Learning to Be Rich: How Expectations Amplify Wealth Inequality*

Pau Belda
Janko Heineken
Adrian Ifrim

October 2026
Latest version here

Households persistently disagree about how the stock market will perform. In surveys, stable individual differences account for most of the variation in expected returns and are only weakly related to observable household characteristics, yet average optimism rises steeply 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 aggregate stock returns from the returns on their own portfolios, so idiosyncratic investment histories induce persistent disagreement. Analytically, we show that disagreement lowers average long-run wealth growth while thickening the upper tail of the wealth distribution, as a small number of relatively optimistic households become exceptionally wealthy. Quantitatively, heterogeneous beliefs raise the top-decile wealth share by about 10 percentage points and help account for the rise in U.S. wealth concentration since the 1980s. Methodologically, we show that 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. Within this literature, a central open question is why returns rise with wealth. We propose that part of the answer lies in households' subjective expectations about stock market returns.

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 future stock returns 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 a quantitatively important contribution to the concentration of wealth at the top.

The theory is motivated by two facts about household return expectations.[3] First, consistent with Giglio et al. (2021), disagreement is predominantly persistent: in the RAND data, individual fixed effects account for 57 percent of the variation in expected returns, compared with 13 percent for time effects, and these fixed effects are only weakly related to observables, with wealth alone explaining less than 9 percent of their cross-sectional variation. 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. Wealth therefore ex-

[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] 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.

plains 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 are uncertain about how asset prices are determined and treat the return on their own portfolio as informative about future market performance. Because portfolios differ in composition and are exposed to idiosyncratic return shocks, households living through the same aggregate history receive different private signals and form different beliefs. The process is self-referential: beliefs determine portfolio exposure, exposure determines realized portfolio returns, and realized returns are the signals from which households subsequently learn. When learning is slow, small differences in composition or luck cumulate into increasingly divergent return histories, sustaining persistent disagreement about the same market. At the same time, favorable histories raise wealth, optimism, and risky exposure together, so average beliefs become ordered 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 0.35.[4]

Persistent disagreement thickens the upper tail of the wealth distribution even as it lowers average long-run wealth growth. We establish this analytically in a version of the model in which beliefs are frozen at their persistent component and asset prices are held fixed. Accurate beliefs maximize a household's expected log wealth growth, and the loss from belief errors is quadratic: pessimists hold too little risk and optimists too much, so a mean-preserving spread in beliefs lowers average wealth growth. The consequences for extreme wealth are asymmetric. Within a group of households holding the same belief, idiosyncratic return shocks generate a Pareto upper tail whose thickness depends on both the drift of wealth accumulation and the exposure to these multiplicative shocks. A small dose of optimism raises exposure to first order while reducing growth only to second order, so the thickest tail is generated neither by the most accurate households nor by the most optimistic ones, but by moderately optimistic households; beyond that point the growth loss from overexposure dominates and the tail thins again. Because the upper tail of the population is governed by the belief group with the heaviest conditional tail, a mean-preserving increase in disagreement around the true expected return thickens the social tail while lowering average growth. Disagreement thus raises wealth concentration through rare, exceptionally wealthy winners among moderately

[4] 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.

optimistic households.[5]

General equilibrium adds a feedback from beliefs to the experiences that shape them. Greater optimism raises the demand for stocks and hence their price, so optimistic households partly generate the capital gains that confirm their optimism. 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 2016 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 model performs well along the dimensions that motivate the theory. It generates a large equity premium, assigns 74 percent of financial wealth to the top decile and 40 percent to the top 1 percent (76 and 38 percent in the data), and reproduces the steep gradients in portfolio composition and beliefs: expected stock returns fall from 8.6 percent in the top 1 percent to 5.1 percent in the bottom half, compared with 7.9 and 4.8 percent in the data. The model also passes a more demanding test, one not targeted in the calibration. Applying the fixed-effect decomposition of Section 2 to simulated beliefs, individual effects explain 65 percent of the variation in expected returns, against 57 percent in RAND, and wealth explains 8.5 percent of the variation in these individual effects, against 8.9 percent in the data. The model thus generates the same combination of large within-wealth disagreement and steep wealth–belief gradient that the data display. To quantify the contribution of beliefs, we compare the baseline with an economy in which the learning gain is set to zero, so that all households keep a common belief about stock returns, holding preferences, labor-income risk, and aggregate risk fixed. The top-decile wealth share falls from 74 to 63 percent: heterogeneous beliefs raise top-decile wealth concentration by about 11 percentage points, an effect measured conditional on preference heterogeneity. The mechanism is not higher average risk-taking—every wealth group holds a larger risky share in the common-belief economy—but endogenous sorting. Learning aligns beliefs, portfolio exposure, realized returns, and wealth across households, and this covariance, compounded over time, concentrates wealth at the top.

[5]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.

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 makes the model both consistent with survey evidence and substantially easier to solve. We adopt Internal Rationality (Adam and Marcet, 2011): households maximize expected utility under a well-specified subjective probability measure over the variables they take as given, including asset prices, but need not know the equilibrium mapping from the cross-sectional distribution of households to prices. Prices therefore enter households' state vectors directly and are forecast with their subjective models, which is what allows realized portfolio returns to inform beliefs about future stock returns while households remain fully rational given those beliefs. 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 prices directly from their 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 prices as state variables 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 price 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.

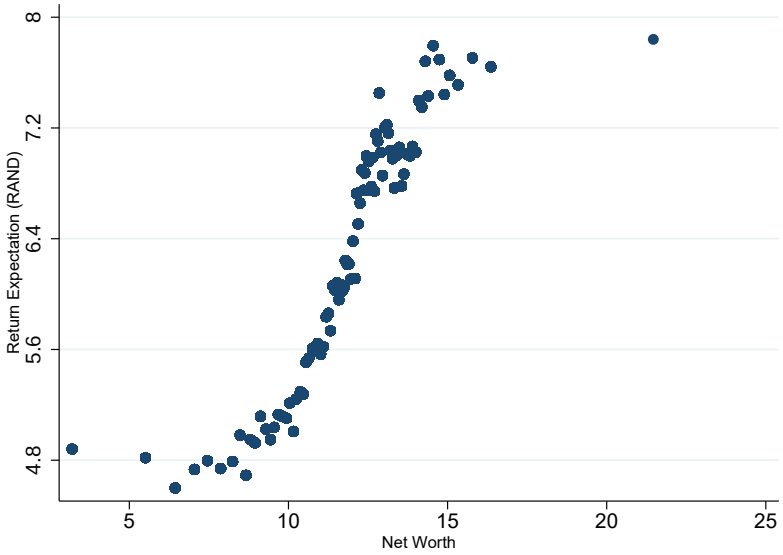
Panel B examines whether this optimism is associated with larger or smaller forecast errors. For each observation, we match the survey forecast to the realized one-year market return over the corresponding forecast horizon and define

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

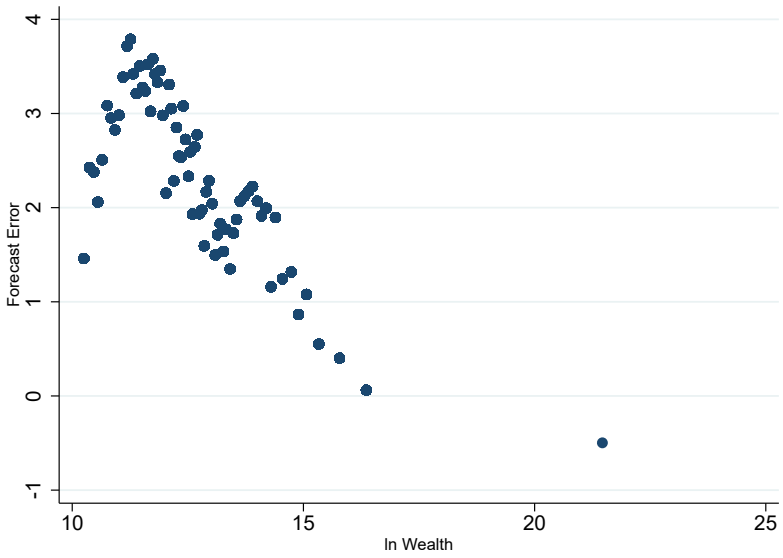
With this sign convention, a positive forecast error means that the household was too pessimistic and a negative forecast error means that it was too optimistic. Forecast errors decline with wealth. Poorer households therefore tend to be more pessimistic relative to the market outcomes realized over 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

$$A_t^i \geq 0, \qquad 0 \leq v_t^i \leq \bar{v}. \tag{14}$$

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

The constraint $A_t^i \geq 0$ rules out negative net worth. The portfolio bound rules out short positions in risky assets while allowing for leverage when $\bar{v} > 1$. These restrictions reflect the limited use of short selling among households and prevent extreme positions from driving the wealth distribution mechanically.

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.

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 broad 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 observed by all households. $R_{t+1}^{S,i}$ is the risky return actually experienced by household i. Hence two households exposed to the same aggregate stock market can experience different investment outcomes.

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.

The analytical model in Section 4 is a direct specialization of (19). There we write the aggregate component as

$$R^S_t = \zeta^0 \varepsilon^c_t,$$

so that, under the objective measure,

$$R^{S,i}_t = \zeta^0 \varepsilon^c_t \eta^i_t,$$

and impose lognormality on both the aggregate and idiosyncratic components.

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.[6]

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.

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 beliefs. 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*. Household i maximizes expected utility under a coherent subjective probability measure $\mathscr{P}^i$, but it need not know the true equilibrium law of motion for asset prices. Belief heterogeneity is therefore compatible with otherwise standard dynamic optimization.

Formally, household i's decision problem is defined on a probability space $(\Omega, \mathscr{B}, \mathscr{P}^i)$. A typical history contains the payoff-relevant variables external to the

[6]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.

household's choices,

$$\omega = \left\{D_t, W_t^i, P_t, Q_t, \eta_t^i\right\}_{t=0}^{\infty}. \tag{28}$$

The probability measure $\mathscr{P}^i$ is therefore a primitive needed to make the optimization problem in (12) well defined, in the same sense that preferences and constraints are primitives. Internal Rationality imposes standard optimization conditional on $\mathscr{P}^i$; it does not require $\mathscr{P}^i$ to coincide with the objective probability law generated in equilibrium. In particular, Euler equations and portfolio optimality conditions are evaluated using the household's own conditional distribution of future asset returns.

We discipline where imperfect knowledge enters. Households know the laws governing dividends, wages, and the idiosyncratic return shock η_t^i. They therefore agree about these sources of fundamental and idiosyncratic risk. The informational friction concerns the equilibrium process governing market prices and, consequently, the conditional distribution of the common stock-market return. This distinction is economically natural. Learning one's labor-income process or the dispersion of individual portfolio outcomes does not require knowing how the distribution of heterogeneous households maps into equilibrium asset prices. By contrast, survey evidence on stock-market expectations is difficult to reconcile with the assumption that ordinary households all possess the same correct model of that mapping.

Rational Expectations is nested as the special case in which the subjective conditional distributions implied by $\mathscr{P}^i$ coincide with the objective equilibrium conditional distributions. Thus, Internal Rationality does not replace optimization with an ad hoc behavioral rule. It relaxes a specific informational requirement: households need not know the equilibrium law of motion for prices. Section 5 gives the recursive formulation of this distinction and shows why it changes the state variables required to solve the household problem.

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. Over time, they observe common stock-market outcomes but also their own realized portfolio performance. Because personal portfolio returns depend on both endogenous portfolio composition and idiosyncratic risky returns, households exposed to the same aggregate market accumulate different histories of investment outcomes.

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 observing exactly the same aggregate market can receive different signals because they hold different portfolios and 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 the memory-retrieval process itself. Instead, it captures the economically relevant implication for wealth dynamics: household-specific investment experience affects beliefs about an aggregate risky return.

Perceived expected-return state. Household i perceives the conditional mean of the stock-market return as a latent, slowly moving state. Its subjective model contains a forecasting equation for next-period market returns,

$$\log R_{t+1}^S = Z_t^i + \varepsilon_{t+1}^{S,i}, \tag{29}$$

$$Z_t^i = Z_{t-1}^i + \varepsilon_t^{Z,i}. \tag{30}$$

Equation (29) is the household's *return forecasting equation*: conditional on Z_t^i, it specifies the subjective predictive distribution of the market return that will be realized at $t+1$. In particular,

$$\mathbb{E}_t^{\mathscr{P}^i}\left[\log R_{t+1}^S \mid Z_t^i\right] = Z_t^i.$$

It is not the measurement equation used to infer Z_t^i at date t, since R_{t+1}^S has not yet been observed. Information about Z_t^i instead arrives through the experience signal introduced below.

The purpose of this specification is not to impose a particular objective law of motion for expected stock returns. Objectively, the common stock return is generated in equilibrium by dividends and asset prices. Equations (29)–(30) instead describe the forecasting model household i uses under its subjective measure.

The random-walk specification captures the idea that investors do not view the expected return as an immutable parameter that can eventually be learned with arbitrary precision. They perceive the investment environment as evolving, so recent information continues to affect beliefs even after a long history of observations. This delivers persistent but revisable expectations rather than either dogmatic heterogeneous priors or eventual convergence to a common belief.

Portfolio experience as a signal. The central assumption is that households treat realized investment experience as informative about the latent state governing future market returns. At date t, household i constructs an experience signal from

the market return realized during period t and the return on its own portfolio. Under its subjective model, the household treats this statistic as a noisy contemporaneous observation of Z_t^i:

$$\mathscr{S}_t^i = Z_t^i + \nu_t^i. \tag{31}$$

Under $\mathscr{P}^i$, $\varepsilon_t^{Z,i}$ and ν_t^i are mean-zero Gaussian innovations with constant perceived variances σ_Z^2 and σ_ν^2, respectively.

Equation (31) is therefore the *measurement equation* of the household's filtering problem. In contrast to (29), which maps Z_t^i into a distribution for a future return, (31) maps an observation available at date t into information about the current latent state.

The observed signal is constructed as

$$\mathscr{S}_t^i = \mu_L \log R_t^{A,i} + (1-\mu_L)\log R_t^S, \qquad \mu_L \in [0,1]. \tag{32}$$

Thus, all realized-return information used for belief updating is summarized by the single statistic $\mathscr{S}_t^i$. In particular, the household does not treat the realized market return R_t^S as a second, independent measurement of Z_t^i in addition to (31); R_t^S enters the filtering problem only through its contribution to $\mathscr{S}_t^i$.

The parameter μ_L controls how much private investment experience matters for beliefs. If $\mu_L = 0$, the signal depends only on the common market return. Starting from a common prior, this channel therefore generates no cross-sectional disagreement. When $\mu_L > 0$, the signal becomes household-specific because

$$R_t^{A,i} = v_{t-1}^i R_t^S \eta_t^i + (1 - v_{t-1}^i) R_t^B. \tag{33}$$

Personal experience therefore differs for two conceptually distinct reasons. Households choose different exposures v_{t-1}^i, and households with the same exposure can receive different idiosyncratic risky returns η_t^i.

The latter feature is important. Without idiosyncratic risky returns, households holding identical portfolios and observing the same aggregate history would receive identical signals. With η_t^i, independent investment outcomes can initiate disagreement even from identical priors and identical initial portfolio positions. Portfolio choices subsequently make this disagreement endogenous and persistent.

The personal-experience component of (32) uses the return on the household's entire portfolio rather than the return on its risky holdings alone. This is a deliberate behavioral assumption. A household that holds most of its wealth in the safe asset experiences returns close to the safe rate and, under (32), reads that experience as evidence that stock returns are low, whereas a household with a large risky position experiences the stock market more directly. We adopt this specification for two reasons. First, it is the pattern in the survey evidence discussed below: in the UBS/Gallup data, expectations about the aggregate market respond to the return investors earned on their own portfolios, and Jiang et al. (2025) find that recalled own returns explain market expectations better than the market's actual returns do. Second, it is what makes portfolio histories, rather than types, the source of persistent household-specific belief components. Because exposure is itself a choice

that responds to beliefs, the signal carries each household's own history forward, and this is what allows the model to generate the individual fixed effects documented in Section 2.

Belief updating. Let

$$\zeta_t^i \equiv \mathbb{E}_{t-1}^{\mathscr{P}^i}[Z_t^i] \tag{34}$$

denote the prior belief about the latent expected-return state carried into period t. During period t, returns are realized and household i observes the experience signal $\mathscr{S}_t^i$. It then uses the measurement equation (31) to update its estimate of Z_t^i. The future market return R_{t+1}^S in (29) is not part of this date-t information set; rather, the updated estimate of Z_t^i determines the household's forecast of that return.

Under the perceived linear-Gaussian system, the steady-state Kalman filter therefore gives

$$\zeta_{t+1}^i = (1-g_L)\zeta_t^i + g_L\mathscr{S}_t^i = (1-g_L)\zeta_t^i + g_L\left[\mu_L \log R_t^{A,i} + (1-\mu_L)\log R_t^S\right], \qquad g_L \in (0,1). \tag{35}$$

Here $\zeta_{t+1}^i = \mathbb{E}_t^{\mathscr{P}^i}[Z_{t+1}^i]$: because the latent state follows a random walk, the posterior estimate of Z_t^i after observing $\mathscr{S}_t^i$ is also the prior mean of Z_{t+1}^i. All households start from a common prior,

$$\zeta_0^i = \zeta_0 \qquad \text{for all } i. \tag{36}$$

The gain g_L determines the persistence of beliefs. A small gain places more weight on the accumulated history embedded in ζ_t^i and therefore makes household differences highly persistent. A larger gain gives greater weight to the latest investment experience.

The learning rule does not imply that households literally believe every difference in personal performance is a permanent difference in aggregate expected returns. Equation (31) explicitly treats personal experience as a noisy measurement of the latent market state. Idiosyncratic return risk, portfolio composition, and other household-specific forces therefore create noise around the information relevant for forecasting aggregate returns. Households optimally filter this noisy signal according to their perceived model. The key departure from full information is that they regard their own investment experience as informative about the aggregate investment opportunity set.

The feedback from experience to wealth. Equations (23) and (35) create a feedback between beliefs, portfolios, and wealth.

Consider a household that receives an unusually high idiosyncratic risky return or happens to hold a portfolio that performs particularly well. Its realized portfolio return $R_t^{A,i}$ rises. When $\mu_L > 0$, this raises ζ_{t+1}^i and therefore its subjective expected stock return. Greater optimism increases its desired risky exposure. The household is consequently more exposed to risky returns in subsequent periods, making future investment outcomes more consequential for both its wealth and its beliefs. The same mechanism operates in the opposite direction after poor experiences.

There are therefore two forms of persistence. First, constant-gain learning makes the belief state itself persistent. Second, beliefs affect portfolio exposure, so past belief innovations alter the distribution of future signals. A favorable sequence of returns can make a household both wealthier and more optimistic, while an unfavorable sequence can make another household poorer and more pessimistic. This self-referential component is the mechanism we call *learning to be rich*.

Importantly, the mechanism does not require a permanent optimism type. Persistent differences in beliefs can emerge from common priors because households accumulate different investment histories. Nor does the mechanism require households to observe different aggregate histories. The aggregate stock-market return is common. What differs is how that common market is experienced through household portfolios.

Evidence on learning from own returns. The UBS–Gallup investor survey analyzed by Vissing-Jørgensen (2003) provides direct evidence on the key empirical link in equation (35). 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. The relationship is therefore not simply that aggregate market conditions move realized and expected returns together.

To examine this more directly, we estimate

$$\log \mathbb{E}_{it}\big(R^{m}_{t+1}\big) = \alpha_t + \gamma_m \log R^{A,i}_{t-1} + X'_{it}\lambda_m + u_{it}, \tag{37}$$

where $R^{A,i}_{t-1}$ is respondent i's past own-portfolio return, R^{m}_{t+1} 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 learning mechanism predicts

$$\gamma_m > 0 :$$

conditional on the common information available at a given date, investors who have experienced higher own-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:

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

TABLE 3. Own Portfolio Experience and Expected Market Returns

	(1)	(2)	(3)
$\log R_{t-1}^{A,i}$	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. $^{***}p < 0.01$.

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 μ_L. Realized portfolio returns are endogenous to previous portfolio choices and therefore to earlier beliefs. In the model, that endogeneity is not a nuisance to be eliminated: it is part of the mechanism. Beliefs determine risky exposure; risky exposure affects subsequent portfolio returns; and those returns become inputs into future beliefs.

The mechanism also provides a way to reconcile the two central empirical facts from Section 2. Individual fixed effects account for most of the panel variation in expected returns, while current wealth and standard observables explain only a small fraction of those fixed effects. At the same time, average optimism rises sharply with wealth. In the model, persistent individual heterogeneity reflects accumulated household-specific histories of portfolio returns rather than a fixed optimism parameter. Because those histories contain substantial idiosyncratic return variation, current wealth need not explain most of the cross-sectional dispersion in beliefs. Nevertheless, successful histories simultaneously raise wealth, optimism, and risky exposure, producing a systematic wealth–belief gradient in the cross-section.

The quantitative analysis asks whether this feedback is strong enough to jointly reproduce these two dimensions of the data and how much the resulting belief heterogeneity contributes to top wealth concentration.

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 of household allocations and portfolio choices

$$\left\{c_t^i, A_t^i, v_t^i, S_t^i, B_t^i\right\}_{i,t\geq 0},$$

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

1. *Household optimality.* Taking current prices 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.* Household belief states evolve according to (35), from the common prior in (36).
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 standard competitive equilibrium in allocations and prices: households are price takers, their choices are individually optimal, and markets clear. The only departure from the full-information rational-expectations concept is that equilibrium does not require each household's subjective law for future prices to coincide with the objective law generated by market clearing. The latter restriction is restored in the Rational Expectations special case. Section 5 recasts this equilibrium recursively and makes precise the associated distinction between the objective equilibrium state and the state variables perceived by households.

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. 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: a mean-preserving spread in beliefs can lower average long-run wealth growth while simultaneously thickening the upper tail of the wealth distribution.

4.1. Analytical benchmark

We simplify the environment in Section 3 in four 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, $\bar{W} > 0$.

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: common stock-market shocks generate different exposures across households, but cannot by themselves generate independent multiplicative return histories within a cross-section.

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

$$R_t^{S,i} = \zeta^i \varepsilon_t^{P,i}, \qquad \log \varepsilon_t^{P,i} \overset{i.i.d.}{\sim} \mathcal{N}\left(-\frac{1}{2}\sigma^2, \sigma^2\right). \tag{38}$$

Because $\mathbb{E}^{\mathscr{P}^i}[\varepsilon_t^{P,i}] = 1$, 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

$$\Delta^i \equiv \log \zeta^i - r_B. \tag{39}$$

The corresponding objective quantity is

$$\Delta^0 \equiv \log \zeta^0 - r_B, \qquad \Delta^0 > 0. \tag{40}$$

To distinguish aggregate from cross-sectional return risk, write the same risky return under the objective measure as

$$R_t^{S,i} = \zeta^0 \varepsilon_t^c \eta_t^i, \tag{41}$$

where

$$\log \varepsilon_t^c \sim \mathcal{N}\left(-\frac{1}{2}\sigma_c^2, \sigma_c^2\right), \qquad \log \eta_t^i \overset{i.i.d.}{\sim} \mathcal{N}\left(-\frac{1}{2}\sigma_\eta^2, \sigma_\eta^2\right), \tag{42}$$

with ε_t^c common to all households and η_t^i independent across households and over time. We assume the two innovations are independent and define

$$\sigma^2 \equiv \sigma_c^2 + \sigma_\eta^2. \tag{43}$$

Thus 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,

$$\frac{\bar{W}}{R_t^{A,i} A_{t-1}^i + \bar{W}} \simeq 0, \tag{44}$$

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

$$v^i \simeq \frac{\log\left(\mathbb{E}_t^{\mathscr{P}^i}[R_{t+1}^S]\right) - \log R^B}{\gamma\sigma^2} = \frac{\Delta^i}{\gamma\sigma^2}. \tag{45}$$

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

$$\kappa^i \simeq \left[\beta\mathbb{E}^{\mathscr{P}^i}\left[(R_{t+1}^{A,i})^{1-\gamma}\right]\right]^{1/\gamma}. \tag{46}$$

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

$$\log R_{t+1}^{A,i} \simeq r_B + v^i\Delta^i - \frac{1}{2}(v^i)^2\sigma^2 + v^i\sigma\varepsilon_{t+1}, \qquad \varepsilon_{t+1} \sim \mathcal{N}(0,1). \tag{47}$$

Hence, using the lognormal moment formula,

$$\log\mathbb{E}^{\mathscr{P}^i}\left[(R_{t+1}^{A,i})^{1-\gamma}\right] = (1-\gamma)\left[r_B + v^i\Delta^i - \frac{\gamma}{2}(v^i)^2\sigma^2\right] \tag{48}$$

$$= (1-\gamma)\left[r_B + \frac{(\Delta^i)^2}{2\gamma\sigma^2}\right], \tag{49}$$

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

$$\log\kappa^i \simeq \frac{1}{\gamma}\left[\log\beta + (1-\gamma)\left(r_B + \frac{(\Delta^i)^2}{2\gamma\sigma^2}\right)\right]. \tag{50}$$

For $\gamma > 1$,

$$\frac{\partial\log\kappa^i}{\partial\Delta^i} = -\frac{\gamma-1}{\gamma^2\sigma^2}\Delta^i < 0. \tag{51}$$

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.[7]

4.2. Subjective beliefs and typical wealth growth

Under the proportional policy rules, financial wealth evolves according to

$$A_t^i = \kappa^i R_t^{A,i} A_{t-1}^i + \kappa^i\bar{W}. \tag{52}$$

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

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

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,

$$\begin{aligned}\log R_t^{A,i} \simeq r_B + v^i\Delta^0 - \frac{1}{2}(v^i)^2\sigma^2 \\ + v^i\sigma_c\varepsilon_t + v^i\sigma_\eta\eta_t^i,\end{aligned} \tag{53}$$

where ε_t and η_t^i are standardized aggregate and idiosyncratic innovations. At high wealth, the additive wage component is negligible for growth. Objective expected log wealth growth is therefore

$$\begin{aligned}\lambda^i &\equiv \mathbb{E}^0\left[\log\left(\kappa^i R_t^{A,i}\right)\right] \\ &= \log\kappa^i + r_B + v^i\Delta^0 - \frac{1}{2}(v^i)^2\sigma^2.\end{aligned} \tag{54}$$

Substituting (45) and (50) gives

$$\lambda^i = \frac{1}{\gamma}\left[\log\beta + r_B + \frac{(\Delta^0)^2}{2\sigma^2} - \frac{(\Delta^i-\Delta^0)^2}{2\sigma^2}\right]. \tag{55}$$

Belief errors therefore reduce objective expected log wealth growth quadratically:

$$\lambda^i - \lambda^0 = -\frac{(\Delta^i-\Delta^0)^2}{2\gamma\sigma^2}, \qquad \lambda^0 \equiv \lambda(\Delta^0). \tag{56}$$

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

$$\begin{aligned}\frac{\partial\lambda^i}{\partial\Delta^i} &= \underbrace{\Delta^0\frac{\partial v^i}{\partial\Delta^i}}_{\substack{\text{return gain}\\(+)}} + \underbrace{\left(-\sigma^2 v^i\frac{\partial v^i}{\partial\Delta^i}\right)}_{\substack{\text{volatility drag}\\(-)}} + \underbrace{\frac{\partial\log\kappa^i}{\partial\Delta^i}}_{\substack{\text{lower saving}\\(-),\ \gamma>1}} \\ &= \frac{\Delta^0-\Delta^i}{\gamma\sigma^2}.\end{aligned} \tag{57}$$

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 (51): for $\gamma > 1$, greater perceived investment opportunities induce households to consume more of current resources.

It follows immediately that

$$\frac{\partial\lambda(\Delta)}{\partial\Delta}\begin{cases}>0, & \Delta<\Delta^0,\\ =0, & \Delta=\Delta^0,\\ <0, & \Delta>\Delta^0.\end{cases} \tag{58}$$

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

$$\lambda^i = \mathbb{E}^0\left[\log\left(\kappa^i R_t^{A,i}\right)\right] < 0. \tag{59}$$

Equivalently,

$$\underbrace{-\log\kappa^i}_{\text{consumption drag}} + \underbrace{\frac{1}{2}(v^i)^2\sigma^2}_{\text{volatility drag}} > \underbrace{r_B + v^i\Delta^0}_{\text{return component}}. \tag{60}$$

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 (59) 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,

$$\lambda(\Delta^0) < 0, \tag{61}$$

implies stationarity for every other belief:

$$\lambda(\Delta) \le \lambda(\Delta^0) < 0. \tag{62}$$

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. At a given date, all households have experienced the same sequence of aggregate stock-market shocks but different sequences of idiosyncratic return innovations. Let

$$\mathscr{A}_t \equiv \{\varepsilon_1, \ldots, \varepsilon_t\} \tag{63}$$

denote the realized aggregate return history.

Conditional on $\mathscr{A}_t$, the wealth distribution across households is a genuine cross-sectional object. Because $\mathscr{A}_t$ itself changes over time, the stationary cross-sectional wealth distribution is a random distribution: aggregate shocks shift the location and shape of the cross-section even after the economy has forgotten its initial conditions.

For a household with belief Δ, suppressing the dependence on $\mathscr{A}_t$ for notational

simplicity, write the stationary upper tail as

$$1 - F_t(a \mid \Delta) \sim C_t(\Delta) a^{-\xi(\Delta)}, \qquad a \to \infty, \tag{64}$$

where $F_t(a \mid \Delta) = \Pr(A_t \le a \mid \Delta)$. The scale $C_t(\Delta)$ is random because it depends on the aggregate return history. The Pareto coefficient $\xi(\Delta)$, by contrast, is constant almost surely.

Define the multiplicative component of wealth accumulation as

$$M_t^i \equiv \kappa^i R_t^{A,i}. \tag{65}$$

From (53),

$$\log M_t^i = \lambda(\Delta^i) + v^i \sigma_c \varepsilon_t + v^i \sigma_\eta \eta_t^i. \tag{66}$$

We first characterize the wealth distribution *within a given belief class*. Fix a household type with subjective expected excess return Δ. Households in this class choose the same risky share $v(\Delta)$ and are exposed to the same aggregate shock, but experience different idiosyncratic return innovations. Conditional on the realized aggregate shock,

$$\log \mathbb{E}_\eta \left[(M_t^i)^\xi \mid \varepsilon_t \right] = \xi [\lambda(\Delta) + v(\Delta)\sigma_c \varepsilon_t] + \frac{1}{2} \xi^2 v(\Delta)^2 \sigma_\eta^2. \tag{67}$$

Thus, for a fixed Δ, aggregate shocks move all households in the class through the same common component, whereas idiosyncratic shocks generate dispersion in their multiplicative wealth histories.

The Pareto coefficient is determined by the long-run Kesten condition along the realized aggregate history. Since the aggregate innovation is mean-zero and ergodic,

$$\frac{1}{T} \sum_{t=1}^{T} \varepsilon_t \longrightarrow 0 \qquad \text{almost surely.} \tag{68}$$

Hence the common component affects the time-varying scale of the wealth distribution within the class, but not its asymptotic Pareto coefficient. The relevant condition is therefore

$$\xi \lambda(\Delta) + \frac{1}{2} \xi^2 v(\Delta)^2 \sigma_\eta^2 = 0. \tag{69}$$

The positive solution is

$$\xi(\Delta) = -\frac{2\lambda(\Delta)}{v(\Delta)^2 \sigma_\eta^2}. \tag{70}$$

Equation (70) gives the Pareto coefficient of the wealth distribution conditional on belief type Δ. Two features are worth emphasizing. First, common aggregate risk does not enter the denominator directly: conditional on Δ, all households have the same exposure $v(\Delta)$ to the aggregate shock, so aggregate innovations do not generate within-class dispersion. Idiosyncratic return risk does. Second, beliefs affect the tail 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 (45) and the growth rate (55) into (70) gives the Pareto coefficient as a closed-form function of beliefs. Define

$$K \equiv -2\sigma^2 \left(\log\beta + r_B\right). \tag{71}$$

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

1. *For every* $\Delta > 0$*, the Pareto coefficient of the stationary wealth distribution of households with belief* Δ *is*

$$\xi(\Delta) = \gamma \frac{\sigma^2}{\sigma_\eta^2} \frac{\Delta^2 - 2\Delta^0\Delta + K}{\Delta^2}. \tag{72}$$

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

$$\Delta^\star \equiv \frac{K}{\Delta^0}. \tag{73}$$

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

$$\Delta^\star > \Delta^0. \tag{74}$$

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

$$\xi(\Delta^\star) = \gamma \frac{\sigma^2}{\sigma_\eta^2}\left(1 - \frac{(\Delta^0)^2}{K}\right) \qquad \textit{and} \qquad \lim_{\Delta\to\infty} \xi(\Delta) = \gamma\frac{\sigma^2}{\sigma_\eta^2}. \tag{75}$$

Proof. See Appendix D. ■

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

$$v(\Delta^0 + h) - v(\Delta^0) = \frac{h}{\gamma\sigma^2}. \tag{76}$$

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

$$\lambda(\Delta^0 + h) - \lambda(\Delta^0) = -\frac{h^2}{2\gamma\sigma^2}. \tag{77}$$

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

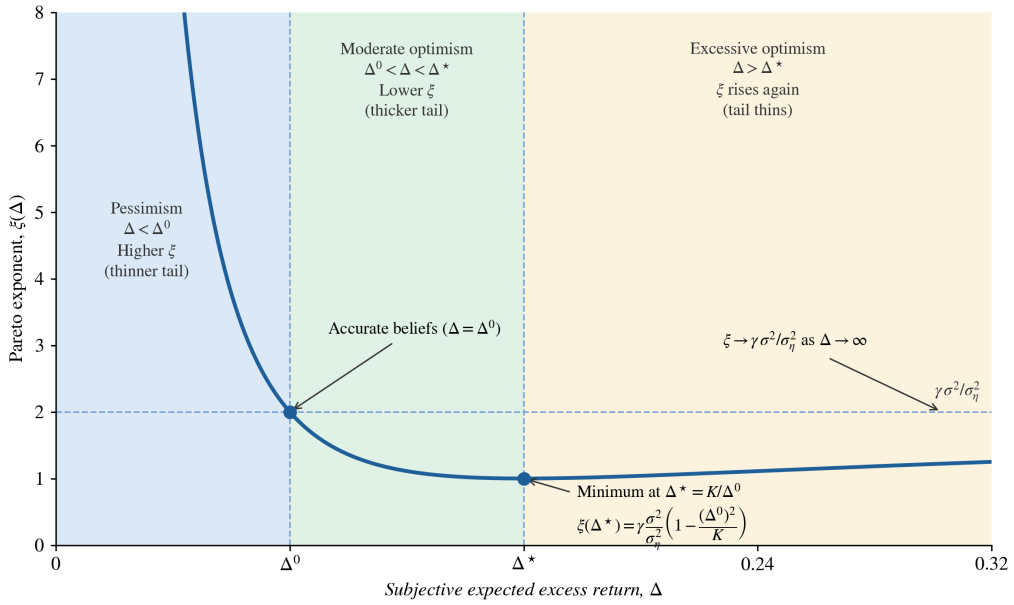


FIGURE 2. **Subjective expected excess return and the Pareto exponent.** The figure plots $\xi(\Delta)$ from Proposition 1 for an illustrative annual parameterization: $\gamma = 1.6$, $\sigma_c = 0.10$, $\sigma_\eta = 0.20$, $\Delta^0 = 0.08$, $r_B = 0.0075$, and β chosen so that $K = 2(\Delta^0)^2$ ($\beta = 0.873$). With these values $\xi(\Delta^0) = \gamma\sigma^2/\sigma_\eta^2 = 2$, $\Delta^\star = 2\Delta^0$, and $\xi(\Delta^\star) = 1$.

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 conditional wealth tail is generated neither by the accurately informed household nor by the most optimistic household, but by an interior degree of optimism. Figure 2 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

$$F_t^{XS}(a) = \int F_t(a \mid \Delta) dG_\Delta(\Delta). \tag{78}$$

Each belief class contributes a conditional 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,

$$\xi^{XS} = \inf_{\Delta \in \text{supp}(G_\Delta)} \xi(\Delta). \tag{79}$$

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

$$\Delta^\star \in \text{supp}(G_\Delta), \tag{80}$$

then

$$\xi^{XS} = \xi(\Delta^\star). \tag{81}$$

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:

$$\Delta = \Delta^0 + dz, \qquad z \sim H, \qquad \text{supp}(H) = [-1, 1], \tag{82}$$

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

$$\mathbb{E}_{G_d}[\Delta] = \Delta^0, \qquad \text{supp}(G_d) = [\Delta^0 - d, \Delta^0 + d]. \tag{83}$$

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 $\lambda(\Delta^0) < 0$. Consider the family of symmetric belief distributions in* (82)*. Then:*

1. *Average objective expected log wealth growth satisfies*

$$\mathbb{E}_{G_d}[\lambda(\Delta)] = \lambda(\Delta^0) - \frac{\text{Var}_{G_d}(\Delta)}{2\gamma\sigma^2} = \lambda(\Delta^0) - \frac{d^2\,\text{Var}(z)}{2\gamma\sigma^2}. \tag{84}$$

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

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

$$\begin{aligned} &\xi(\Delta^0 - h) - \xi(\Delta^0 + h) \\ &\quad = 4\gamma \frac{\sigma^2}{\sigma_\eta^2} \frac{\Delta^0 h \left[K - (\Delta^0)^2 + h^2\right]}{\left[(\Delta^0)^2 - h^2\right]^2} > 0. \end{aligned} \tag{85}$$

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

3. *Define*

$$d^\star \equiv \Delta^\star - \Delta^0 = \frac{K}{\Delta^0} - \Delta^0 > 0. \tag{86}$$

   The social Pareto coefficient is

$$\xi^{XS}(d) = \xi\left(\min\{\Delta^0 + d, \Delta^\star\}\right). \tag{87}$$

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

$d^\star < \Delta^0$, then

$$\xi^{XS}(d) = \begin{cases} \xi(\Delta^0 + d), & 0 \le d < d^\star, \\ \xi(\Delta^\star), & d^\star \le d < \Delta^0. \end{cases} \tag{88}$$

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.

Greater disagreement can therefore simultaneously reduce average long-run wealth growth and increase top wealth inequality. 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.

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. This is the relevant case in the quantitative model of Section 6: there, 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, so all households are pessimistic relative to the objective return and lie to the left of Δ^0, and a fortiori of $\Delta^\star$. The operative statement is that households that are less pessimistic than their peers take more risk and populate the upper tail. 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. 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 making explicit the distinction between the state of the economy and the state relevant for an individual household. Let Γ_t denote the cross-sectional distribution of predetermined household states, including financial wealth, inherited portfolio positions, labor-income states, preference types, and belief states. Aggregate dividends D_t are the aggregate exogenous state. Given household policy functions, asset-market clearing determines the two current asset prices, so there exists an equilibrium pricing function

$$(P_t, Q_t) = p(\Gamma_t, D_t). \tag{89}$$

For the individual problem, it is useful to write household i's beginning-of-period position as

$$\boldsymbol{x}_t^i = \left(W_t^i, A_{t-1}^i, v_{t-1}^i, D_t, P_t, Q_t, \zeta_t^i\right), \tag{90}$$

where fixed preference types are suppressed to lighten notation. The realized idiosyncratic risky-return innovation is part of the current household history and determines the realized portfolio return $R_t^{A,i}$ before the date-t saving and portfolio choices are made. Equivalently, it may be appended as a transitory component of $\boldsymbol{x}_t^i$ without changing any of the arguments below. The budget constraint is

$$c_t^i + A_t^i = W_t^i + \left[v_{t-1}^i R_t^{S,i} + (1 - v_{t-1}^i) R_t^B\right] A_{t-1}^i, \tag{91}$$

with $R_t^{S,i} = R_t^S \eta_t^i$ and $R_{t+1}^B = 1/Q_t$.

For given household policy functions, equilibrium prices solve the two asset-market clearing conditions,

$$\begin{aligned} \sum_i \mu^i h^s\left(W_t^i, A_{t-1}^i, v_{t-1}^i, D_t, P_t, Q_t, \zeta_t^i\right) &= 1, \\ \sum_i \mu^i h^b\left(W_t^i, A_{t-1}^i, v_{t-1}^i, D_t, P_t, Q_t, \zeta_t^i\right) &= \bar{B}_t, \end{aligned} \tag{92}$$

which gives the pricing function in (89).

State variables under Rational Expectations and Internal Rationality. The distinction between Rational Expectations and Internal Rationality is most transparent in recursive form. Under full-information rational expectations, households know the equilibrium pricing map p. Forecasting future prices therefore requires forecasting the future cross-sectional distribution. Suppressing fixed preference types, the relevant perceived state can be represented as

$$(\boldsymbol{x}_t^i)^{RE} = \left(W_t^i, A_{t-1}^i, v_{t-1}^i, D_t, \Gamma_t\right), \tag{93}$$

with current prices obtained from $p(\Gamma_t, D_t)$. The household problem consequently requires a perceived law of motion for the high-dimensional object Γ_t. This is the standard distribution-forecasting problem in heterogeneous-agent models with aggregate risk.

Under Internal Rationality, households do not use the equilibrium pricing map to forecast future prices. They observe current prices and use their subjective probability measures $\mathscr{P}^i$ to forecast prices and returns directly. The relevant state is therefore

$$(\boldsymbol{x}_t^i)^{IR} = \left(W_t^i, A_{t-1}^i, v_{t-1}^i, D_t, P_t, Q_t, \zeta_t^i\right). \tag{94}$$

The distribution Γ_t still determines actual equilibrium prices through (89), but it does not enter the household's optimization problem. Households need neither to observe the distribution nor to forecast its law of motion.

Recursive competitive equilibrium. We can now state the recursive counterpart of the competitive equilibrium defined in Section 3.4. Given the collection of subjective probability measures $\{\mathscr{P}^i\}_i$, the objective laws of motion for exogenous shocks, and an initial distribution Γ_0, a *recursive competitive equilibrium under Internal Rationality* consists of household policy functions

$$\left(h_c^i, h_A^i, h_v^i\right)\left((\boldsymbol{x}_t^i)^{IR}\right), \tag{95}$$

an equilibrium pricing function

$$(P_t, Q_t) = p(\Gamma_t, D_t), \tag{96}$$

and the induced transition law for the cross-sectional distribution, such that:

1. *Household optimality.* For every household state, the policy functions solve the household problem in Section 3 under $\mathscr{P}^i$, taking observed current prices as given.
2. *Market clearing.* For every (Γ_t, D_t), the prices $p(\Gamma_t, D_t)$ make the stock and bond market-clearing conditions in (92) hold.
3. *Distributional consistency.* The transition from Γ_t to Γ_{t+1} is the one induced by the household policy functions, the objective realizations of wages, dividends, and idiosyncratic risky returns, and the belief-updating rule in (35).

The equilibrium consistency requirement is therefore standard on the *objective* side of the economy: household choices aggregate into market-clearing prices and those choices and shocks generate the next cross-sectional distribution. What Internal Rationality removes is the additional requirement that the subjective forecasting model used by each household reproduce the equilibrium law $p(\Gamma_t, D_t)$. Households optimize under their subjective measures; they do not need to know the objective pricing function that their joint choices generate.

Two-stage decomposition. This recursive equilibrium immediately yields the computational separation used below. Conditional on the subjective probability mea-

sures, household policies can first be computed without imposing market clearing:

$$\mathscr{P}^i \quad \Longrightarrow \quad h^i\big((\boldsymbol{x}_t^i)^{IR}\big). \tag{97}$$

Given these policy functions and a cross-sectional distribution, equilibrium prices can then be obtained by aggregating individual asset demands and imposing market clearing:

$$\Big\{h^i\big((\boldsymbol{x}_t^i)^{IR}\big)\Big\}_i, \Gamma_t, D_t \quad \Longrightarrow \quad (P_t, Q_t) = p(\Gamma_t, D_t). \tag{98}$$

The second step recovers the equilibrium pricing function over the region of the state space visited by the economy. It does not require feeding that function back into the household forecasting problem. 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 price process. 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 policy functions jointly with a neural network,

$$\Big(\kappa_t^i, v_t^i\Big) = \mathscr{N}_\theta\Big((\boldsymbol{x}_t^i)^{IR}\Big), \tag{99}$$

where θ collects the network parameters. Given current resources, the saving policy determines A_t^i and therefore consumption through the budget constraint, while v_t^i determines the allocation of savings between stocks and bonds.

Economic theory provides a useful starting point for the approximation. We initialize the network close to the analytical Merton–Samuelson solution. In the benchmark problem with constant investment opportunities, the optimal risky share is characterized by the usual trade-off between the subjective expected excess return and risk, while the optimal saving rate has an analytical consumption–saving counterpart. The neural network therefore does not have to learn the policy functions from an arbitrary initialization. Instead, training learns the nonlinear correction to this benchmark generated by income risk, time-varying beliefs, endogenous prices, and occasionally binding portfolio constraints.

Training states are generated by simulating household histories under their subjective probability measures $\mathscr{P}^i$. In particular, we draw sequences for wages, dividends, belief states, and asset prices according to the processes perceived by the household and propagate the individual state using the candidate policy functions. Market clearing is *not* imposed at this stage. This is precisely where Internal Rationality delivers the computational separation: evaluating household optimality only requires the stochastic processes that households perceive, not the cross-sectional distribution that will ultimately generate equilibrium prices.

Simulation alone tends to concentrate training observations in the region of the state space visited under the current policy. This can provide too little information close to occasionally binding constraints. We therefore augment the simulated histories with off-path collocation states, with additional mass close to the saving constraint. These observations improve the approximation of the policy functions in regions that may be visited in equilibrium following sufficiently large shocks, even if

they occur infrequently in the subjective simulations used for training.

The network is trained directly on the household optimality conditions. Let $\mathscr{E}_A(\boldsymbol{x};\theta)$ denote the Euler residual associated with the consumption–saving choice and let $\mathscr{E}_v(\boldsymbol{x};\theta)$ denote the corresponding portfolio optimality residual. Because savings and portfolio positions are subject to

$$A_t^i \geq 0, \qquad 0 \leq v_t^i \leq \bar{v}, \tag{100}$$

the first-order conditions take the form of Euler–KKT conditions rather than equalities everywhere.

We impose these conditions using the Fischer–Burmeister function

$$\phi(a,b) = \sqrt{a^2+b^2} - a - b, \tag{101}$$

which satisfies $\phi(a,b) = 0$ if and only if $a \geq 0$, $b \geq 0$, and $ab = 0$. Applying this transformation to the relevant Euler gaps and constraint slacks converts the complementarity conditions into differentiable residuals that can be minimized with standard gradient-based methods. Schematically, the training objective is

$$\mathscr{L}(\theta) = \mathbb{E}_{\boldsymbol{x}\sim\mathscr{D}}\left[\Phi_A(\boldsymbol{x};\theta)^2 + \Phi_{v,0}(\boldsymbol{x};\theta)^2 + \Phi_{v,\bar{v}}(\boldsymbol{x};\theta)^2 + \lambda\,\Pi(\boldsymbol{x};\theta)\right], \tag{102}$$

where $\mathscr{D}$ combines simulated and collocation states, Φ_A, $\Phi_{v,0}$ and $\Phi_{v,\bar{v}}$ are the Fischer–Burmeister residuals associated with the saving and portfolio constraints, and Π penalizes policies that imply infeasible current or future resources.

At the end of Stage 1, we therefore obtain household policy functions that satisfy optimality under the agents' subjective probability measures. No equilibrium price function and no perceived law of motion for the cross-sectional distribution has been 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 asset prices.

We simulate a large population of households, assigning each household its type and belief state, and drawing the corresponding idiosyncratic innovations. At the beginning of date t, the collection of individual predetermined states and belief states defines the cross-sectional distribution Γ_t. For any candidate pair of prices (P,Q), the Stage-1 policy network generates the saving and portfolio choices of every household and therefore their demands for stocks and bonds.

Define the excess-demand vector

$$\mathscr{M}(P,Q;\Gamma_t,D_t) = \begin{pmatrix} \sum_i \mu^i h^s\left(W_t^i, A_{t-1}^i, v_{t-1}^i, D_t, P, Q, \zeta_t^i\right) - 1 \\ \sum_i \mu^i h^b\left(W_t^i, A_{t-1}^i, v_{t-1}^i, D_t, P, Q, \zeta_t^i\right) - \bar{B}_t \end{pmatrix}. \tag{103}$$

Equilibrium prices are the solution to

$$\mathcal{M}(P_t, Q_t; \Gamma_t, D_t) = 0. \tag{104}$$

We solve this two-dimensional system at every date of the aggregate simulation.

Once (P_t, Q_t) have been obtained, the trained network determines consumption, saving, and portfolio allocations. Household wealth and belief states are then updated using the realized returns and exogenous shocks, producing the next-period distribution Γ_{t+1}. Repeating these steps generates the general-equilibrium path:

$$\Gamma_t \longrightarrow (P_t, Q_t) \longrightarrow \{c_t^i, A_t^i, v_t^i\}_i \longrightarrow \Gamma_{t+1}. \tag{105}$$

The sequence of market-clearing solutions provides evaluations of the equilibrium pricing function

$$(P_t, Q_t) = p(\Gamma_t, D_t) \tag{106}$$

over the region of the state space visited by the economy. Crucially, households never need to know this function. It is an equilibrium object recovered from their optimal demands rather than an object they must learn or forecast in order to make those demands.

The distinction from the standard rational-expectations approach is therefore more than a reduction in the number of state variables fed into a numerical approximation. Under full-information rational expectations, the household problem and the equilibrium forecasting rule have to be solved jointly: a conjectured law for aggregate states determines household choices, those choices generate a new law of motion for the distribution, and the two must be made mutually consistent. Under Internal Rationality, the economic model itself separates these tasks. Stage 1 enforces individual optimality under subjective beliefs; Stage 2 enforces general-equilibrium consistency through market clearing. Deep learning provides a flexible way of solving the nonlinear household problem, while Internal Rationality is what makes the two-stage decomposition possible.

6. Quantitative Analysis

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 much heterogeneous expectations contribute 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. We convert these probabilities into expected returns using equation (5), following von Gaudecker and Wogrolly (2022). Second, 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

$$Y_{it} = \alpha + Z_{it}'\beta + \varepsilon_{it}, \qquad Y_{it} \in \left\{ \mathbb{E}_{it}^{S}[R_{t+1}], R_{t+1} - \mathbb{E}_{it}^{S}[R_{t+1}] \right\}. \tag{107}$$

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

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 (107) 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. *** $p < 0.01$.

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 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, private businesses, or other illiquid 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,

$$g_{t+1} \equiv \frac{D_{t+1}}{D_t} = \bar{g}_D \exp\left(-\frac{1}{2}\sigma_D^2 + \sigma_D \varepsilon_{t+1}^D\right), \qquad \varepsilon_{t+1}^D \sim \mathcal{N}(0,1). \tag{108}$$

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

$$R_{t+1}^S = \frac{P_{t+1} + D_{t+1}}{P_t}, \qquad R_{t+1}^B = \frac{1}{Q_t}. \tag{109}$$

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,

$$\tau_t^i \mapsto \left(\varepsilon(\tau_t^i), \beta(\tau_t^i), \gamma(\tau_t^i)\right). \tag{110}$$

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, which makes the quantitative exercise conservative with respect to expectations: 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),

$$P_W^A = \begin{pmatrix} 0.9972 & 0.0028 \\ 0.0252 & 0.9748 \end{pmatrix}, \qquad P_\gamma^A = \begin{pmatrix} 0.993 & 0.007 \\ 0.001 & 0.999 \end{pmatrix}, \tag{111}$$

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,

$$P_W^Q \simeq \begin{pmatrix} 0.999293 & 0.000707 \\ 0.006367 & 0.993633 \end{pmatrix}, \qquad P_\gamma^Q \simeq \begin{pmatrix} 0.998245 & 0.001755 \\ 0.000251 & 0.999749 \end{pmatrix}. \tag{112}$$

These matrices define each type's quarterly exit hazard,

$$m(\tau) = \phi \left[1 - \pi^W(\tau)\pi^\gamma(\tau) \right], \qquad \phi = 0.55, \tag{113}$$

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,

$$\begin{aligned} &\Pi(\tau, \tau') = m(\tau)\, q(\tau') \quad \text{for } \tau' \neq \tau, \qquad \Pi(\tau, \tau) = 1 - m(\tau) + m(\tau)\, q(\tau), \\ &\quad q(\tau') \propto \pi(\tau')\, m(\tau'), \end{aligned} \tag{114}$$

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

TABLE 5. **Household Types**

Type	Share (%)	Wage / mean	β (quarterly)	γ	Exit prob. (%, quarterly)
1	39.1	0.52	0.952	13.00	0.05
2	12.2	2.52	0.973	2.60	0.45
3	22.0	1.05	0.956	5.50	0.05
4	16.7	0.37	0.951	9.50	0.05
5	9.0	0.73	0.954	7.50	0.05
6	1.0	12.90	0.975	1.45	0.45

Notes. Each type determines a household's wage level, discount factor, and relative risk aversion. Wages are expressed relative to the population mean; the mean quarterly wage is 5.24 times the quarterly aggregate dividend. The last column reports the quarterly probability of leaving the type implied by equation (113). Beliefs are not assigned by type and evolve separately through learning from portfolio experience.

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,

$$\left(\tau_t^i, A_{t-1}^i, v_{t-1}^i, \zeta_t^i\right). \tag{115}$$

Households perceive quarterly log stock returns as normally distributed around their belief with standard deviation $\sigma = 0.10$, or 20 percent per year, and perceive the safe return as centered on the historical mean real safe rate of 0.75 percent per year. The common prior ζ_0 is the population-weighted mean of the expected-return schedule in the RAND–SCF data, 5.7 percent per year, and beliefs are restricted numerically to a support between 3.2 and 9.1 percent per year, which contains every value in that schedule. The gain g_L is the only free structural learning parameter; it is chosen so that the stationary expected-return schedule by wealth group is close to the data.

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,

TABLE 6. **Quantitative Targets**

Moment	Data	Moment	Data
Capital-income share	27.0%	Wealth–income ratio	491%
Mean risky share	30.5%	Wealth-weighted risky share	70.8%
Saving rate	6.7%	Mean real stock return	9.29%
Mean real safe rate	0.75%		

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.

capital-income shares, risky shares, expected stock returns, and portfolio returns—to assess the joint cross-sectional fit.

Two features of this strategy are worth emphasizing. First, all portfolio choices, asset prices, and the income and wealth distributions are endogenous. Second, forecast errors are deliberately excluded from the set of targeted moments.

6.3. Quantitative Performance

We evaluate the stationary economy along three dimensions: aggregate quantities, cross-sectional distributions, and the dynamics of beliefs. Some of the corresponding moments enter the calibration and should therefore be interpreted as measures of joint fit rather than out-of-sample validation. The belief-persistence moments are more demanding. In particular, the model is not directly calibrated to reproduce the fraction of belief variation accounted for by individual fixed effects or the small explanatory power of wealth for those fixed effects.

6.3.1. Aggregate Moments

Table 8 compares aggregate model outcomes with their empirical counterparts. The model reproduces the broad composition of household portfolios. The average risky share is 25.35 percent, compared with 30.50 percent in the data, while risky assets account for 69.40 percent of aggregate wealth, almost exactly the empirical 70.80 percent. This difference between the unweighted and wealth-weighted risky shares is itself informative: wealthy households hold much larger risky positions than the average household, a feature that the model reproduces endogenously.

The model also comes close to the aggregate saving rate and to the moments of dividend growth. Its principal aggregate shortcoming is the level of wealth relative to income. The model produces a wealth–income ratio of 236 percent, compared with 491 percent in the data, and consequently understates the aggregate capital-income share. It also generates somewhat high average asset returns: the mean stock return is 12.77 percent rather than 9.29 percent, and the mean safe rate is 1.22 percent rather than 0.75 percent. We view these discrepancies as useful discipline. The model's success in generating wealth concentration below does not come from matching the aggregate stock of wealth or from assuming an unusually low safe return.

TABLE 7. **Parameter Values**

Parameter	Description	Value
Panel A. Types, income, and assets (quarterly)		
$\beta(\tau), \gamma(\tau), \varepsilon(\tau)$	Discount factor, risk aversion, and wage by type, with population shares $\pi(\tau)$	Table 5
$\bar{w}$	Mean labor income relative to the aggregate dividend	5.24
$\bar{g}_D$	Mean gross dividend growth (2.84 percent per year)	1.00703
σ_D	Volatility of log dividend growth (6.4 percent per year)	0.0311
$\bar{S}$	Supply of the risky asset	1
$\bar{b}$	Supply of the safe asset relative to the dividend	19.0
$\bar{v}$	Upper bound on the risky share	1.10
ϕ	Scale of type exit hazards, equation (113)	0.55
Panel B. Beliefs and learning		
σ	Perceived volatility of quarterly log stock returns (20 percent per year)	0.10
R^B	Perceived mean gross safe return (0.75 percent per year)	1.00187
ζ_0	Common prior, gross quarterly expected stock return (5.7 percent per year)	1.0139
μ_L	Weight on own portfolio return in the signal	0.50
g_L	Constant learning gain	0.01
$[\zeta_{\min}, \zeta_{\max}]$	Numerical support for beliefs (3.2 to 9.1 percent per year)	[1.008, 1.022]
Panel C. Solution and simulation		
N	Simulated households	50,000
T	Simulated quarters (of which burn-in)	20,000 (5,000)
	Policy network	3 hidden layers of 96 units
	Quadrature nodes, stock return × dividend growth	32×3
	Pseudo-RAND panels for Table 11	$200 \times 2{,}000 \times 36$

Notes. The model is quarterly; annual equivalents are given in parentheses. The policy network maps the household state to the saving rate and risky share with SiLU activations and is initialized at the Merton–Samuelson solution. The pseudo-RAND panels follow 2,000 randomly sampled households for 36 quarters in each of 200 non-overlapping windows of the stationary simulation.

TABLE 8. **Aggregate Quantitative Performance**

Moment	Model	Data
Household aggregates		
Capital-income share (%)	19.78	27.00
Wage-income share (%)	80.22	73.00
Wealth–income ratio (%)	236.42	491.00
Average risky asset share (%)	25.35	30.50
Risky assets / aggregate wealth (%)	69.40	70.80
Saving rate (%)	5.53	6.70
Aggregate returns, annual (%)		
Mean stock return	12.77	9.29
Mean safe rate	1.22	0.75
Mean dividend growth	2.77	2.84
SD dividend growth	6.32	6.40

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 produces substantial wealth concentration. The top 1 percent holds 40.22 percent of financial wealth, compared with 37.60 percent in the SCF. The next 9 percent holds 33.94 percent, compared with 38.10 percent. Hence the model assigns 74.16 percent of wealth to the top decile, close to the empirical 75.70 percent, although it places somewhat too much of this wealth inside the top 1 percent. The main discrepancy appears further down the distribution: the bottom half holds 8.65 percent of model wealth but only 1.40 percent in the data. The model therefore captures concentration at the top considerably better than the very thin wealth holdings at the bottom.

The income distribution is considerably less concentrated than wealth in both the model and the data. Wage shares are also close to their empirical counterparts. This distinction matters for the mechanism. Extreme wealth concentration is not being generated by assigning an equally extreme share of labor income to wealthy households. The model instead transforms more moderate earnings inequality into much greater wealth inequality through saving, portfolio choice, realized returns, and belief dynamics.

The model also reproduces the strong portfolio gradient. The risky share falls from 89.32 percent for the top 1 percent to 13.97 percent for the bottom half, compared with 77.30 and 18.10 percent, respectively, in the data. The model slightly overstates risky exposure at the top and understates it in the middle and bottom, but captures the sharp monotonic relationship between wealth and financial risk-taking.

Most importantly for the paper's mechanism, the model jointly generates the wealth gradient in beliefs and the wealth gradient in portfolio returns. Expected stock returns decline from 8.60 percent for the top 1 percent to 5.09 percent for

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	40.22	33.94	17.18	8.65	37.60	38.10	22.80	1.40
Income share	18.35	25.86	33.57	22.22	18.40	24.70	35.20	21.60
Wage share	9.58	23.80	39.55	27.07	12.30	21.60	38.50	27.50
Risky asset share	89.32	70.11	33.53	13.97	77.30	62.30	39.60	18.10
Expected stock return	8.60	8.14	6.72	5.09	7.88	6.78	5.39	4.77
Portfolio return	12.09	8.97	4.72	2.53	8.30	6.12	3.63	3.40

Notes. Wealth, income, and wage entries are shares of the corresponding aggregate. Risky asset shares are averages within wealth groups. Expected stock returns and realized portfolio returns are annual percentages. The empirical wealth concept is financial wealth, excluding housing and other nonfinancial assets.

the bottom half, close to the empirical decline from 7.88 to 4.77 percent. Average portfolio returns fall even more sharply, from 12.09 to 2.53 percent. The model therefore generates the empirical ordering

$$\text{wealth} \longleftrightarrow \text{optimism} \longleftrightarrow \text{risky exposure} \longleftrightarrow \text{portfolio returns},$$

rather than matching each of these objects in isolation.

Two features of Table 9 connect these results to Section 4. First, every group's expected stock return lies below the model's mean stock return of 12.77 percent (Table 8), so households are pessimistic relative to the objective return and lie on the decreasing branch of the Pareto coefficient in Proposition 1; the comparative static at work is that less pessimistic households hold more risky assets and populate the upper tail. Second, the wealthiest households are those that are least pessimistic: the upper tail is populated by the relatively optimistic side of the belief distribution, as Proposition 2 implies.

The contribution of heterogeneous expectations. To measure the distributional role of belief heterogeneity, we compare the baseline economy with a counterfactual in which the weight on own portfolio returns is set to zero, $\mu_L = 0$. Households continue to update their beliefs, but only in response to the common market return. Because they start from the same prior and receive the same signal, they hold identical expected stock returns at every date. The type process for preferences and labor income, aggregate risk, and the solution method are otherwise unchanged. The comparison therefore removes disagreement generated by heterogeneous portfolio experiences without eliminating heterogeneity in patience, risk aversion, or labor income.

Table 10 shows that disagreement has a large effect on top concentration. Under heterogeneous beliefs, the top 1 percent holds 40.22 percent of wealth and the next 9 percent holds 33.94 percent. Under common beliefs, the corresponding shares are 35.42 and 27.61 percent. Thus, heterogeneous expectations raise the wealth share of the top decile from 63.03 to 74.16 percent, an increase of approximately 11 percentage points.

The mechanism is not a mechanical increase in average risk-taking. Indeed, every

TABLE 10. **The Distributional Effect of Heterogeneous Expectations**

	Heterogeneous beliefs				Homogeneous beliefs			
	Top 1%	Next 9%	Middle 40%	Bottom 50%	Top 1%	Next 9%	Middle 40%	Bottom 50%
Wealth share	40.22	33.94	17.18	8.65	35.42	27.61	23.83	13.14
Income share	18.35	25.86	33.57	22.22	23.64	25.03	34.84	16.49
Wage share	9.58	23.80	39.55	27.07	11.88	23.41	41.06	23.64
Risky asset share	89.32	70.11	33.53	13.97	97.60	79.26	51.16	17.87
Expected stock return	8.60	8.14	6.72	5.09	9.07	9.07	9.07	9.07
Portfolio return	12.09	8.97	4.72	2.53	11.95	4.49	2.51	0.68

Notes. The homogeneous-beliefs economy sets the learning gain to zero, so that all households keep the common prior, while preserving the type process for preferences and labor income. All entries are percentages.

wealth group holds a *larger* risky share in the common-belief economy, because the common expected return of 9.07 percent exceeds the average belief in the heterogeneous economy. What heterogeneous learning changes is the cross-sectional sorting of households by beliefs, portfolio exposures, realized returns, and eventually wealth. Under common beliefs, expected returns are identical across households and much of the dispersion in realized portfolio performance disappears. With learning, differences in portfolio experience generate differences in beliefs; these beliefs affect subsequent portfolio decisions, which in turn generate further differences in returns and wealth. The relevant amplification object is therefore the endogenous covariance between beliefs, realized returns, and wealth, rather than simply the mean risky share.

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,

$$\mathbb{E}_{it}[R_{t+1}] = \alpha_i + \mu_t + \varepsilon_{it}. \tag{116}$$

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 model closely reproduces the empirical decomposition. Individual fixed effects explain 65.4 percent of the variation in simulated expected returns, compared with 57.3 percent in RAND. More strikingly, wealth explains only 8.5 percent of the cross-sectional variation in these estimated individual effects, essentially the same

TABLE 11. **Persistent Disagreement: Data and Model**

	Data	Model
R^2: individual fixed effects	57.3%	65.4%
R^2: wealth $\rightarrow \widehat{\alpha}_i$	8.9%	8.5%

as the 8.9 percent in the data.

These two moments summarize an important feature of the mechanism. The model generates a steep relationship between *average* beliefs and wealth while simultaneously leaving most household-level disagreement unexplained by wealth. This is precisely the apparently contradictory pattern documented in RAND: persistent individual heterogeneity is large, wealth has low explanatory power for the individual fixed effects, yet average optimism rises sharply across wealth groups.

Learning from portfolio experience reconciles these facts 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 are therefore jointly generated rather than separately imposed.

6.4. 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 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 3 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

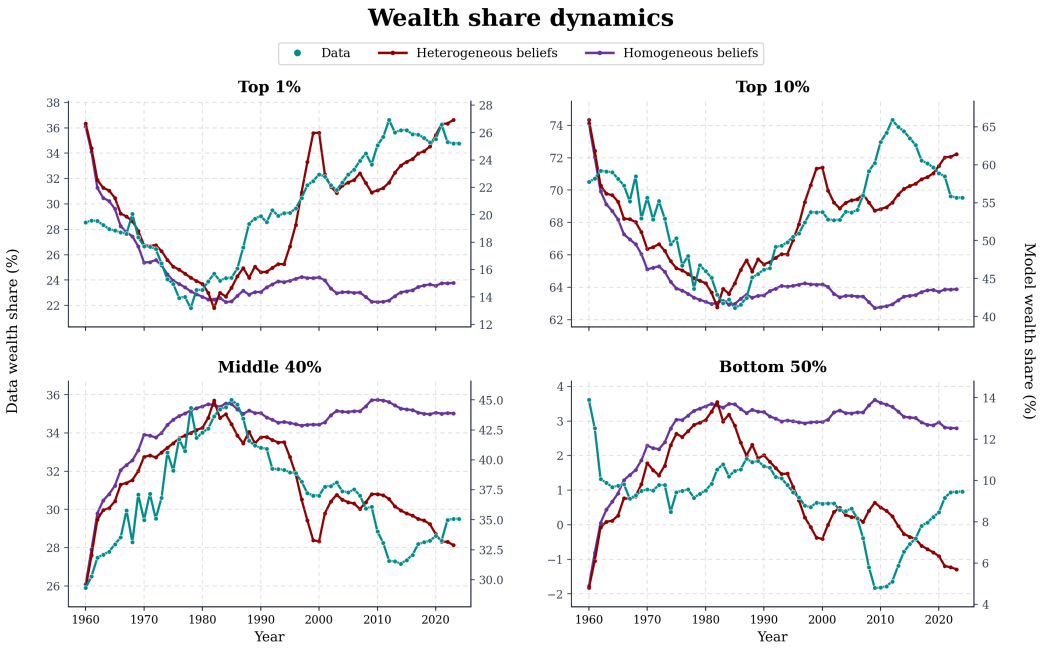


FIGURE 3. **Historical Transitional Dynamics, 1960–2024.** The figure compares observed U.S. wealth shares with model-implied wealth shares under heterogeneous and homogeneous beliefs, conditional on the realized U.S. asset-price and dividend history.

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 economy, the mechanism raises the top-decile wealth share by approximately 11 percentage points. 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. Quantitatively, the model reproduces key features of the joint distribution of wealth, portfolios, and expectations, including the predominance of persistent individual heterogeneity in beliefs despite the relatively limited explanatory power of observable wealth. Removing heterogeneous expectations substantially reduces wealth concentration, both in the stationary economy and along the historical U.S. asset-price path.

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 prices 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 12 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 12. 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:

$$R_{t+k} - \mathbb{E}_{it}[R_{t+k}] = \alpha_i + bR_t + \varepsilon_{it}. \tag{117}$$

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

$$R_{t+k} - \mathbb{E}_{it}[R_{t+k}] = \alpha_i + d\left(\overline{\mathbb{E}}_t[R_{t+k}] - \overline{\mathbb{E}}_{t-1}[R_{t+k}]\right) + \varepsilon_{it}. \tag{118}$$

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 13 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 13. 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,

$$p_{j,it} = \Lambda\Big(\alpha_j + \beta_{W,j} \log W_{it} + \beta_{V,j} V_{it} + \Gamma_j' X_{it}\Big), \tag{119}$$

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.

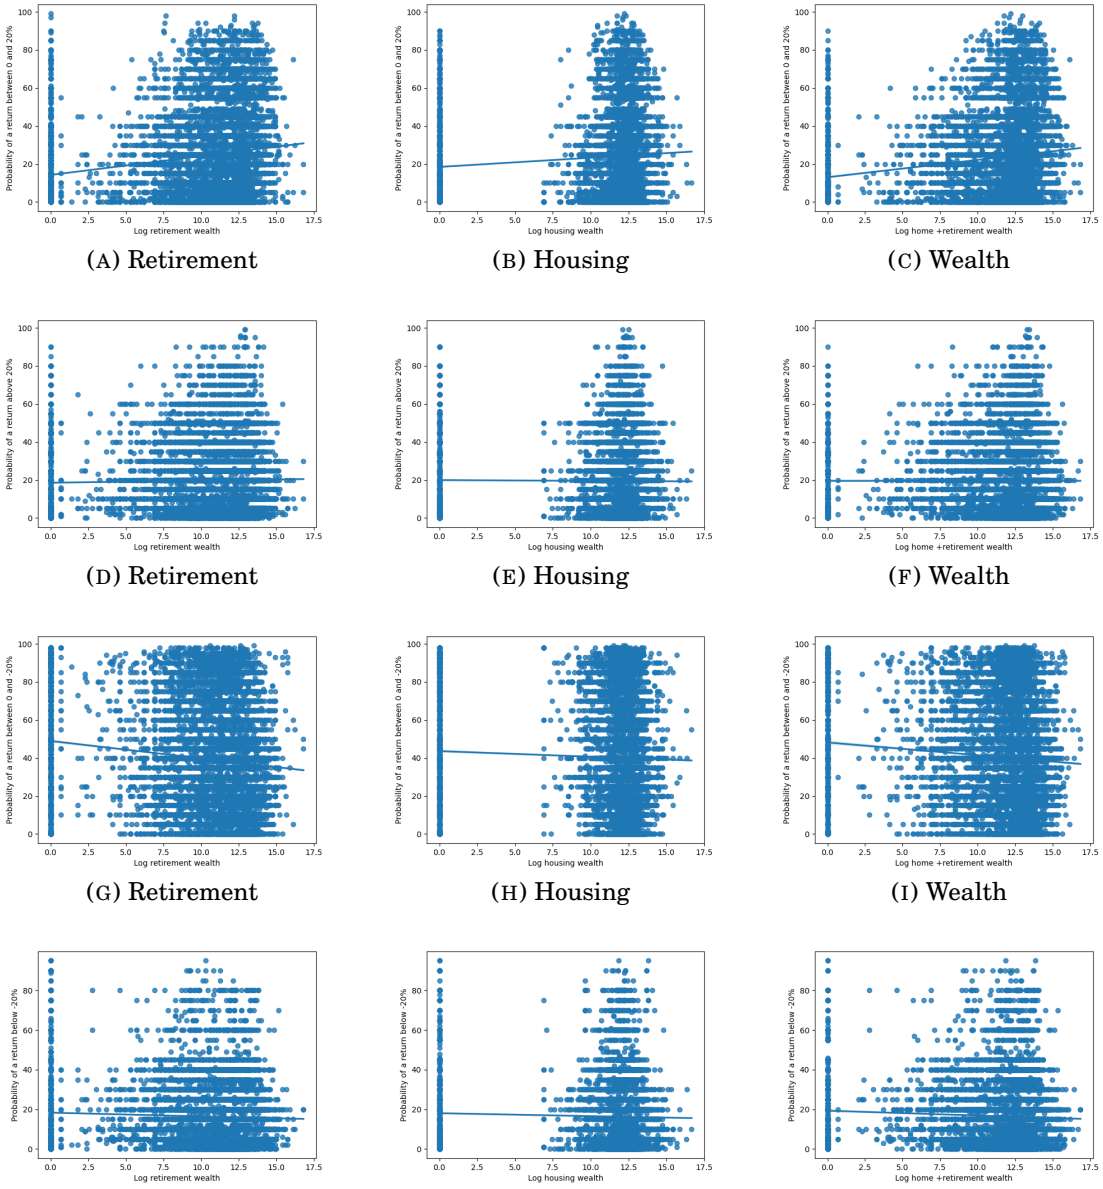
Because the four probability equations are estimated separately, their fitted values are not mechanically constrained to sum to one. We therefore interpret the estimates bin by bin. The predicted-probability counterfactuals reported in Table 2 provide the economically interpretable summary of these regressions.

TABLE 14. 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 (119). 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 4 and 5 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.



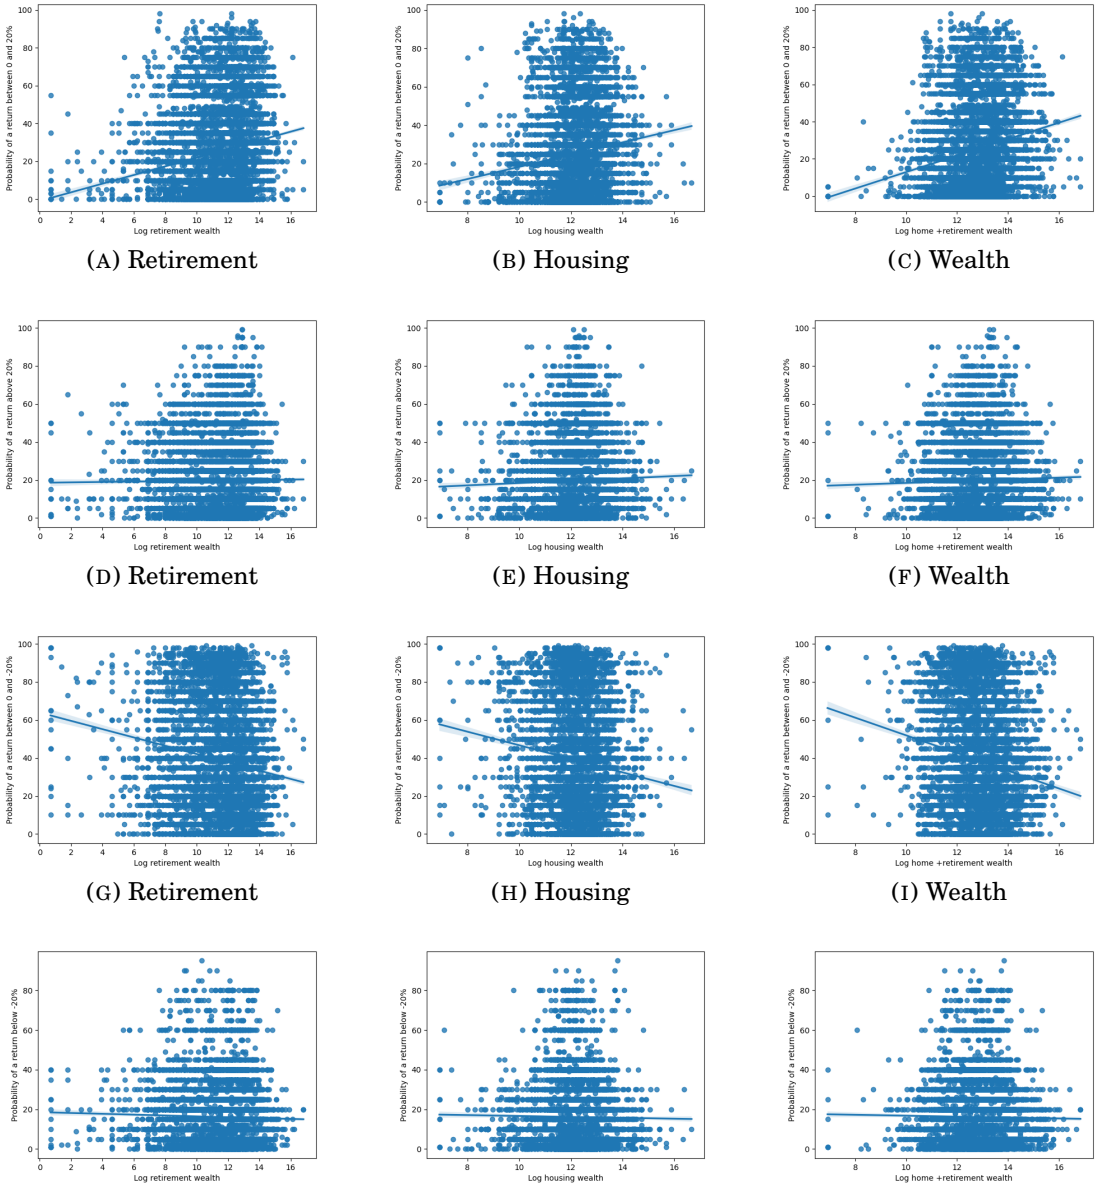
(A) Retirement (B) Housing (C) Wealth

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

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

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

FIGURE 4. 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 5. 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. Substituting $v(\Delta) = \Delta/(\gamma\sigma^2)$ from (45) and $\lambda(\Delta)$ from (55) into (70),

$$\begin{aligned}\xi(\Delta) &= -\frac{2}{\gamma}\left[\log\beta + r_B + \frac{2\Delta^0\Delta - \Delta^2}{2\sigma^2}\right]\frac{\gamma^2\sigma^4}{\Delta^2\sigma_\eta^2} \\ &= \gamma\frac{\sigma^2}{\sigma_\eta^2}\frac{-2\sigma^2(\log\beta + r_B) + \Delta^2 - 2\Delta^0\Delta}{\Delta^2},\end{aligned} \tag{120}$$

which is (72) by the definition of K in (71). Differentiating,

$$\frac{\partial\xi(\Delta)}{\partial\Delta} = 2\gamma\frac{\sigma^2}{\sigma_\eta^2}\frac{\Delta^0\Delta - K}{\Delta^3}, \tag{121}$$

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, (61), reads

$$\lambda(\Delta^0) = \frac{1}{\gamma}\left[\log\beta + r_B + \frac{(\Delta^0)^2}{2\sigma^2}\right] < 0, \tag{122}$$

which by (71) is equivalent to $K > (\Delta^0)^2$, so that $\Delta^\star = K/\Delta^0 > \Delta^0$. Evaluating (72) at $\Delta^\star$,

$$\xi(\Delta^\star) = \gamma\frac{\sigma^2}{\sigma_\eta^2}\frac{K^2/(\Delta^0)^2 - 2K + K}{K^2/(\Delta^0)^2} = \gamma\frac{\sigma^2}{\sigma_\eta^2}\left(1 - \frac{(\Delta^0)^2}{K}\right), \tag{123}$$

and the limit as $\Delta \to \infty$ follows directly from (72). ■

Proof of Proposition 2. Equation (56) implies

$$\lambda(\Delta) = \lambda(\Delta^0) - \frac{(\Delta - \Delta^0)^2}{2\gamma\sigma^2}. \tag{124}$$

Taking expectations under G_d and using $\mathbb{E}_{G_d}[\Delta] = \Delta^0$ yields (84).

For the second result, evaluate (72) at $\Delta^0 - h$ and $\Delta^0 + h$. Their difference is

$$\xi(\Delta^0 - h) - \xi(\Delta^0 + h) = 4\gamma\frac{\sigma^2}{\sigma_\eta^2}\frac{\Delta^0 h\left[K - (\Delta^0)^2 + h^2\right]}{\left[(\Delta^0)^2 - h^2\right]^2}. \tag{125}$$

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

Finally, (121) 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. Equation (79) then gives (87) and (88). ■