrium

The second stage embeds the trained household policies in the heterogeneous-agent economy and determines equilibrium stock and safe asset prices. We simulate a large population of households, assigning inherited account balances, risky shares, labor-income types and prior beliefs, and drawing the objective idiosyncratic investment shocks. At the beginning of date t, the distribution Γ_t and the previous prices (P_{t-1}, Q_{t-1}) are predetermined.

From candidate prices to observed states. For candidate prices (P,Q), the objective aggregate and individual stock returns realized at date t are

$$R_t^S(P) = \frac{P + D_t}{P_{t-1}}, \qquad R_t^{S,i}(P) = R_t^S(P)\eta_t^i. \tag{116}$$

The safe return on inherited positions is $R_t^B = 1/Q_{t-1}$, whereas a new safe investment earns the known return $1/Q$. Thus the state passed to the *already trained* household policy is

$$\boldsymbol{x}_t^i(P,Q) = \left(W_t^i, A_{t-1}^i, v_{t-1}^i, D_t, Q_{t-1}, Q, \zeta_t^i, \frac{P + D_t}{P_{t-1}}\eta_t^i\right). \tag{117}$$

The price solver uses the *objective* return identity in (116). Households do not use that identity when solving their subjective problem: for them the personal risky return is an observed state, with its future conditional law determined by Section 3.2. In particular, there is no direct stock-price argument in h_A or h_v: P enters through the observed personal return, while the observable safe-asset price Q enters the household state directly. Neither dependence requires knowledge of the equilibrium pricing map.

Clearing conditions. The dollar amounts allocated to risky and safe assets are $h_v(\boldsymbol{x}_t^i)h_A(\boldsymbol{x}_t^i)$ and $[1-h_v(\boldsymbol{x}_t^i)]h_A(\boldsymbol{x}_t^i)$. The intermediary's stock and safe-tree demands in units are therefore

$$S_t^i(P,Q)=\frac{h_v(\boldsymbol{x}_t^i(P,Q))h_A(\boldsymbol{x}_t^i(P,Q))}{P},\qquad N_t^{B,i}(P,Q)=\frac{[1-h_v(\boldsymbol{x}_t^i(P,Q))]h_A(\boldsymbol{x}_t^i(P,Q))}{Q}. \tag{118}$$

With population weights μ^i, define excess demand equivalently in expenditure units:

$$\mathscr{M}(P,Q;\Gamma_t,D_t,P_{t-1},Q_{t-1})=\begin{pmatrix}\sum_i \mu^i h_v(\boldsymbol{x}_t^i(P,Q))h_A(\boldsymbol{x}_t^i(P,Q))-P\\ \sum_i \mu^i[1-h_v(\boldsymbol{x}_t^i(P,Q))]h_A(\boldsymbol{x}_t^i(P,Q))-Q\bar{B}_t\end{pmatrix}. \tag{119}$$

The first component is equivalent to $\sum_i \mu^i S_t^i=1$, and the second to $\sum_i \mu^i N_t^{B,i}=\bar{B}_t$. Equilibrium prices solve the two-dimensional system

$$\mathscr{M}(P_t,Q_t;\Gamma_t,D_t,P_{t-1},Q_{t-1})=0. \tag{120}$$

In particular, stock-market clearing can be written

$$P_t=\sum_i \mu^i h_v(\boldsymbol{x}_t^i(P_t,Q_t))h_A(\boldsymbol{x}_t^i(P_t,Q_t)). \tag{121}$$

The stock price is *determined* by equity demand despite not being an *argument* of household policies. The identity $v_t^i=P_tS_t^i/A_t^i$ converts dollar allocations into stock units, and market clearing pins down the common conversion price.

The indirect price dependence is essential. Holding an individual's observed state fixed, its policy has no direct price derivative, $\partial h_v/\partial P|_{\boldsymbol{x}}=0$. But across candidate prices in Stage 2, the agent's realized observation changes:

$$\frac{\partial R_t^{S,i}(P)}{\partial P}=\frac{\eta_t^i}{P_{t-1}},\qquad \frac{\partial(1/Q)}{\partial Q}=-\frac{1}{Q^2}. \tag{122}$$

Therefore $h_A(\boldsymbol{x}_t^i(P,Q))$ and $h_v(\boldsymbol{x}_t^i(P,Q))$ generally vary with P and Q in the market-clearing system, even though the household never knows the mapping from prices to those observations. The household's Bayesian belief update is likewise recomputed at each candidate P from the personal-return signal. Thus the two price-clearing equations remain simultaneous nonlinear equations; no price can generally be read off from fixed household orders without solving the contemporaneous system.

Once (P_t, Q_t) have been obtained, their corresponding personal returns and quotes determine household consumption, saving and asset allocations, and the Bayesian recursion determines the next prior beliefs. Together with the exogenous shocks these choices generate Γ_{t+1}. Repeating these steps constructs the equilibrium path,

$$(\Gamma_t, D_t, P_{t-1}, Q_{t-1}) \longrightarrow (P_t, Q_t) \longrightarrow \{R_t^{S,i}, c_t^i, A_t^i, v_t^i\}_i \longrightarrow \Gamma_{t+1}. \tag{123}$$

The price solver recovers the equilibrium mapping (102) only along the states visited by the economy. Existence and uniqueness of the fixed point are not implied by the accounting identities alone and must be checked numerically. The price function is never fed back into the household forecasting problem: Stage 1 enforces individual optimality under subjective observations, while Stage 2 enforces market clearing of the resulting *physical* stock and safe-tree orders. This is the two-stage computational advantage of Internal Rationality.

6. Quantitative Analysis

Status of quantitative results. The quantitative results reported in this section are preliminary and provisional. We are refining the computational algorithm and improving the calibration of the model. The results presented here describe the current quantitative implementation and will be updated as these refinements progress.

We now turn to the quantitative implications of the model. The exercise has three objectives. First, we ask whether a model disciplined jointly by household beliefs and portfolio data can reproduce the main aggregate and cross-sectional features of the U.S. wealth distribution. Second, we evaluate whether the learning mechanism generates the form of persistent disagreement documented in Section 2, rather than merely reproducing the average relationship between wealth and expectations. Third, we quantify how learning from own investment experiences contributes to wealth concentration, both in the stationary distribution and along the historical U.S. asset-price path.

A central difficulty is that no single data set contains all the objects needed for this exercise. RAND measures household beliefs repeatedly over time but contains only limited information on financial portfolios. The Survey of Consumer Finances (SCF), by contrast, provides a detailed description of household balance sheets and portfolio composition but does not elicit expected stock returns. Section 6.1 constructs a statistical bridge between the two surveys. Section 6.2 describes the calibration. Section 6.3 evaluates the stationary quantitative performance of the model. Section 6.4 studies historical transitional dynamics from 1960 to 2024.

6.1. A Joint RAND–SCF Dataset

Why combine the two surveys? The empirical analysis in Section 2 uses RAND because its panel dimension is essential for measuring persistent disagreement. RAND is less well suited for quantitative portfolio analysis. First, the survey elicits a coarse subjective probability distribution rather than a point forecast, as RAND

does not provide a complete household balance sheet. In particular, it does not observe directly held equities and bonds and provides limited information on the asset composition of retirement accounts. These omissions are consequential for our purposes because the mechanism operates precisely through the interaction of beliefs, risky portfolio exposure, and wealth accumulation.

The 2016 SCF has the opposite strengths. It contains detailed information on financial assets, retirement accounts, business wealth, debt, income, and household risk-taking, and it oversamples wealthy households whose portfolios account for a large share of aggregate equity holdings. It does not, however, contain stock-return expectations. We therefore use the overlapping variables in RAND and the SCF to impute expectations into the SCF.

This procedure is best viewed as a *statistical bridge*. The resulting RAND–SCF file is not a second data set in which beliefs are directly observed. Rather, it allows us to map the conditional expectation patterns measured in RAND into the much richer balance-sheet information in the SCF and thereby construct quantitative targets involving beliefs, portfolios, income, and wealth jointly.

Imputation. Let Z_{it} collect household characteristics observed in both surveys. Our preferred specification includes homeownership, an indicator for holding a retirement account, retirement-account wealth, the equity share of the retirement account, home value, wage income, and age. In RAND we estimate

$$Y_{it} = \alpha + Z_{it}'\beta + \varepsilon_{it}, \qquad Y_{it} \in \left\{ \mathbb{E}_{it}^{S}[R_{t+1}], R_{t+1} - \mathbb{E}_{it}^{S}[R_{t+1}] \right\}. \tag{124}$$

We then apply the estimated expected-return coefficients to the corresponding observables in the 2016 SCF.

Table 4 reports the first-stage estimates. Variables describing financial-market participation are particularly informative. Retirement account wealth and, especially, the equity share held inside retirement accounts are positively associated with expected stock returns and negatively associated with forecast errors. By contrast, housing variables contain little information about stock-return beliefs. This distinction is economically natural: beliefs about equity returns are more tightly related to the part of household wealth that is actually exposed to financial markets than to housing wealth.

The identifying assumption behind this exercise is a transportability assumption: conditional on the common observables Z, the systematic component of beliefs is the same in RAND and the SCF. We do not require residual belief heterogeneity to vanish. Indeed, Section 2 shows that most individual variation in beliefs is *not* explained by observables. The imputation is intended to recover the conditional mean of beliefs relevant for calibration, not to assign the true latent belief of every SCF household.

The comparison of wealth distributions in the asset categories observed by both surveys provides a useful check on this exercise. RAND underrepresents the bottom of the wealth distribution, but the two surveys are substantially closer among middle- and high-wealth households and, importantly, RAND contains substantial support

TABLE 4. **Imputing Return Expectations from RAND to the SCF**

	Expected return	Forecast error
Constant	−8.17	6.30***
Homeowner	−2.10	0.70
Retirement account	1.66	−0.27
Retirement account value	0.40***	−0.10***
Equity share in retirement acct.	0.04***	−0.01***
Home value	0.07	−0.02
Wage income	0.50	−0.21
Age	0.002	−0.01

Notes. The table reports estimates of equation (124) in the RAND sample. The expected-return specification is used to impute stock-return beliefs to households in the 2016 SCF. The forecast-error regression is reported as a diagnostic and is not used to construct the baseline imputed expected returns. *** $p < 0.01$.

in the right tail. This is the part of the distribution that is most relevant for equity demand and top wealth concentration.

The integrated data preserve three features that are central for the quantitative exercise. Expected stock returns rise with wealth; risky portfolio shares rise with expected returns; and forecast errors fall with wealth. The first of these relationships is already present directly in RAND and is therefore not an artifact of the imputation. The RAND–SCF bridge allows us to study the same belief gradient together with the substantially richer portfolio information in the SCF.

The SCF also contains a measure of households' willingness to take financial risk. Risk tolerance, wealth, and imputed expected returns are positively associated, but they are not interchangeable objects. This motivates an important feature of the quantitative model below. We allow for persistent heterogeneity in preferences, including risk aversion, rather than asking heterogeneous beliefs to account for all observed portfolio heterogeneity. Beliefs are then generated separately through learning from investment experience. The quantitative effect attributed to expectations is therefore measured *conditional on* preference heterogeneity.

6.2. Calibration

We calibrate the quarterly model to the United States and evaluate its stationary distribution against the 2016 SCF. The asset side of the model contains a risky financial asset and a safe asset. Accordingly, the empirical counterpart of model wealth is financial wealth: we exclude housing and other nonfinancial assets from the moments used to discipline portfolio allocations. This restriction is important. The model is designed to study how beliefs affect financial risk-taking and the accumulation of financial wealth, not to reproduce household demand for housing, or other physical assets.

The calibration separates three sources of heterogeneity. Aggregate dividend risk determines the return environment. A persistent six-state Markov process generates heterogeneity in labor income, patience, and risk aversion. Finally, beliefs

evolve endogenously through the learning mechanism of Section 3.3. In particular, the household type does not determine beliefs, and households do not begin life as exogenously assigned optimists or pessimists. Table 7 collects the parameter values.

Dividends and asset returns. Dividend growth is lognormal,

$$g_{t+1} \equiv \frac{D_{t+1}}{D_t} = \bar{g}_D \exp\left(-\frac{1}{2}\sigma_D^2 + \sigma_D \varepsilon_{t+1}^D\right), \qquad \varepsilon_{t+1}^D \sim \mathcal{N}(0,1). \tag{125}$$

We choose $\bar{g}_D$ and σ_D to reproduce an annual mean dividend-growth rate of 2.84 percent and an annual volatility of 6.40 percent. Given equilibrium prices, realized stock and safe returns are

$$R_{t+1}^S = \frac{P_{t+1} + D_{t+1}}{P_t}, \qquad R_{t+1}^B = \frac{1}{Q_t}. \tag{126}$$

The corresponding empirical return anchors are an average real stock return of 9.29 percent and an average real safe rate of 0.75 percent per year.

Using dividend growth rather than directly imposing a stochastic stock-return process is economically important. Stock returns are equilibrium objects in the stationary model: prices adjust to household saving and portfolio demand, and belief heterogeneity can therefore affect both quantities and equilibrium asset prices. The dividend process disciplines the fundamental cash-flow risk without mechanically imposing the return distribution that the model is intended to generate. Because dividends grow, we solve the model in dividend-normalized units: labor income and the supply of the safe asset are proportional to D_t, and household wealth is expressed relative to the current dividend.

Labor income, patience, and risk aversion. Each household is in one of six types, $\tau_t^i \in \{1,\dots,6\}$, which jointly determines its wage level, discount factor, and coefficient of relative risk aversion,

$$\tau_t^i \mapsto \left(\varepsilon(\tau_t^i), \beta(\tau_t^i), \gamma(\tau_t^i)\right). \tag{127}$$

Table 5 reports the support and the stationary population share of each type. The support bundles the three characteristics: the two highest-wage types are also the most patient and the most risk tolerant. This bundling is deliberate. In the RAND–SCF data, wage rank and risk-tolerance rank are strongly positively correlated (a weighted rank correlation of 0.76), and we choose the rank orderings and population shares to reproduce this correlation and the SCF wage shares by wealth group. Heterogeneity in patience helps reproduce dispersion in saving propensities, while heterogeneity in risk aversion generates portfolio heterogeneity even in the absence of disagreement. Wealth sorting by type is therefore already present without any heterogeneity in beliefs, so the model does not attribute every difference in risky holdings to beliefs. In the common-belief counterfactual below, the entire type process is left unchanged.

The type follows a Markov chain with transition matrix Π. We discipline its persistence with the annual two-state transition matrices estimated by Fernández-Villaverde and Levintal (2024) for labor income (bottom 90 versus top 10 percent) and for risk aversion (low versus high),

$$P_W^A = \begin{pmatrix} 0.9972 & 0.0028 \\ 0.0252 & 0.9748 \end{pmatrix}, \qquad P_\gamma^A = \begin{pmatrix} 0.993 & 0.007 \\ 0.001 & 0.999 \end{pmatrix}, \tag{128}$$

whose stationary distributions place 10 percent of households in the high-income class and 12.5 percent in the low-risk-aversion class. Because the model is quarterly, we take the exact stochastic fourth root of each matrix,

$$P_W^Q \simeq \begin{pmatrix} 0.999293 & 0.000707 \\ 0.006367 & 0.993633 \end{pmatrix}, \qquad P_\gamma^Q \simeq \begin{pmatrix} 0.998245 & 0.001755 \\ 0.000251 & 0.999749 \end{pmatrix}. \tag{129}$$

These matrices define each type's quarterly exit hazard,

$$m(\tau) = \phi \left[1 - \pi^W(\tau)\pi^\gamma(\tau) \right], \qquad \phi = 0.55, \tag{130}$$

where $\pi^W(\tau)$ and $\pi^\gamma(\tau)$ are the quarterly probabilities of remaining in the income and risk-aversion classes to which type τ belongs (types 2 and 6 belong to the high-income, low-risk-aversion classes and the other types to the low-income, high-risk-aversion classes), and ϕ scales the joint hazard so that turnover across the bundled types is not the product of two independent exit events. Conditional on leaving, a household draws a new type with probability proportional to the product of that type's population share and exit hazard,

$$\begin{aligned} &\Pi(\tau, \tau') = m(\tau)\, q(\tau') \quad \text{for } \tau' \neq \tau, \qquad \Pi(\tau, \tau) = 1 - m(\tau) + m(\tau)\, q(\tau), \\ &\quad q(\tau') \propto \pi(\tau')\, m(\tau'), \end{aligned} \tag{131}$$

which preserves the stationary shares $\pi(\tau)$ exactly. The implied quarterly exit probabilities, reported in the last column of Table 5, are 0.45 percent for the two high-wage types and 0.05 percent for the remaining types. Types are therefore highly persistent but not permanent: every household can move between types, and the restriction that matters for the mechanism is that no type carries a belief. Persistent cross-sectional differences in expected returns arise endogenously from households' accumulated investment experiences.

Beliefs. Belief heterogeneity is not assigned through types. All households start from the same prior and update expected stock returns using their own realized portfolio experience according to the learning process of Section 3.3. The household-specific belief ζ_t^i is therefore an endogenous state variable, and the full idiosyncratic state at a given date contains the type, financial wealth, the lagged risky share, and the belief,

$$\left(\tau_t^i, A_{t-1}^i, v_{t-1}^i, \zeta_t^i \right). \tag{132}$$

TABLE 5. **Household Types**

Type	Share (%)	Wage / mean	β (quarterly)	γ
1	39.1	0.52	0.98370	13.00
2	12.2	2.52	0.98909	2.60
3	22.0	1.05	0.98474	5.50
4	16.7	0.37	0.98345	9.50
5	9.0	0.73	0.98422	7.50
6	1.0	12.90	0.98959	1.45

Notes. Each type determines a household's wage level, discount factor, and relative risk aversion. Wages are expressed relative to the population mean; the wage-scale multiplier in the simulation is 6.79. The last column indicates an exit probability that has not been recomputed for the implemented quarterly type transition. Beliefs are not assigned by type and evolve separately through learning from portfolio experience.

TABLE 6. **Quantitative Targets**

Moment	Data	Moment	Data
Saving rate	6.7%	Wealth–income ratio	491%
Mean risky share	30.5%	Wealth-weighted risky share	70.8%
Mean real safe rate	0.75%	Mean real stock return	9.29%

Notes. Distributional moments by wealth group additionally include wealth, income and wage shares, capital-income shares, risky portfolio shares, expected stock returns, and realized portfolio returns. Forecast errors are not targeted.

This separation is important for interpretation. Any persistent relationship between wealth and expected returns in the stationary distribution is generated by the interaction of portfolio choice, realized returns, and learning; it is not imposed by assigning high expected returns to patient, risk-tolerant, or high-wage households.

Targets and estimation. The common wage–dividend scale $\bar{w}$ and the supply of the safe asset $\bar{b}$ are chosen by simulated method of moments. Table 6 summarizes the principal aggregate moments used to discipline the quantitative model. We additionally use distributional moments by wealth group—wealth, income, wages, capital-income shares, risky shares, expected stock returns, and portfolio returns—to assess the joint cross-sectional fit.

6.3. Quantitative Performance

The stationary quantitative results are organized along three dimensions: aggregate quantities, cross-sectional distributions, and the dynamics of beliefs. We also compare the baseline equilibrium with a counterfactual in which the learning signal is the common market return. Some distributional moments enter the calibration, while the fixed-effect statistics provide additional evidence on the structure of simulated beliefs.

TABLE 7. **Parameter Values**

Parameter	Description	Value
Panel A. Types, income, and assets (quarterly)		
$\beta(\tau), \gamma(\tau), \varepsilon(\tau), \pi(\tau)$	Type-specific patience, risk aversion, wages and weights	Table 5
$\bar{\beta}$	Population-weighted geometric mean discount factor (annual)	0.94
$\bar{w}$	Common multiplier on wage-to-dividend ratios	6.79
$\bar{g}_D$	Mean quarterly gross dividend growth	1.01
σ_D	Quarterly log dividend-growth volatility	0.03
$\bar{S}$	Supply of risky trees	1.00
$\bar{b}$	Supply of one-period safe trees	21.50
$\bar{v}$	Maximum risky portfolio share	1.10
$\underline{v}, \underline{a}$	Minimum risky position and saving level	0, 0
ϕ	Type-transition exit-hazard scale	0.55
σ_{HKS}	Wealth-rank-dependent annual idiosyncratic equity-return SD	0.031–0.05
Panel B. Beliefs and learning		
g_L	Quarterly learning gain	0.001
$\sigma_{S,\text{subj}}$	Subjective quarterly log stock-return volatility	0.10
$\bar{R}_B^{\text{subj}}$	Subjective quarterly gross safe return	1.0019
ρ_x, σ_x	Idiosyncratic transitory belief process	0.50, 0.00
σ_{survey}	Annual survey reporting noise (percentage points)	1.00
$[\zeta_{\min}, \zeta_{\max}]$	Numerical quarterly gross-belief support	[1.0080, 1.0220]
$\rho, \bar{Z}, \sigma_Z$	Bayesian latent-state persistence, mean and innovation SD	[0.9950,1.00, 0.0070]
φ_B, σ_B	Subjective safe-quote persistence and volatility	[0.01, 0.01]
Panel C. Numerical solution and simulation		
N	Households	50,000
T	Quarters (burn-in)	20,000 (5,000)
	Neural-network hidden layers and units	3×96
	Gauss–Hermite nodes: stock × dividend growth	32×3
	Pseudo-RAND panels: replications × agents × quarters	$200 \times 2{,}000 \times 36$

Notes. Parameters describe the numerical implementation used for Tables 8–11 and 10.

TABLE 8. **Aggregate Quantitative Performance**

Moment	Model	Data
Household aggregates		
Wealth–income ratio (%)	235.89	491.00
Average risky asset share (%)	24.73	30.50
Risky assets / aggregate wealth (%)	69.43	70.80
Saving rate (%)	1.69	6.70
Aggregate returns and growth, annual (%)		
Mean stock return	11.54	9.29
Mean safe rate	−0.25	0.75
Equity premium (pp)	11.79	8.54
SD stock log return, annualized	5.93	13.10
Mean dividend growth	2.77	2.84
SD dividend growth	6.32	6.40

Notes. Annual returns are measured from nonoverlapping four-quarter compounded returns; the reported stock-volatility measure is the annualized standard deviation of log returns. A comparable empirical figure for that precise measure is not supplied in the current quantitative inputs.

6.3.1. Aggregate Moments

Table 8 evaluates the aggregate quantitative performance of the baseline stationary equilibrium. The model performs particularly well in reproducing aggregate portfolio composition: risky assets account for 69.4 percent of total wealth, remarkably close to the 70.8 percent observed in the data. It also closely matches the mean and volatility of dividend growth. However, several discrepancies remain. Most notably, the model generates substantially less aggregate wealth accumulation, with a wealth–income ratio and saving rate well below their empirical counterparts. On the asset-pricing side, the model produces a sizeable equity premium, although it somewhat overstates its magnitude, while stock return volatility remains less than half its empirical value. These discrepancies suggest that the model captures key features of portfolio allocation and average asset returns more successfully than the aggregate level of wealth accumulation and the dynamics of asset prices.

6.3.2. Cross-Sectional Distributions

Table 9 reports the central quantitative comparison. We divide households into the top 1 percent, the next 9 percent, the middle 40 percent, and the bottom 50 percent of the financial-wealth distribution.

The model places 39.31 percent of financial wealth with the top 1 percent and 34.18 percent with the next 9 percent, for a top-decile share of 73.51 percent (75.70 percent in the data). The bottom half owns 5.23 percent, compared with 1.40 percent in the data. Income and wage shares are substantially less concentrated than wealth, as shown in Table 9.

Portfolios and beliefs exhibit a strong wealth gradient. The risky share increases from 13.50 percent for the bottom half to 93.43 percent for the top 1 percent. Expected

TABLE 9. **Distributional Moments: Model and Data**

	Model				Data			
	Top 1%	Next 9%	Middle 40%	Bottom 50%	Top 1%	Next 9%	Middle 40%	Bottom 50%
Wealth share	39.31	34.18	21.27	5.23	37.60	38.10	22.80	1.40
Income share	17.59	24.07	36.11	22.22	18.40	24.70	35.20	21.60
Wage share	11.46	22.94	40.14	25.45	12.30	21.60	38.50	27.50
Risky asset share	93.43	64.55	33.82	13.50	77.30	62.30	39.60	18.10
Expected stock return	7.93	7.80	6.51	4.90	7.88	6.78	5.39	4.77
Portfolio return	11.57	7.03	3.33	1.07	8.30	6.12	3.63	3.40

Notes. Wealth, income, and wages are shares of their respective totals. Groups are ranked by current wealth. Risky asset shares are fractions of gross financial assets. Expected and portfolio returns are annual percentages.

stock returns rise from 4.90 percent to 7.93 percent across the same groups, compared with 4.77 and 7.88 percent in the data. Realized portfolio returns similarly rise from 1.07 to 11.57 percent. The wealthiest groups are less pessimistic relative to the model's mean stock return of 11.54 percent, and accumulate wealth through larger risky exposures and more favorable portfolio performance.

6.3.3. Belief Dynamics

The previous comparisons concern the cross-sectional mean of beliefs across wealth groups. A more demanding test is whether the model reproduces the *within-group* persistence of disagreement emphasized in Section 2.2.

Recall that all households begin with the same prior. Persistent disagreement is therefore not initialized through household-specific fixed effects. It must emerge endogenously because households experience different portfolio returns and update their beliefs from those experiences. We apply to simulated beliefs the same fixed-effect decomposition used in the RAND panel,

$$\mathbb{E}_{it}[R_{t+1}] = \alpha_i + \mu_t + \varepsilon_{it}. \tag{133}$$

To keep the exercise comparable to the data, we mimic the structure of the RAND sample rather than using the full simulated population: we draw 200 pseudo-panels, each following 2,000 randomly sampled households for nine years (36 quarters) over non-overlapping windows of the stationary simulation, estimate the decomposition in each pseudo-panel, regress the estimated individual effects on average wealth over the same window as in Section 2.2, and report averages across pseudo-panels.

The simulated survey-observed belief panel displays substantial persistence. Individual fixed effects explain 63.5 percent of the variation in expected returns, compared with 57.3 percent in RAND. Time fixed effects explain 0.05 percent in the model, compared with 12.8 percent in RAND. Wealth explains 45.9 percent of the estimated individual effects, compared with 8.9 percent in the data. These results capture persistent disagreement but imply a much tighter individual-level association between expectations and wealth than observed in RAND.

The comparison confirms that the model produces a steep relationship between average beliefs and wealth as well as persistent cross-sectional disagreement. The individual-level wealth gradient is stronger than in RAND, while simulated common

TABLE 10. **Persistent Disagreement: Data and Model**

	Data	Model
R^2: individual fixed effects	57.3%	63.5%
R^2: time fixed effects	12.8%	0.05%
R^2: log mean wealth $\rightarrow \widehat{\alpha}_i$	8.9%	45.9%

time effects are smaller.

The coexistence of persistent disagreement and a wealth gradient arises because wealth is not the primitive source of disagreement. Households with similar wealth can arrive there with different histories of portfolio exposure and realized returns and hence hold different beliefs. At the same time, those histories affect wealth accumulation, so the conditional distribution of beliefs gradually becomes ordered by wealth. Persistent disagreement and a wealth–belief gradient therefore emerge jointly rather than being assigned by type.

6.4. Measuring the effects of learning from own experience.

We compare the baseline economy, in which households update beliefs from their own realized portfolio returns, with an economy in which all households use the common traded stock-market return as the learning signal. Both simulations share the same initial prior, learning gain, learning intercept, preference-type process, private payoff shocks, asset supplies, wage calibration and exogenous innovations. A separate household policy is solved for each learning rule, and asset prices, returns, beliefs and wealth distributions adjust endogenously in each equilibrium. This comparison identifies the *total general-equilibrium effect of learning from own portfolio experience*; it does not hold the cross-sectional average of beliefs fixed.

Table 11 shows a large distributional effect. Learning from own experience raises the top-decile wealth share from 61.77 to 73.51 percent, or 11.74 percentage points. The top-1-percent share rises from 31.16 to 39.36 percent, while the bottom-half share falls from 11.82 to 5.23 percent. With common-market learning, the mean expected return increases from 5.84 to 9.08 percent and cross-sectional dispersion in structural stock-return beliefs disappears. Risky shares are larger across all wealth groups in the counterfactual; thus the concentration result does not simply reflect greater average equity participation in the baseline. Both the endogenous sorting of portfolios and the changes in aggregate beliefs and equilibrium returns are part of the experiment.

The differences extend beyond the household distribution, as changes in the information available to investors also affect equilibrium asset prices and returns. Under the common market signal, households no longer learn from their own portfolio experiences, eliminating the persistently pessimistic beliefs of investors with poor return histories. As a result, more households perceive stocks as attractive investments, increasing aggregate demand for equities and bidding up their prices. Table 12 shows that the equilibrium price–dividend ratio rises from 48.90 to 59.08, while financial wealth per household increases from 70.43 to 81.69 dividend units.

TABLE 11. **Distributional Effects of Learning from Own Investment Experience**

	Own portfolio returns				Common market return			
	Top 1%	Next 9%	Middle 40%	Bottom 50%	Top 1%	Next 9%	Middle 40%	Bottom 50%
Wealth share	39.31	34.18	21.27	5.23	31.13	30.64	26.42	11.82
Income share	17.59	24.07	36.11	22.22	21.04	24.35	34.13	20.48
Consumption share	17.92	23.80	35.95	22.33	21.70	24.59	33.58	20.13
Wage share	11.46	22.94	40.14	25.45	12.14	23.20	39.20	25.47
Risky asset share	93.43	64.55	33.82	13.50	98.15	77.05	48.59	19.91
Saving share κ	95.41	92.91	79.74	68.72	94.96	94.16	89.97	87.81
Expected stock return	7.93	7.80	6.51	4.90	9.08	9.08	9.08	9.08
Portfolio return	11.57	7.03	3.33	1.07	12.44	3.30	−5.70	−13.24

Notes. Shares of wealth, income, consumption and wages are percentages of their respective aggregates. Saving shares and risky shares are percentages; return entries are annual percentages. Wealth groups are re-ranked each period in each equilibrium. The mean learning intercept and initial prior are shared, but equilibrium mean beliefs are endogenous.

Higher equity valuations, in turn, compress expected stock returns, which decline from 11.54 to 9.96 percent. The adjustment in the safe asset market is even more pronounced, with the equilibrium safe rate falling from −0.25 to −18.00 percent, raising the equity premium from 11.79 to 27.96 percentage points. Despite higher aggregate financial wealth, average consumption declines from 7.34 to 6.26 dividend units, while wealth inequality falls substantially, with the Gini coefficient decreasing from 0.803 to 0.678. These results highlight how heterogeneous learning experiences shape not only the distribution of wealth but also equilibrium asset valuations and risk premia.

6.4.1. Decomposing the effects of learning from experience.

Learning from own investment experience affects wealth inequality through two channels. First, heterogeneous experiences generate differences in beliefs, which influence households' saving and portfolio decisions. Second, these decisions affect aggregate asset demand and hence equilibrium stock and safe returns. To assess the contribution of each channel, we conduct a partial-equilibrium decomposition based on the two general-equilibrium economies: one in which households learn from their own portfolio returns (H), and another in which they learn from the common market return (F).

We first record the equilibrium stock and safe return histories generated by each economy, denoted by R^H and R^F. We then simulate households under both return histories, using either the heterogeneous-learning rule or the common-market-learning rule. This produces four experiments. In the two diagonal cases, (F, R^F) and (H, R^H), the simulations reproduce their respective general-equilibrium wealth distributions up to numerical precision. In the remaining cases, we hold the aggregate return history fixed while changing how households learn. Saving, portfolio choices, and beliefs remain endogenous throughout, but asset markets need not clear in the partial-equilibrium counterfactuals.

Table 13 reports the results. Learning from own experience increases the top-decile wealth share from 61.77 to 73.51 percent, a total general-equilibrium effect of 11.74 percentage points. However, the contribution of each channel depends

TABLE 12. **Aggregate Effects of Removing Learning from Own Experience**

Moment	Own returns	Market signal	Difference
Top 1% wealth share (%)	39.36	31.16	−8.20
Top 10% wealth share (%)	73.51	61.77	−11.74
Bottom 50% wealth share (%)	5.23	11.82	+6.59
Gini coefficient	0.803	0.678	−0.125
Mean expected stock return (annual %)	5.84	9.08	+3.24
SD expected stock return (annual pp)	1.30	0.00	−1.30
Mean price–dividend ratio	48.90	59.08	+10.18
Mean safe-tree price	1.001	1.051	+0.050
Mean stock return (annual %)	11.54	9.96	−1.58
Mean safe return (annual %)	−0.25	−18.00	−17.75
Equity premium (annual pp)	**11.79**	**27.96**	+16.17
SD stock log return (annualized %)	5.93	5.81	−0.12
Mean financial wealth (dividend units)	70.43	81.69	+11.26
Mean consumption (dividend units)	7.34	6.26	−1.08
Mean cash income (dividend units)	7.46	6.39	−1.07
Saving rate (%)	1.69	2.07	+0.38
Wealth–income ratio (%)	235.89	319.59	+83.70
Wealth-weighted risky share (%)	69.43	72.32	+2.89

Notes. Difference is the common-market-signal economy minus the own-experience economy, in the units shown. Returns are computed from nonoverlapping four-quarter compounded outcomes. All asset prices, returns, portfolio choices, wealth and beliefs are endogenous in both equilibria.

strongly on the other. Holding the common-learning return history fixed, introducing heterogeneous learning reduces the top-decile share by 0.44 percentage points. By contrast, under heterogeneous learning, replacing the common-learning returns with the baseline equilibrium returns increases it by 12.18 percentage points. The same complementarity appears in the Gini coefficient, which increases from 0.678 to 0.803 across the two general-equilibrium benchmarks.

The results indicate substantial interaction between household learning and equilibrium asset returns. Indeed, the top-decile wealth share is lower in both mixed experiments than in the heterogeneous-learning benchmark. Averaging the contributions across the two possible decomposition orders assigns 8.58 percentage points of the total 11.74-point increase to household learning and portfolio decisions, and 3.16 percentage points to changes in aggregate returns. These figures are accounting contributions rather than independently identified causal effects: the two channels reinforce each other, and the return channel captures changes in both stock and safe returns, not only the equity premium. Moreover, the experiment does not hold average beliefs fixed. Nevertheless, it shows that the distributional consequences of learning from experience depend importantly on its general-equilibrium effects on asset returns.

TABLE 13. **Decomposing the Effects of Learning from Experience**

	Common learning (F)		Own-experience learning (H)	
	R^F	R^H	R^F	R^H
Gini coefficient	0.678	0.636	0.644	0.803
Top 1% wealth share (%)	31.16	26.56	36.05	39.36
Top 10% wealth share (%)	61.77	55.91	61.33	73.51
Bottom 50% wealth share (%)	11.82	13.09	14.78	5.23

Notes. Columns combine the household learning rule with the aggregate return history generated in each general-equilibrium economy. Saving, portfolio choice, and learning remain endogenous. The first and last columns report the general-equilibrium benchmarks from Section 6.3.4. The intermediate columns report partial-equilibrium counterfactuals from a separate simulation.

6.4.2. The effect of the equity premium on wealth inequality.

The preceding general-equilibrium experiment changes both the distribution of beliefs and the asset returns that households face. To isolate the role of the equity premium in the accumulation mechanism, we conduct a partial-equilibrium experiment conditional on the aggregate return history generated by the baseline model. We first record the realized quarterly stock and safe returns in the general-equilibrium simulation and feed these same returns back into the household problem, without imposing asset-market clearing. With the original household policies, initial conditions, and innovations, this replay reproduces the general-equilibrium paths of household wealth, portfolios, and beliefs up to numerical precision. It therefore provides a benchmark against which to change the return process while leaving the household side of the model unchanged.

In the counterfactual, we eliminate the *average* realized equity premium while preserving stock-market risk. Specifically, we replace the aggregate gross stock return R^S_{t+1} with

$$\widetilde{R}^S_{t+1} = \lambda R^S_{t+1}, \tag{134}$$

where the constant λ is chosen so that mean annual stock and safe returns, computed from nonoverlapping four-quarter returns, coincide. The realized safe-return sequence is unchanged. Because the same constant multiplies every quarterly gross stock return, the volatility and serial dependence of log stock returns are preserved; the counterfactual does not make stocks riskless. We also keep the same dividend-growth sequence, household type transitions, and underlying idiosyncratic innovations, as well as all preference and learning parameters. Households continue to use the original subjective decision rules, but their saving, portfolio choices, wealth, and beliefs evolve with their counterfactual investment experiences. Thus, the experiment measures the partial-equilibrium effect of the realized equity premium, including households' endogenous responses, rather than the effect of holding portfolios fixed.

Table 14 reports the results. The benchmark return history features an annual equity premium of 13.31 percentage points; the counterfactual reduces it to zero. With no premium, the Gini coefficient is 0.666, compared with 0.814 in the benchmark.

TABLE 14. **The Equity Premium and Wealth Inequality**

Moment	Zero premium	GE premium	Difference
Gini coefficient	0.666	0.814	+0.148
Top 1% wealth share (%)	34.97	39.65	+4.68
Top 10% wealth share (%)	63.04	76.07	+13.03
Bottom 50% wealth share (%)	13.54	5.02	−8.52

Notes. Both columns replay the same general-equilibrium aggregate return innovations, dividend growth, and underlying household shocks, with endogenous saving, portfolio choice, and learning. The zero-premium economy rescales quarterly gross stock returns by a constant to equate mean annual stock and safe returns, preserving stock log-return volatility. Difference is the GE premium column minus the zero-premium column; differences in wealth shares are percentage points. Asset prices do not clear markets in the counterfactual.

The wealth share of the top 10 percent falls from 76.07 to 63.04 percent, a difference of 13.03 percentage points. The top-1-percent share falls by 4.68 percentage points, while the bottom-half share increases from 5.02 to 13.54 percent. The equity premium therefore makes a substantial contribution to wealth concentration, especially when comparing the top decile with the rest of the distribution. Nevertheless, the top decile still owns 63.04 percent of wealth when the premium is zero, reflecting the other sources of heterogeneity retained in the model.

The mechanism operates through heterogeneous exposure to aggregate equity returns. Wealthier households allocate a larger share of savings to stocks and therefore benefit disproportionately when stocks earn a positive premium. These differences in realized portfolio returns affect subsequent wealth and, through learning from personal investment experience, future beliefs and portfolio choices. Consistent with this feedback, the wealth-weighted risky share rises from 29.91 percent in the zero-premium economy to 69.59 percent in the benchmark. The comparison thus captures both the direct return advantage of equity holders and the subsequent adjustment of household portfolios and beliefs.

6.5. Historical Transitional Dynamics

The stationary distribution establishes that heterogeneous expectations can substantially increase long-run wealth concentration. It does not establish whether the mechanism is quantitatively relevant along the particular sequence of returns experienced by U.S. households. We therefore conduct a historical transition exercise covering 1960–2024.

The exercise is deliberately conditional on the realized aggregate asset-market history. We feed the model observed U.S. stock prices and dividends, together with the calibrated price moments entering households' subjective models, and ask how the same aggregate sequence is propagated through heterogeneous household portfolios and beliefs. We initialize the economy from the observed wealth distribution at the beginning of the sample. Initial subjective expected capital gains are chosen to reproduce the wealth gradient measured in RAND. Thereafter, portfolio choices and wealth accumulation are endogenous, and beliefs evolve through the learning

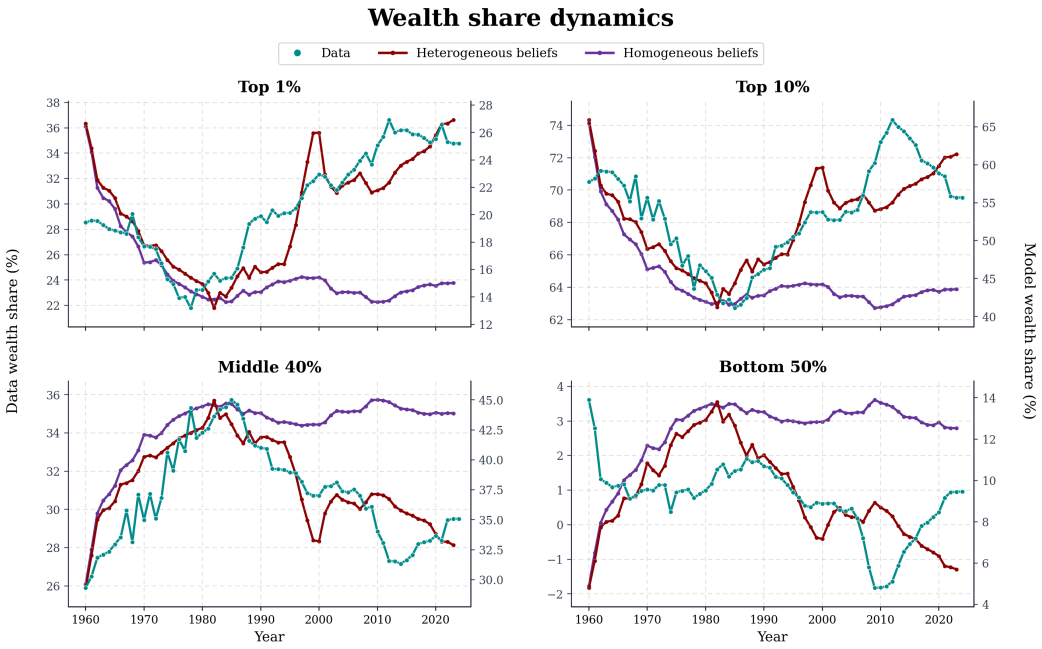

FIGURE 4. **Historical Transitional Dynamics, 1960–2024.** The figure compares observed U.S. wealth shares with model-implied wealth shares under heterogeneous and homogeneous beliefs, conditional on the realized U.S. asset-price and dividend history.

mechanism.

This exercise is not intended to explain the historical path of U.S. stock prices. By conditioning on realized prices, we instead hold the aggregate return history fixed and isolate the distributional propagation mechanism. The relevant comparison is between households exposed to the same market history under heterogeneous learning and under common beliefs.

Figure 4 reports the resulting wealth-share dynamics. The common-belief economy generates comparatively limited movement in the upper tail. Heterogeneous learning produces substantially larger changes in top wealth shares and moves the model in the direction of the historical increase in U.S. wealth concentration. The mechanism is especially visible in the top decile: market experiences generate persistent differences in expected returns and portfolio exposure, and these differences compound over long horizons. Correspondingly, wealth shifts away from the middle and lower parts of the distribution.

The historical exercise does not reproduce every movement in U.S. wealth shares, nor should it. Taxes, entrepreneurship, housing, inheritances, changes in labor income inequality, and other forces that are absent from the model also changed substantially over this period. The purpose of the exercise is narrower: to ask how much movement in the wealth distribution can arise from heterogeneous responses to a common realized asset-market history.

Taken together, the stationary and historical exercises point to the same economic mechanism. Heterogeneous expectations do not create inequality by simply making every household save more or hold more stocks. They change *who* takes financial risk, *when* they take it, and how realized portfolio outcomes feed back into

subsequent expectations. This endogenous sorting creates persistent differences in portfolio returns that compound over time. In the stationary experiment, learning from own portfolio returns raises the top-decile wealth share by 11.74 percentage points relative to learning from the common stock-market return. Along the realized postwar U.S. market history, it also generates substantially more movement in top wealth shares than the corresponding common-belief economy.

7. Conclusion

This paper studies heterogeneous expectations as a source of wealth inequality. Using the RAND American Life Panel, we document large and persistent disagreement about stock returns and a systematic relationship between beliefs, portfolios, and wealth: wealthier households are more optimistic, hold more risky assets, and, on average, hold more accurate beliefs. We develop a heterogeneous-agent model in which these patterns arise from learning from one's own investment experience. Households with different portfolios experience different returns; these experiences shape subsequent beliefs, which in turn affect portfolio choices and future returns. The resulting feedback between beliefs, risk taking, and wealth accumulation generates persistent differences in portfolio returns and amplifies wealth concentration. In the stationary quantitative economy, replacing personal portfolio returns with a common market-return learning signal lowers the top-decile wealth share by 11.74 percentage points and raises the equilibrium equity premium by 16.17 percentage points. This comparison captures the general-equilibrium contribution of learning from own experience, allowing mean beliefs and asset prices to adjust endogenously.

The analysis also has a methodological implication. Allowing households to hold subjective beliefs need not make heterogeneous-agent models with aggregate risk harder to solve. Under Internal Rationality, households forecast personal investment returns and safe-rate quotes directly rather than forecasting the high-dimensional distribution of households from which equilibrium prices emerge. This separates individual optimization from market clearing and motivates the Two-Stage Deep Learning algorithm developed in the paper. In the first stage, household policies are solved under their subjective probability measures; in the second, those policies are aggregated and equilibrium prices are obtained from market clearing. The approach accommodates aggregate risk, portfolio choice, occasionally binding constraints, and rich forms of household heterogeneity without requiring households to know or forecast a law of motion for the cross-sectional distribution. This provides a flexible route for incorporating empirically realistic expectations into quantitative heterogeneous-agent models.

Important questions remain. More and better data are needed to measure expectations over longer horizons and over longer periods of households' lives, and to understand how beliefs interact with preferences rather than treating the two as conceptually separate primitives. Like most quantitative heterogeneous-agent models, our framework still takes some ex-ante heterogeneity—in discount factors and risk aversion, for example—as given and studies the distributions that subsequently emerge. We have tried to go one layer deeper by providing a mechanism through

which persistent heterogeneity in beliefs can emerge endogenously from households' experiences. But a deeper account of inequality requires understanding where differences in beliefs and preferences ultimately come from, how they interact, whether they can be understood as manifestations of broader cultural or social differences, and how these differences are reproduced across time and generations. Progress on these questions will require combining richer longitudinal data with models in which preferences, beliefs, and economic circumstances jointly evolve. Computationally, the framework developed here is sufficiently flexible to accommodate richer environments, including HANK models. It can also be used to study the distributional consequences of policies that operate directly on wealth accumulation and exposure to aggregate risk, including wealth taxation and sovereign wealth funds. Such applications may help connect the deeper origins of household heterogeneity to the design of policies aimed at limiting its amplification into persistent wealth inequality.

References

Acedański, J. (2017). "Heterogeneous expectations and the distribution of wealth". In: *Journal of Macroeconomics* 53, pp. 162–175.

Adam, K., J. Beutel, A. Marcet, and S. Merkel (2015). "Can a Financial Transaction Tax Prevent Stock Price Booms?" In: *Journal of Monetary Economics* 76, S90–S109.

Adam, K. and A. Marcet (2011). "Internal rationality, imperfect market knowledge and asset prices". In: *Journal of Economic Theory* 146.3, pp. 1224–1252.

Adam, K., A. Marcet, and J. Beutel (2017). "Stock price booms and expected capital gains". In: *American Economic Review* 107.8, pp. 2352–2408.

Adam, K., A. Marcet, and J. P. Nicolini (2016). "Stock Market Volatility and Learning". In: *The Journal of Finance* 71.1, pp. 33–82.

Aiyagari, S. R. (1994). "Uninsured Idiosyncratic Risk and Aggregate Saving". In: *The Quarterly Journal of Economics* 109.3, pp. 659–684.

Bach, L., L. E. Calvet, and P. Sodini (2020). "Rich pickings? Risk, return, and skill in household wealth". In: *American Economic Review* 110.9, pp. 2703–2747.

Belda, P. (2026). *Capital Gains Taxation and Asset Price Volatility*. Staff Working Paper 1200. Bank of England.

Benhabib, J. and A. Bisin (2018). "Skewed Wealth Distributions: Theory and Empirics". In: *Journal of Economic Literature* 56.4, pp. 1261–1291.

Benhabib, J., A. Bisin, and M. Luo (2019). "Wealth Distribution and Social Mobility in the US: A Quantitative Approach". In: *American Economic Review* 109.5, pp. 1623–1647.

Bordalo, P., N. Gennaioli, Y. Ma, and A. Shleifer (2020). "Overreaction in macroeconomic expectations". In: *American Economic Review* 110.9, pp. 2748–2782.

Broer, T., A. Kohlhas, K. Mitman, and K. Schlafmann (2022). "Expectation and Wealth Heterogeneity in the Macroeconomy". Working paper.

Carceles-Poveda, E. and C. Giannitsarou (2008). "Asset Pricing with Adaptive Learning". In: *Review of Economic Dynamics* 11.3, pp. 629–651.

Castañeda, A., J. Díaz-Giménez, and J.-V. Ríos-Rull (2003). "Accounting for the US Earnings and Wealth Inequality". In: *Journal of Political Economy* 111.4, pp. 818–857.

Chiu, K. Y. (2024). "Heterogeneous Beliefs about Stock Returns and Wealth Inequality". Job Market Paper.

Cioffi, R. A., S. Hurtado, and G. Nuño (2026). "Asset Pricing with Heterogeneous Agents under Limited Information". Working paper.

D'Acunto, F., U. Malmendier, J. Ospina, and M. Weber (2021). "Exposure to grocery prices and inflation expectations". In: *Journal of Political Economy* 129.5, pp. 1615–1639.

De Nardi, M. and G. Fella (2017). "Saving and Wealth Inequality". In: *Review of Economic Dynamics* 26, pp. 280–300.

De Nardi, M., G. Fella, and G. Paz Pardo (2016). *The Implications of Richer Earnings Dynamics for Consumption and Wealth*. NBER Working Paper 21917. National Bureau of Economic Research.

Fagereng, A., L. Guiso, D. Malacrino, and L. Pistaferri (2020). "Heterogeneity and persistence in returns to wealth". In: *Econometrica* 88.1, pp. 115–170.

Fernández-Villaverde, J. and O. Levintal (2024). *The Distributional Effects of Asset Returns*. NBER Working Paper 32182. National Bureau of Economic Research.

Fidrmuc, J., C. Hainz, and W. Hölzl (2024). "Individual credit market experience and beliefs about bank lending policy: evidence from a firm survey". In: *The Scandinavian Journal of Economics* 126.2, pp. 387–414.

Gabaix, X., J.-M. Lasry, P.-L. Lions, and B. Moll (2016). "The Dynamics of Inequality". In: *Econometrica* 84.6, pp. 2071–2111.

Giglio, S., M. Maggiori, J. Stroebel, and S. Utkus (2021). "Five facts about beliefs and portfolios". In: *American Economic Review* 111.5, pp. 1481–1522.

Gomez, M. (2025). "Wealth inequality and asset prices". In: *Review of Economic Studies* 92.6, pp. 3924–3967.

Hubmer, J., P. Krusell, and A. A. Smith Jr. (2021). "Sources of US Wealth Inequality: Past, Present, and Future". In: *NBER Macroeconomics Annual* 35.1, pp. 391–455.

Huggett, M. (1993). "The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies". In: *Journal of Economic Dynamics and Control* 17.5–6, pp. 953–969.

Jiang, Z., H. Liu, C. Peng, and H. Yan (2025). "Investor memory and biased beliefs: Evidence from the field". In: *The Quarterly Journal of Economics* 140.4, pp. 2749–2804.

Kohlhas, A. N. and A. Walther (2021). "Asymmetric attention". In: *American Economic Review* 111.9, pp. 2879–2925.

Kuchler, T. and B. Zafar (2019). "Personal experiences and expectations about aggregate outcomes". In: *The Journal of Finance* 74.5, pp. 2491–2542.

Lei, X. (2026). "How do stock market experiences shape wealth inequality?" In: *Journal of Economic Dynamics and Control*, p. 105337.

Malmendier, U. and S. Nagel (2011). "Depression babies: Do macroeconomic experiences affect risk taking?" In: *The Quarterly Journal of Economics* 126.1, pp. 373–416.

— (2016). "Learning from inflation experiences". In: *The Quarterly Journal of Economics* 131.1, pp. 53–87.

Martin, I. and D. Papadimitriou (2022). "Sentiment and speculation in a market with heterogeneous beliefs". In: *American Economic Review* 112.8, pp. 2465–2517.

Stachurski, J. and A. A. Toda (2019). "An impossibility theorem for wealth in heterogeneous-agent models with limited heterogeneity". In: *Journal of Economic Theory* 182, pp. 1–24.

Timmermann, A. (1996). "Excess Volatility and Predictability of Stock Prices in Autoregressive Dividend Models with Learning". In: *Review of Economic Studies* 63.4, pp. 523–557.

Vissing-Jørgensen, A. (2003). "Perspectives on Behavioral Finance: Does "Irrationality" Disappear with Wealth? Evidence from Expectations and Actions". In: *NBER Macroeconomics Annual* 18, pp. 139–194.

Von Gaudecker, H.-M. and A. Wogrolly (2022). "Heterogeneity in households' stock market beliefs: Levels, dynamics, and epistemic uncertainty". In: *Journal of Econometrics* 231.1, pp. 232–247.

Winkler, F. (2020). "The Role of Learning for Asset Prices and Business Cycles". In: *Journal of Monetary Economics* 114, pp. 42–58.

Xavier, I. (2021). "Wealth inequality in the US: the role of heterogeneous returns". In: *Available at SSRN 3915439*.

Yang, Y., C. Wang, A. Schaab, and B. Moll (2025). "Structural reinforcement learning for heterogeneous agent macroeconomics". In: *arXiv preprint arXiv:2512.18892*.

APPENDIX

A. RAND Data and Measurement Details

This appendix provides additional details on the RAND sample and the construction of the belief variables used in Section 2.

Survey questions. Respondents are asked about the value one year ahead of mutual-fund shares invested in blue-chip stocks such as those in the Dow Jones Industrial Average. The three questions relevant for our analysis elicit: (i) the probability that the shares will be worth more than they are today; (ii) the probability that they will have increased in value by more than 20 percent; and (iii) the probability that they will have fallen in value by more than 20 percent. The three answers are sufficient to recover the four probabilities in equation (4). We exclude observations whenever the implied probability of one of the four mutually exclusive regions is negative.

The survey questions refer to changes in the value of the fund over the following year. We therefore treat the elicited object as the one-year stock-market capital gain and construct realized returns for the forecast-error exercises on the corresponding horizon and return basis.

Expected-return construction. The probability questions identify intervals rather than a complete continuous distribution. To obtain the scalar measure m_{it} used in the fixed-effect and forecast-error exercises, we follow von Gaudecker and Wogrolly (2022) and associate each interval with a representative gross return x_j. The resulting expected return is the probability-weighted average in equation (5). The representative returns are disciplined by the historical distribution of stock-market returns. This procedure is a measurement device: it maps a coarse elicited distribution into a scalar expectation while preserving the ordering of households implied by their probability assessments. Because the fractional-logit exercise uses $p_{j,it}$ directly, its results do not depend on this mapping.

Panel length. The ALP is unbalanced. For analyses that require persistent individual variation, we restrict attention to respondents with repeated observations. Table 15 reports the number of households satisfying alternative minimum-observation requirements. Requiring at least 10 observations leaves 2,321 households. The regression samples in Section 2 can be smaller because each exercise imposes its own requirements for internally consistent probability answers and nonmissing wealth, portfolio, and control variables.

Wealth in the fixed-effect exercise. Equation (7) uses the respondent's average wealth across the RAND observations entering the household-level exercise. This averaging is deliberate. The estimated fixed effect $\widehat{\alpha}_i$ is a single persistent component for each respondent, so the appropriate descriptive wealth counterpart is also household-level. Using average wealth reduces the influence of transitory wealth

TABLE 15. RAND Panel Size by Minimum Number of Observations

Minimum observations	3	4	5	6	10	30	50
Number of households	3,027	2,969	2,882	2,764	2,321	785	303

Notes. The table reports the number of RAND ALP respondents satisfying alternative minimum numbers of observations in the 2008–2017 panel. Exercise-specific samples additionally require the variables used in the corresponding regression.

movements and avoids treating repeated observations of wealth as independent explanations for the same estimated fixed effect. It does not remove the possibility of reverse causality from beliefs to wealth, and we accordingly interpret the resulting wealth gradient as descriptive.

B. Forecast-Error Predictability

The main text emphasizes persistent cross-sectional disagreement because the paper studies long-run portfolio choice and wealth accumulation. For completeness, this appendix shows that RAND expectations also exhibit the familiar time-series predictability of forecast errors documented in other survey data.

We consider two tests. First, we ask whether the current market return predicts the subsequent forecast error:

$$R_{t+k} - \mathbb{E}_{it}[R_{t+k}] = \alpha_i + bR_t + \varepsilon_{it}. \tag{135}$$

Second, we ask whether revisions in the average survey forecast predict subsequent individual forecast errors:

$$R_{t+k} - \mathbb{E}_{it}[R_{t+k}] = \alpha_i + d\left(\overline{\mathbb{E}}_t[R_{t+k}] - \overline{\mathbb{E}}_{t-1}[R_{t+k}]\right) + \varepsilon_{it}. \tag{136}$$

The horizon is one year, corresponding to $k = 4$ quarterly waves. Both specifications include individual fixed effects. Under full-information rational expectations, information available at date t should not systematically predict the forecast error, implying $b = d = 0$.

Table 16 reports the results. In column (1), the coefficient on the current market return is −0.13 and statistically significant at the 1 percent level. High current returns are therefore followed by lower realized-minus-expected forecast errors, consistent with households raising their expectations after strong recent market performance. In column (2), the coefficient on the average forecast revision is −0.50 and statistically significant at the 1 percent level: upward revisions in the survey forecast are followed, on average, by more negative forecast errors. When the two predictors are included jointly in column (3), the coefficients are −0.14 and −0.51, respectively. The estimates are therefore not driven by the two predictors proxying for the same variation.

The signs are consistent with extrapolation and overreaction, in line with a broad literature on survey expectations, including Adam, Marcet, and Beutel (2017); Bordalo et al. (2020); Kohlhas and Walther (2021). We treat these results as a

TABLE 16. Forecast-Error Predictability in RAND

	(1)	(2)	(3)
R_t	-0.13^{***}	—	-0.14^{***}
$\bar{\mathbb{E}}_t[R_{t+k}] - \bar{\mathbb{E}}_{t-1}[R_{t+k}]$	—	-0.50^{***}	-0.51^{***}
Individual fixed effects	Yes	Yes	Yes
Observations	64,749	64,749	64,749

Notes. The dependent variable is the individual one-year-ahead forecast error, defined as realized minus expected return. R_t is the current market return. The second regressor is the change in the average survey expectation between adjacent survey waves. The forecast horizon is $k = 4$ quarterly waves. All specifications include individual fixed effects. $^{***}p < 0.01$.

validation and benchmarking exercise rather than as the primary empirical object of the paper. The model is disciplined mainly by the persistence and cross-sectional structure of disagreement documented in Section 2.

C. Fractional-logit regressions for subjective return probabilities

Section 2.3 summarizes the economic magnitude of the wealth–belief relationship using changes in predicted probabilities. For completeness, this appendix reports the coefficient estimates underlying those counterfactuals.

Let $p_{j,it}$ denote the probability assigned by household i in survey wave t to return region

$$j \in \{R < -20\%,\ -20\% < R < 0\%,\ 0\% < R < 20\%,\ R > 20\%\}.$$

For each region, we estimate a separate fractional-logit model,

$$p_{j,it} = \Lambda\Big(\alpha_j + \beta_{W,j} \log W_{it} + \beta_{V,j} V_{it} + \Gamma_j' X_{it}\Big), \tag{137}$$

where $\Lambda(z) = \exp(z)/[1 + \exp(z)]$. As throughout the RAND analysis, W_{it} is the sum of housing and retirement-account wealth and V_{it} is the share of retirement-account wealth invested in stocks (Section 2.1). The vector X_{it} contains homeownership, retirement-account status, housing wealth, income, and age.

The coefficient estimates reinforce the pattern summarized in Table 2. Higher wealth is associated with a lower probability assigned to moderate losses and a higher probability assigned to moderate gains. Its association with the tails of the subjective distribution is considerably smaller: wealth is essentially unrelated to the probability of a return below −20 percent and is only weakly positively associated with the probability of a return above 20 percent. Greater risky exposure exhibits a similar pattern, with a strong shift from moderate losses toward moderate gains and smaller positive associations with the two extreme-return regions.

Because the four probability equations are estimated separately, their fitted values are not mechanically constrained to sum to one. We therefore interpret the estimates bin by bin. The predicted-probability counterfactuals reported in Table 2 provide the economically interpretable summary of these regressions.

TABLE 17. Fractional-Logit Estimates for Subjective Return Probabilities

	$R < -20\%$	$-20\% < R < 0\%$	$0\% < R < 20\%$	$R > 20\%$
Homeowner	0.249	0.460^{**}	-0.401^{**}	-0.329^{*}
	(0.159)	(0.191)	(0.186)	(0.179)
Retirement account	−0.378	1.080^{**}	-0.946^{*}	−0.091
	(0.339)	(0.450)	(0.498)	(0.177)
Log wealth	0.001	-0.088^{***}	0.103^{***}	0.014^{**}
	(0.006)	(0.007)	(0.008)	(0.006)
Risky share	0.0006^{**}	-0.0074^{***}	0.0069^{***}	0.0020^{***}
	(0.0003)	(0.0004)	(0.0003)	(0.0003)
Log housing wealth	−0.020	-0.033^{**}	0.030^{**}	0.023
	(0.013)	(0.016)	(0.015)	(0.015)
Log income	0.022	-0.103^{***}	0.117^{***}	−0.009
	(0.022)	(0.026)	(0.028)	(0.022)
Age	-0.0072^{***}	0.0115^{***}	-0.0038^{***}	-0.0059^{***}
	(0.0007)	(0.0009)	(0.0009)	(0.0008)
Observations	12,025	12,025	12,025	12,025

Notes. Each column reports estimates from a separate fractional-logit regression of the form in equation (137). The dependent variable is the probability assigned to the return interval indicated by the column heading. Wealth is the sum of housing and retirement-account wealth and risky share is the equity share of retirement accounts (Section 2.1). HC1 heteroskedasticity-robust standard errors are reported in parentheses. Because fractional-logit coefficients are on the log-odds scale, their magnitudes should not be interpreted directly as percentage-point changes in probabilities. Table 2 therefore reports the corresponding average changes in predicted probabilities from doubling wealth and increasing the risky share by 10 percentage points. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Household-level scatter plots. Figures 5 and 6 display the household-level relationships underlying these regressions. For each of the four return regions, they plot the reported probability against the logarithm of retirement-account wealth, of housing wealth, and of their sum, together with a linear fit, first for the full sample and then excluding households with zero holdings of the respective asset. The fitted lines show the same pattern as the regression coefficients: the probability of a moderate gain rises with wealth, the probability of a moderate loss falls, and the two tail probabilities are nearly flat.

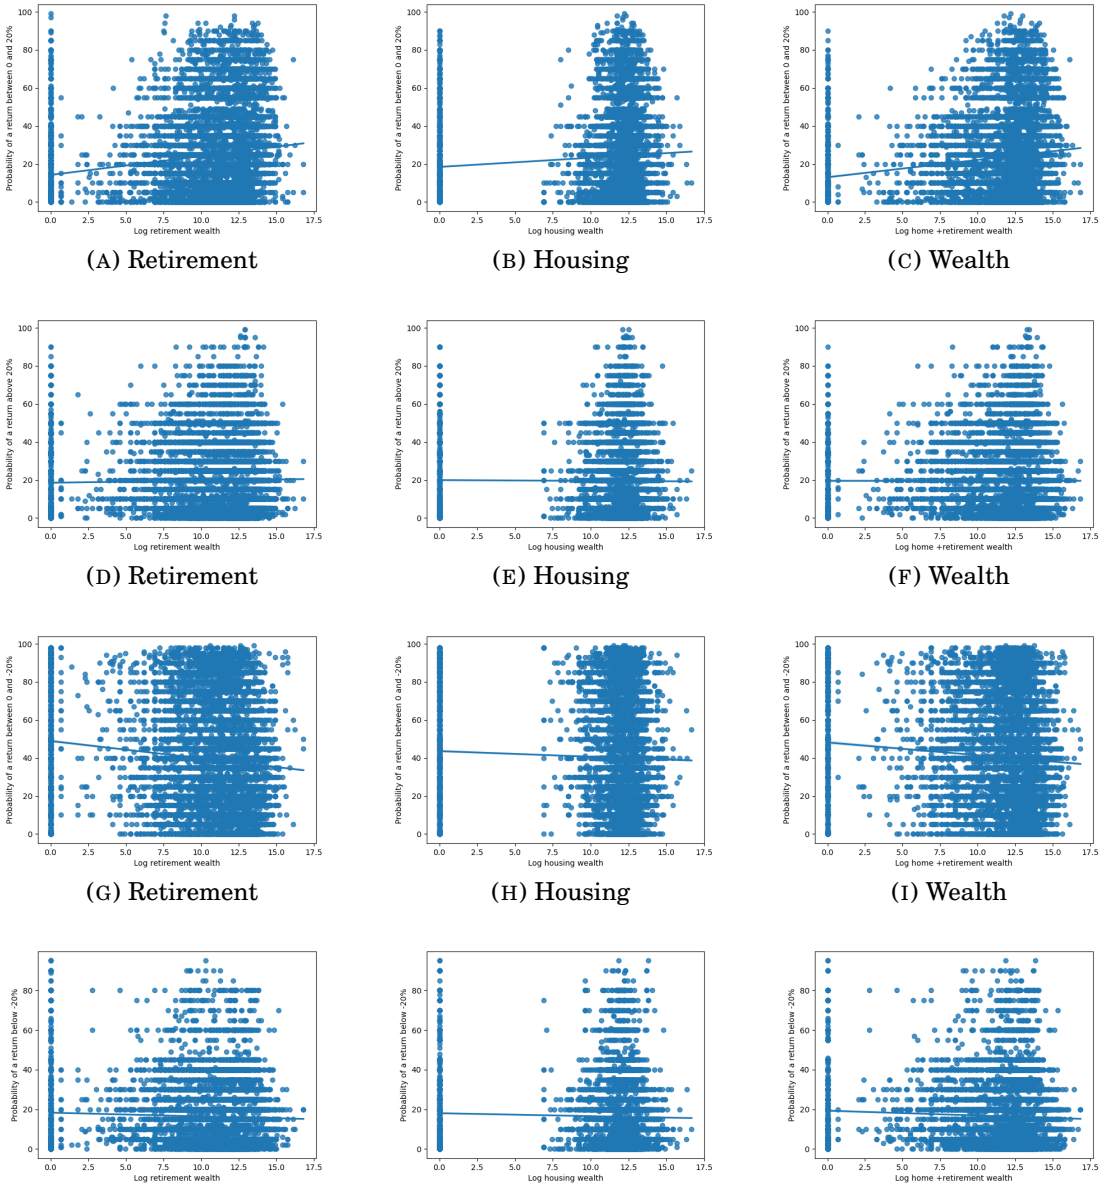

(A) Retirement (B) Housing (C) Wealth

(D) Retirement (E) Housing (F) Wealth

(G) Retirement (H) Housing (I) Wealth

(J) Retirement (K) Housing (L) Wealth

FIGURE 5. Asset holdings and stock-market return expectations, full sample. "Retirement" denotes the log value of the retirement account, "Housing" log housing value, and wealth the log of the sum of the two. Rows correspond, from top to bottom, to $0 < R < 20\%$, $R > 20\%$, $-20\% < R < 0$, and $R < -20\%$.

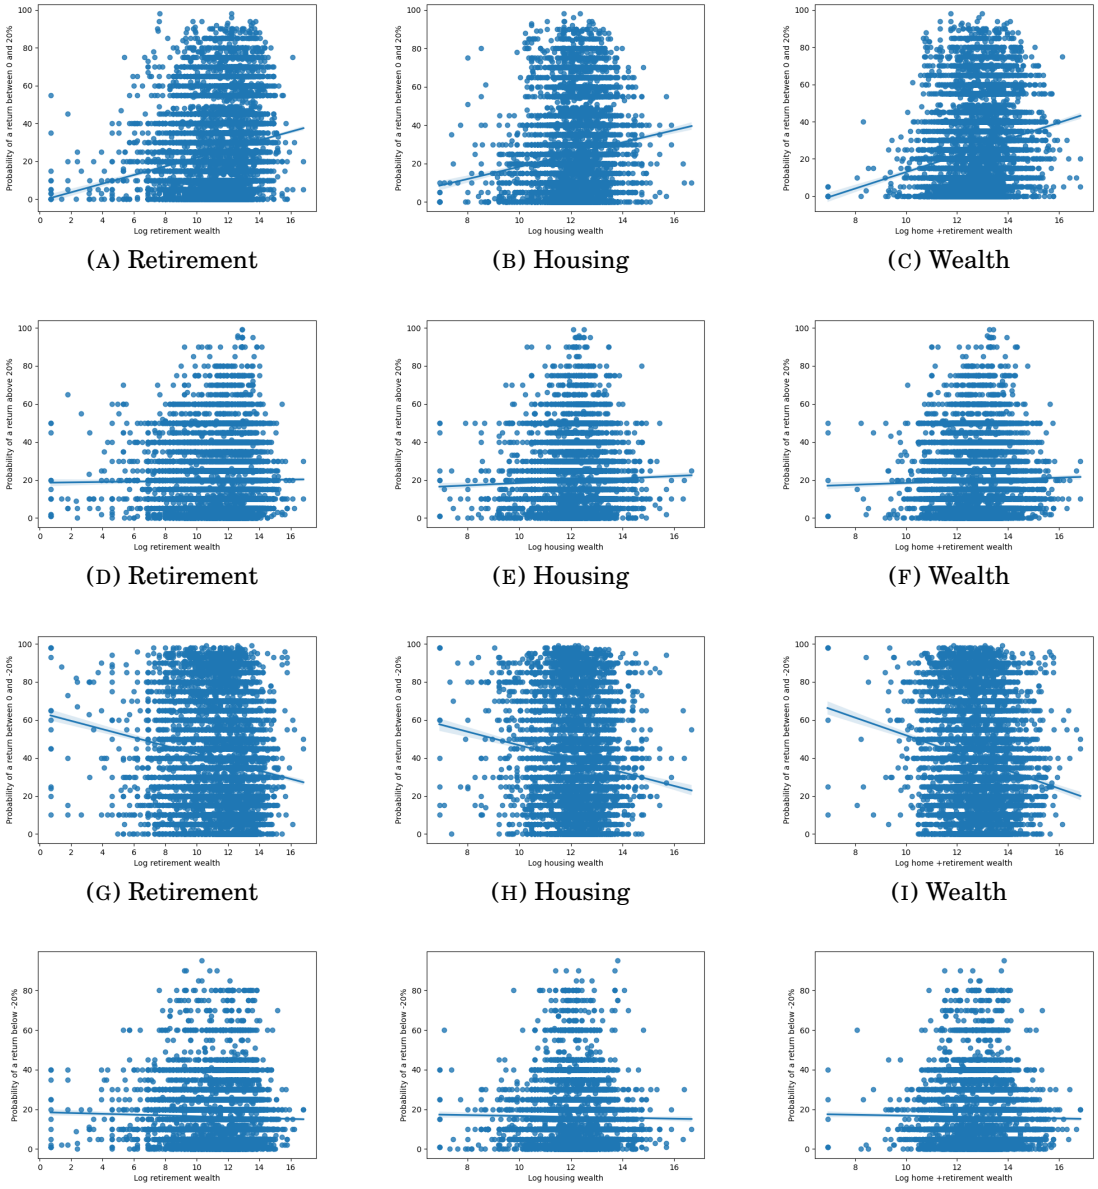

(A) Retirement (B) Housing (C) Wealth

(D) Retirement (E) Housing (F) Wealth

(G) Retirement (H) Housing (I) Wealth

(J) Retirement (K) Housing (L) Wealth

FIGURE 6. Asset holdings and stock-market return expectations, excluding observations with zero assets (intensive margin). "Retirement" denotes the log value of the retirement account, "Housing" log housing value, and wealth the log of the sum of the two. Rows correspond, from top to bottom, to $0 < R < 20\%$, $R > 20\%$, $-20\% < R < 0$, and $R < -20\%$.

D. Proofs for Section 4

Proof of Proposition 1. Fix $\Delta > 0$. Within the analytical lognormal approximation, the wealth recursion (66) has i.i.d. positive multipliers $M_t = \kappa R_t^A$ and a constant positive additive term $B = \kappa \bar{W}$. By (79), $\log M_t \sim \mathcal{N}(\lambda(\Delta), v(\Delta)^2 \sigma_\eta^2)$. The condition $\lambda(\Delta) < 0$ establishes existence of a unique stationary solution to the affine recursion. Since $v(\Delta) > 0$ and $\sigma_\eta > 0$, the multiplier has a nondegenerate lognormal law, $\Pr(M_t > 1) > 0$, and $\log M_t$ is nonarithmetic. Its positive moments are finite and the positive root ξ of $\mathbb{E}[M_t^\xi] = 1$ satisfies $\mathbb{E}[M_t^\xi \log^+ M_t] < \infty$. These are the standard Kesten–Goldie conditions for a Pareto upper tail with a positive finite coefficient, as in (77); the result applies to the approximating affine recursion.

Substituting $v(\Delta) = \Delta/(\gamma\sigma^2)$ from (59) and $\lambda(\Delta)$ from (69) into (82), and using $\sigma^2 = \sigma_\eta^2$, gives

$$\begin{aligned} \xi(\Delta) &= -\frac{2}{\gamma}\left[\log\beta + r_B + \frac{2\Delta^0\Delta - \Delta^2}{2\sigma^2}\right]\frac{\gamma^2\sigma^4}{\Delta^2\sigma_\eta^2} \\ &= \gamma\,\frac{-2\sigma^2(\log\beta + r_B) + \Delta^2 - 2\Delta^0\Delta}{\Delta^2}, \end{aligned} \tag{138}$$

which is (84) by the definition of K in (83). Differentiating,

$$\frac{\partial\xi(\Delta)}{\partial\Delta} = 2\gamma\frac{\Delta^0\Delta - K}{\Delta^3}, \tag{139}$$

which for $\Delta > 0$ is negative when $\Delta < K/\Delta^0$ and positive when $\Delta > K/\Delta^0$; this proves the second part with $\Delta^\star = K/\Delta^0$. For the third part, stationarity of the accurately informed household, (75), reads

$$\lambda(\Delta^0) = \frac{1}{\gamma}\left[\log\beta + r_B + \frac{(\Delta^0)^2}{2\sigma^2}\right] < 0, \tag{140}$$

which by (83) is equivalent to $K > (\Delta^0)^2$, so that $\Delta^\star = K/\Delta^0 > \Delta^0$. Evaluating (84) at $\Delta^\star$,

$$\xi(\Delta^\star) = \gamma\frac{K^2/(\Delta^0)^2 - 2K + K}{K^2/(\Delta^0)^2} = \gamma\left(1 - \frac{(\Delta^0)^2}{K}\right), \tag{141}$$

and the limit as $\Delta \to \infty$ follows directly from (84). Finally, a Pareto upper tail has finite mean if and only if its exponent exceeds one. Since the global minimum is at $\Delta^\star$, $\xi(\Delta) > 1$ for all $\Delta > 0$ is equivalent to $\gamma[1 - (\Delta^0)^2/K] > 1$, or to (88). ■

Proof of Proposition 2. Equation (70) implies

$$\lambda(\Delta) = \lambda(\Delta^0) - \frac{(\Delta - \Delta^0)^2}{2\gamma\sigma^2}. \tag{142}$$

Taking expectations under G_d and using $\mathbb{E}_{G_d}[\Delta] = \Delta^0$ yields (97).

For the second result, evaluate (84) at $\Delta^0 - h$ and $\Delta^0 + h$. Their difference is

$$\xi(\Delta^0 - h) - \xi(\Delta^0 + h) = 4\gamma\frac{\Delta^0 h\left[K - (\Delta^0)^2 + h^2\right]}{\left[(\Delta^0)^2 - h^2\right]^2}. \tag{143}$$

Stationarity of the accurately informed household implies $K > (\Delta^0)^2$, so this expression is strictly positive.

Finally, (139) shows that $\xi(\Delta)$ is strictly decreasing up to $\Delta^\star$ and strictly increasing thereafter. Since $\Delta^\star > \Delta^0$, the lower end of the symmetric support, $\Delta^0 - d$, always lies to the left of $\Delta^\star$. Therefore the minimum of $\xi(\Delta)$ over $[\Delta^0 - d, \Delta^0 + d]$ is attained at $\Delta^0 + d$ until the upper endpoint reaches $\Delta^\star$, and at $\Delta^\star$ thereafter. For the nested supports in (96), Equation (92) then gives (100) and (101). ■