Expectations and Attention to Experience

Heiner Beckmeyer* and Ahmed Guecioueur†

February 14, 2025

Abstract

We construct an empirical model of households' perceived links between the stock market and macroeconomy. Our approach enriches experience-based learning with households' own heterogeneous beliefs about economic conditions. The machine learning-estimated model successfully explains a substantial fraction of cross-sectional variation in households' subjective perceptions of the market Sharpe Ratio. Our empirical model replicates the recency bias in return extrapolation, and additionally uncovers an outsize impact of distant past recession periods on households' stock market beliefs. Interpreting the model, we find households are more influenced by lifetime experiences that they believe will be most similar to short-run changes in macroeconomic conditions.

Keywords: Stock Returns, Subjective Sharpe Ratios, Machine Learning, Survey Microdata
JEL Codes: G40, G17, E70, E31

*University of Münster. Email: heiner.beckmeyer@wiwi.uni-muenster.de
†University of Washington, Foster School of Business. Email: ahmedgc@uw.edu

First version: February 16, 2023. We would like to thank Philip Bond, Stefano Cassella, Josefina Cenzón, Theodoros Evgeniou, Roberto Gomez Cram, Naveen Gondhi, Avraham Kamara, Maurizio Montone (discussant), Rodolfo Prieto, Joshua D. Rauh, Stephan Siegel, Mark M. Westerfield, Trevor Young (discussant) and Florian Zimmermann, as well as participants at the 2024 Research in Behavioral Finance Conference, the 12th Helsinki Finance Summit on Investor Behavior, the 2023 Bamberg Behavioral Macroeconomics Workshop, the 2023 China International Conference in Finance, the 2023 Financial Management Association European Conference, the 2023 Financial Markets and Corporate Governance Conference, INSEAD, the University of Washington and the University of Münster for their helpful comments. Guecioueur gratefully acknowledges the support of the Henry Grunfeld Foundation.

1 Introduction

Linking stock market returns with business cycles has been a focal point for research since the seminal work of Grossman and Shiller (1981). While numerous theoretical models couple expected returns to the state of the economy (Cochrane, 2017), empirical studies of households' mental models of the stock market are typically silent on the role of macroeconomic conditions.[1] Gaining a fuller understanding of the subjective models that households employ is crucial for illuminating asset allocation decisions and laying a foundation for asset pricing theories.

In this paper, we construct a general, empirical model of how households arrive at their heterogeneous stock return expectations. Crucially, we incorporate the role played by households' perceived macroeconomic conditions. Our approach thus enriches the existing understanding of households' extrapolative and experience-based expectation formation process (e.g., Malmendier and Nagel, 2011; Barberis, Greenwood, Jin, and Shleifer, 2015; Cassella and Gulen, 2018). We begin by extracting households' subjective perceptions of the stock market's Sharpe ratio from a rich source of survey data. We then use machine learning techniques to construct an empirical model of the expectation formation process that conditions on households' macroeconomic perceptions and characteristics (Das, Kuhnen, and Nagel, 2020).

We use our novel framework to uncover a number of cross-sectional differences in households' mental models of the stock market, irrespective of the objective state of the economy. First, households extrapolate risk-adjusted returns more strongly from past recession periods in their experience than from past expansionary periods. Second, this effect is stronger for households that are more pessimistic about future macroeconomic conditions. More generally, households extrapolate more strongly from past returns that occurred during historical macroeconomic periods that are more similar to their own subjective beliefs about *future* macroeconomic conditions. Finally, households appear to arrive at their perceptions through a form of local, heuristic thinking that sustains their divergent models of the stock market (Gennaioli and Shleifer, 2010; Meeuwis, Parker, Schoar, and Simester, 2022).

Our paper begins by introducing a rich source of microdata on subjective expectations about the return of the market portfolio. We show how to recover investors' subjective perceptions of the Sharpe Ratio of the stock market from a long-running and well-known source of survey microdata with sufficient length to cover the business cycle – the Michigan Surveys of Consumers. Our approach relies on a parametric assumption inspired by the prior literature on survey expectations (Dominitz and Manski, 2007; Ben-David, Graham, and Harvey, 2013;

[1] In two notable exceptions, Amromin and Sharpe (2014) and Giglio, Maggiori, Stroebel, and Utkus (2021) find that households perceive a positive relationship between various indicators of macroeconomic expansion and the Sharpe ratio or expected return (respectively) of the market portfolio. We estimate a model of the household expectation formation process that captures this regularity, and others.

Lochstoer and Muir, 2022), and which yields a simple transformation of survey responses to infer each respondent's subjective Sharpe Ratio (SSR). The SSRs that our procedure recovers are procyclical and extrapolative, consistent with previous findings on return expectations (Greenwood and Shleifer, 2014). Importantly, SSRs predict household demand: in keeping with basic portfolio choice intuition, the share of household wealth devoted to stock market risk-taking is increasing in the level of the subjective Sharpe Ratio, and the relationship holds in the cross-section of individual households as well as in aggregate.

Our dataset is uniquely advantageous for the goal of learning how households perceive the link between the stock market and the business cycle. First, we are able to exploit two decades of monthly survey samples that cover the full business cycle. Second, Michigan survey respondents are sampled to be representative of the US population, yielding rich heterogeneity in their beliefs and sociodemographic characteristics. Third, and perhaps most importantly, for each household we observe its beliefs about various macroeconomic attributes (business conditions, interest rates, unemployment, inflation) as well as recover its subjective Sharpe ratio of the market portfolio, as *jointly* perceived by that individual household. We are thus able to use this dataset to estimate the within-household mental model that subjectively links the stock market to the business cycle.

This paper's second methodological contribution is to introduce a machine learning-based methodology for estimating households' mental models of the stock market. Our point of departure is the experience-based learning framework of Malmendier and Nagel (2011), which captures the influence of historical stock market returns on households' expectations today, via a weighted average of their past lifetime experiences. We loosen the assumption of exponentially-decaying weights on past experiences and allow within-household beliefs about macroeconomic conditions to flexibly instrument these weights. We similarly include observed demographic characteristics, such as socioeconomic status (Das, Kuhnen, and Nagel, 2020) as a second set of conditioning variables. To flexibly estimate household-level experience effect weights that best explain individual households' perceived SSRs, we introduce a customized Transformer neural network architecture (Vaswani et al., 2017).[2]

In estimating this model, we are able to explain almost a fifth of both aggregate and cross-sectional variation in households' risk-adjusted expected returns. These results strengthen the case for experience effects as an organizing framework for understanding households' expectation formation process. Remarkably, the estimated experience effects exhibit an exponential-

[2]The economic restriction that experience weights must be non-negative and sum to unity for each household also implicitly restricts the functional forms that are estimated to best fit the data. Note that our estimation framework must meet a technically challenging set of criteria: conforming to this economic restriction, while allowing flexible and yet interpretable experience weights to be extracted, and taking sequential dependencies among time series of experienced realizations into account. Section 3.2 describes our framework further.

like decay on average, although no explicit functional form was specified at the outset: our estimates thus replicate the well-known recency bias in return extrapolation. Importantly, experience weights also exhibit non-monotonicity and substantial heterogeneity, leading us to investigate the presence of additional mechanisms.

We find that estimated experience weights for past returns that occurred during NBER recession periods are on average higher than during periods of expansion; that is, past recessions that a household has experienced exert an outsize influence on its beliefs about future stock returns. We arrive at this conclusion by conducting panel regressions on Transformer-estimated experience weights, which furthermore allows us to assess statistical significance. We also find that this effect strengthens when a respondent is surveyed during a current recession, and when the respondent's beliefs about the macroeconomy are more pessimistic.

To understand the underlying mechanism, we next compute measures of similarity between every respondent's perceived macroeconomic conditions over the next year and macroeconomic conditions during every year that the respondent has lived through. We find that, in general, respondents weigh past experienced returns more highly when past economic conditions at that time are more similar to the respondent's beliefs about the short-run evolution of the macroeconomy. Households thus attempt to predict returns by attending to and retrieving past returns that covary the most with perceived future macroeconomic conditions.

Is this particular behavior rational? Our final test exploits a shock to household expectations employed by Meeuwis, Parker, Schoar, and Simester (2022): the surprise Presidential election result in 2016. We find that households before and after the change of political environment adjust the weights they assign to past experiences in a way that accords with their stated political affiliations. Democrats place a higher weight on past recessions (pre-to-post election), and Republicans by contrast assign a lower weight. Interestingly, this effect is robust to controlling for respondents' own stated perceptions of future macroeconomic conditions. It is therefore likely that households make use of local, heuristic thinking in forming their models of the stock market return (Gennaioli and Shleifer, 2010).

Related Literature We contribute to a recent literature in financial economics that studies the determinants of investors' return expectations using survey data (Dominitz and Manski, 2011; Malmendier and Nagel, 2011; Amromin and Sharpe, 2014; Greenwood and Shleifer, 2014; Giglio, Maggiori, Stroebel, and Utkus, 2021). Our work focuses on households (Gomes, Haliassos, and Ramadorai, 2021), and we produce a measure of households' perceived risk-adjusted return expectations that predicts households' demand for equities.

Our goal is to understand the perceived link between macroeconomic conditions and stock

returns, within a household.[3] Amromin and Sharpe (2014) find that households typically perceive a positive association between economic conditions and stock returns: we additionally incorporate the role of experience effects and use a machine learning approach to fit household survey data in order to uncover the role of attention to past returns. Goetzmann, Watanabe, and Watanabe (2022) and Boons, Ottonello, and Valkanov (2024) focus on the correlations of return and macroeconomic forecasts made by professional economists. Andre, Schirmer, and Wohlfart (2024) qualitatively examine households' mental models of the equilibrium between stock payoffs and returns. Andre, Pizzinelli, Roth, and Wohlfart (2022) analyze the drivers of macroeconomic expectations in general. Our paper sheds further light on cross-household model disagreement (Meeuwis, Parker, Schoar, and Simester, 2022), which is an important driver of trading activity (Cookson and Niessner, 2020).

We add to a small literature in financial economics that applies Transformer neural networks; e.g., Beckmeyer and Wiedemann (2022) analyze the effects of missing firm-level characteristics. To the best of our knowledge, our paper is the first to use a Transformer neural network to produce an empirical model of the behavior of an economic agent.[4]

Our approach is loosely related to a recent stream of literature that assesses the quality of professionals' forecasts of macroeconomic variables (Bianchi, Ludvigson, and Ma, 2022) or stock prices (Cao, Jiang, Wang, and Yang, 2024) by comparing them to machine learning forecasts. However, we estimate a machine learning model upon the forecasts produced by individuals, rather than on the actual realizations: our use of machine learning techniques is to directly predict and understand human behavior, rather than set up a benchmark for comparison. In that sense, our work is closer to Bose et al. (2022), who estimate a neural network on visual salience data produced in the laboratory by human subjects. Another relevant study by Chen and Huang (2023) employs a particular neural network structure to measure the mental contexts used by equity analysts, assuming their earnings forecasts stem from a memory mechanism. We focus on general experience effects that may nest multiple mechanisms, and households' jointly-held expectations of returns and economic conditions.

Finally, we build on a small literature studying investors' beliefs about volatility (Lochstoer and Muir, 2022; Sheikh, 1993). Our measure of households' expected returns captures both the first and second moments of the return distribution and can be mapped to a subjective Sharpe Ratio, allowing us to study an important economic quantity for investors' portfolio choices.

[3]Studying the subjective perceptions of households (instead of the forecasts of professional economists or analysts) means our findings are not driven by strategic forecasting behavior. Another advantage of focusing on households is that their subjective expectations directly determine asset demand.

[4]To the extent that Large Language Models (LLMs) are implemented using Transformer architectures, our paper also loosely relates to a fast-growing stream of literature that uses LLMs to study human behavior (e.g., Bybee (2023) and Fedyk, Kakhbod, Li, and Malmendier (2024)).

Outline The remainder of this paper proceeds as follows. Section 2 describes our survey microdata and method for recovering subjective Sharpe Ratios, and presents a number of key stylized facts. Section 3 explains our machine learning-based procedure for estimating experience effects, and assesses its performance. Section 4 analyzes estimated experience effects to trace out household's perceived links between the stock market and macroeconomic conditions. Section 5 concludes. The Appendix contains all proofs. An Internet Appendix contains further methodological details, additional analyses, and data validation checks.

2 Expectations About the Market Sharpe Ratio

2.1 Recovering Subjective Sharpe Ratios from Survey Microdata

To study how investors form expectations about future stock market returns, we exploit 20 years (2002-2022) of detailed monthly microdata from the Michigan Surveys of Consumers, in which respondents are asked the following question: *"What do you think is the percent chance that a one thousand dollar investment in a diversified stock mutual fund will increase in value in the year ahead, so that it is worth more than one thousand dollars one year from now?"* That is, at each survey date t each respondent i is asked to provide the following subjective expectation about the year-ahead market return:

$$\mathbb{E}_{it}[\mathbb{1}\{r_{t+1} > 0\}] \;=\; \mathbb{P}_{it}(r_{t+1} > 0) \tag{1}$$

In typical portfolio choice models, investors choose portfolios based on the first and second moments of their perceived distributions of future asset returns. In the case of the market return, these two moments $\mathbb{E}_{it}[r_{t+1}]$ and $\mathbb{V}\text{ar}_{it}(r_{t+1})$ are closely linked to the cumulative distribution function $\mathbb{P}_{it}(r_{t+1} \leq x)$ that the investor must estimate when she reports the subjective expectation above, in Eqn. (1). Proposition 1 shows that we can recover each investor's subjective Sharpe Ratio from her stated probability by applying a simple probit transformation:

Proposition 1 (Subjective Sharpe Ratio). *Assume that an individual i perceives the market return to be conditionally normally distributed, $r_{t+1} \sim \mathcal{N}(\mu_{it}, \sigma_{it})$, and the riskless rate to be zero. Then*

$$SSR_{it} = \frac{\mu_{it}}{\sigma_{it}} \;=\; \text{probit}\big(\mathbb{P}_{it}(r_{t+1} > 0)\big) \tag{2}$$

is her subjective perception of the market's conditional Sharpe Ratio over the next period.

Under the above set of assumptions, there is thus a simple one-to-one mapping between the elicited quantity $\mathbb{P}_{it}(r_{t+1} > 0)$ and the investor's perceived Sharpe Ratio of the market

portfolio. An investor's beliefs about the Sharpe Ratio are a key determinant of her portfolio choice decision, under a range of horizons and preference structures: in a choice between a risky and a riskless asset, the investor's optimal portfolio share of the risky asset is typically directly proportional to her perceived Sharpe Ratio of the risky asset. For example, an investor with mean-variance preferences (with risk aversion parameter γ_i) who chooses which fraction x of her wealth to allocate to the risky market portfolio returning r_{t+1} seeks to maximize

$$U_i(x_{it}) = x_{it}\mu_{it} - \frac{\gamma_i}{2}x_{it}^2\sigma_{it}^2, \tag{3}$$

and therefore chooses to allocate a fraction of her wealth

$$x_{it} = \frac{1}{\gamma_i\sigma_{it}} \cdot \frac{\mu_{it}}{\sigma_{it}} \tag{4}$$

to the risky market portfolio, in direct proportion to her perception of the Sharpe Ratio $\frac{\mu_{it}}{\sigma_{it}}$. Proposition 1 therefore allows us to recover a crucial economic quantity.

Discussion Internet Appendix A discusses the assumptions underlying Proposition 1, and empirically verifies their reasonableness for our sample period and setting.

Recovered Subjective Sharpe Ratios (SSRs) Figure 1 plots the cross-sectional mean over time of households' recovered SSRs. We observe procyclical variation over the business cycle. In addition, the dispersion measures shown in Figure 2 indicate the presence of substantial heterogeity in SSRs between households surved in the same year-month.

2.2 Additional Microdata Variables

In addition to the subjective probability of a market up-move, we make use of a variety of respondent attributes recorded by the Michigan Surveys of Consumers across our sample period from June 2002 to May 2022. This period covers the full business cycle, including a financial crisis and the Covid shock. Most importantly, during this time period the Michigan surveys consistently elicited respondents' subjective expectations of a stock market increase (Equation 1), from which we recover respondents' joint perceptions of the first two moments of the market return distribution (i.e. their Subjective Sharpe Ratios).

The Michigan surveys also record information about the demographic characteristics of respondents. We make use of the following demographic variables: age, gender, education level, geographic region, political affiliation, household income, stock holdings and home value.

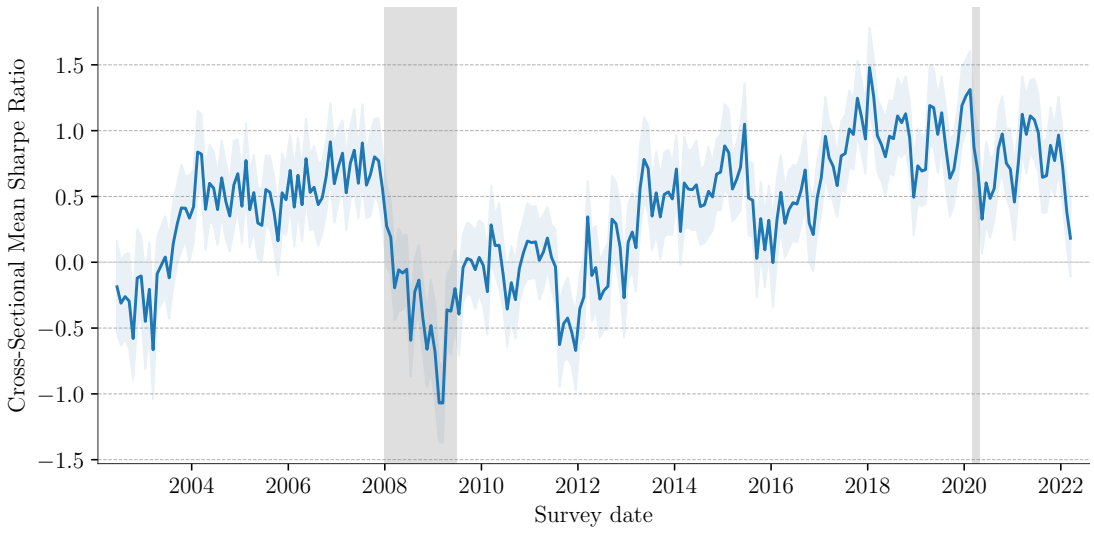


Figure 1: Mean of Survey Subjective Sharpe Ratios

This chart thus depicts the mean subjective Sharpe Ratio perceived by survey respondents. The time series shows the cross-sectional mean of probit-transformed responses on the probability of an up-move in the stock market over the next twelve months (see Proposition 1), at a monthly frequency. Blue-shaded areas are the ±2 standard error range. Gray-shaded areas are NBER recession periods.

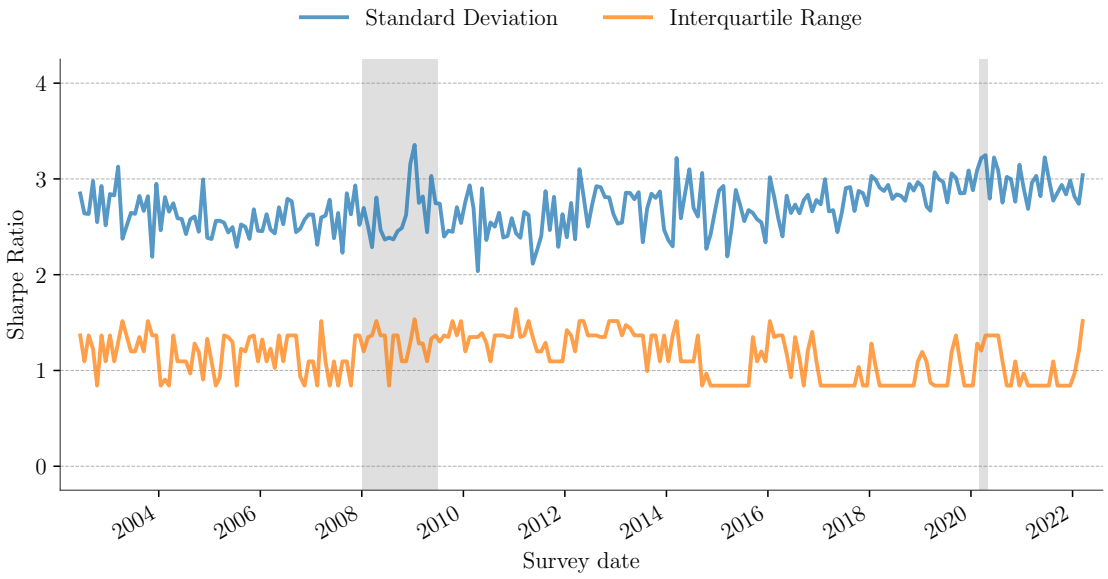


Figure 2: Dispersion of Survey Subjective Sharpe Ratios

This chart depicts cross-sectional heterogeneity in the subjective Sharpe Ratio perceived by survey respondents. Both series show cross-sectional dispersion measures (SD or IQR) of probit-transformed responses on the probability of an up-move in the stock market over the next twelve months (see Proposition 1), at a monthly frequency. Shaded areas are NBER recession periods.

Table 1: Michigan Survey Variables

This table describes and summarizes variables extracted from responses to the Michigan Surveys of Consumers. The 3rd column names the Michigan survey question from which each variable is originally sourced. The last two rows describe categorical variables: for these rows the last two columns show the most frequent category and the number of occurrences of the most frequent category, respectively. The summary statistics in this table are for raw values of the Michigan variables, not for cross-sectional percentiles.

Variable							Percentile				Category		
Name	**Unit**	**Survey Q.**	**N**	**Mean**	**S.D.**	**Min**	**25**	**50**	**75**	**Max**	**Unique**	**Top**	**Count**
Age	Years	AGE	81,388	52.2	15.71	18	40	53	64	97			
Household income	USD	INCOME	78,358	108,302.90	85,855.88	1,200	55,000	85,000	132,000	500,000			
Stock holdings value	USD	INVAMT	67,893	332,204.22	793,601.04	1,000	30,000	100,000	300,000	10,000,000			
Home value	USD	HOMEAMT	55,307	363,983.42	390,612.18	1,000	160,000	255,000	425,000	5,000,000			
Forecast probability of increase in stock market over next year	Percent	PSTK	81,879	54.81	29.24	0	30	55	80	100			
Forecast probability of increase in personal income over next year	Percent	PINC2	81,301	52.46	36.77	0	10	50	85	100			
Forecast 1Y inflation rate	Percent	PX1Q2	65,719	4.35	3.51	1	2	3	5	40			
Forecast 5Y inflation rate	Percent	PX5Q2	74,293	3.23	2.35	1	2	3	4	30			
Forecast direction of business conditions over next year	{-1, 0, +1}	BEXP	80,411	0.11	0.74	-1	0	0	1	1			
Forecast direction of interest rates over next year	{-1, 0, +1}	RATEX	81,135	0.47	0.66	-1	0	1	1	1			
Forecast direction of unemployment over next year	{-1, 0, +1}	UNEMP	81,340	0.05	0.73	-1	0	0	1	1			
Forecast percent increase or decrease in local house prices over next year	Percent	HOMPX1Q2	27,775	6.34	5.13	1	3	5	9	35			
Is female	True/False	SEX	81,879	0.45	0.50	0	0	0	1	1			
Supports Republican party	True/False	POLAFF	36,158	0.30	0.46	0	0	0	1	1			
Supports Democratic party	True/False	POLAFF	36,158	0.32	0.47	0	0	0	1	1			
Supports current President's party	True/False	POLAFF	36,158	0.31	0.46	0	0	0	1	1			
Region of residence	Category	REGION	81,878								4	South	27,306
Education level	Category	EDUC	81,696								6	Grade 13-16 w/ col degree	25,529

In addition, the Michigan surveys record respondents' beliefs concerning general economic conditions, and other economic variables besides the stock market. Since these beliefs tend to be associated within a household (Amromin and Sharpe, 2014), these macroeconomic beliefs are candidates for inclusion in households' mental models of the stock market. We therefore make use of the following belief-related variables: probability of an increase in personal income, direction of business conditions, direction of interest rates, direction of unemployment, and changes in local house prices, all elicited for a horizon of one year.

Our main data processing step is to include only survey responses in which the respondent provides her subjective probability of a positive market return over the next year – our main variable of interest – and in which the respondent reported that she or a member of her household own any stocks.[5] After this filtering step, our survey microdata comprises 81,879 responses. Table 1 presents summaries of the Michigan survey variables used in our study.

2.3 Relationship Between Stock Market and Macroeconomic Beliefs

Since we observe a series of measured expectational values that are held jointly by each household, we now examine the direction of association between these beliefs. Table 2 shows within-household correlations between beliefs. Similarly to Amromin and Sharpe (2014), we find that households are more optimistic about the market's SSR when they are optimistic about macroeconomic conditions, over the same short-run horizon of the year ahead.

3 Estimating an Empirical Model

We now describe a framework that can estimate households' mental models of the stock market. We begin with a classic model of experience-based learning, show how we enrich it for our goals, and then present model estimates.

3.1 Classic Experience-Based Learning

Our point of departure is the Malmendier and Nagel (2011) experience-based learning framework: households form their return expectations based on the past returns that they have experienced in their lifetimes. More precisely, their framework implies that the Subjective Sharpe Ratio perceived by an individual i surveyed at time t, $\text{SSR}_{i,t}$, is formed based on an

[5] To be precise, we require a response of Yes to the question *"Do you (or any member of your family living there) have any investments in the stock market, including any publicly traded stock that is directly owned, stocks in mutual funds, stocks in any of your retirement accounts, including 401(K)s, IRAs, or Keogh accounts?"* (question INVEST). 72% of respondents answered Yes to this question in May 2022.

Table 2: Within-Household Associations Between Stock Market and Macroeconomic Beliefs

This table shows within household correlations between the perceived Subjective Sharpe Ratio (SSR) and each of a series of jointly-held beliefs about macroeconomic conditions. All beliefs are for a consistent 1-year horizon from the survey date onward. The *** asterisks beside each correlation indicate a t-test rejects the null of zero correlation at a 99% significance level for each.

Variable	Correl. to SSR	
Prob. of Increase in Personal Income	0.33	***
Direction of Business Conditions	0.23	***
Direction of Interest Rates	0.04	***
Prob. of Increase in House Prices	-0.10	***
Prob. of Increase in Prices	-0.16	***
Direction of Unemployment	-0.23	***

exponentially-weighted average of all the past lagged Sharpe Ratios SR_{t-k} that have been realized during that individual's lifetime:

$$\text{SSR}_{i,t} = \sum_{k=1}^{\text{age}_{i,t}-1} w_{i,t}(k, \text{age}_{i,t}) \times \text{SR}_{t-k} \tag{5}$$

$$\text{where} \quad w_{i,t}(k, \text{age}_{i,t}) = \frac{(\text{age}_{i,t} - k)^{\lambda}}{\sum_{k=1}^{\text{age}_{i,t}-1} (\text{age}_{i,t} - k)^{\lambda}} \tag{6}$$

In the above, positive scalar weights $w_{i,t}(k, \lambda)$ sum to one for each agent. The hyperparameter λ is common to all agents and controls the relative weighting placed on recent returns versus distant past returns: when $\lambda > 0$, weights decrease as the lag k increases, so that more recent experiences receive more weight than more distant experiences.[6]

While this is a parsimonious specification for an individual's expectation formation process, there is no room for beliefs about macroeconomic conditions to enter into the expectation formation process (Amromin and Sharpe, 2014). Nor do additional demographics, such as socioeconomic status (Das, Kuhnen, and Nagel, 2020), play a direct role. We thus generalize

[6]When $\lambda < 0$, more recent experiences receive less weight than more distant experiences. When $\lambda = 0$, all experiences receive an equal weight. The special case of $\lambda = 1$ corresponds to linearly decaying weights, and $\lambda \to \infty$ converges to all weight placed on the most recent point. Malmendier and Nagel (2011) estimate λ to be roughly between 1 and 2 on their sample, such that respondents place a larger weight on more recent stock market returns. At the same time, because age is the sole demographic input governing cross-sectional differences in weights $\mathbf{w}$ and λ is fixed, younger respondents mechanically place a higher weight on recent market returns than older respondents do.

this framework.

3.2 Model Specification

We now describe our own approach to estimating the weights that individuals place on their experiences, focusing on building economic intuition. A more detailed discussion of each modeling step is provided in Appendix G. Much like the classic exponential weighting framework described in Eqns. (5)-(6), we model the subjective expectation held by each individual i surveyed at time t as a weighted sum of her past experienced Sharpe Ratios SR_{t-k}, with $k \in [1, \text{age}_{i,t} - 1]$. We depart in using the respondent's demographics $\mathbf{d}$ and beliefs about macroeconomic state variables $\mathbf{b}$ to instrument for the across-household variation in experience effects.

We achieve this by jointly modeling each individual's experience weights vector $\mathbf{w}_{i,t}$ (consisting of $(\text{age}_{i,t} - 1)$ experience weights $w_{i,t}$) as a high-dimensional nonlinear function of the respondent's demographics $\mathbf{d}_{i,t}$ and other beliefs $\mathbf{b}_{i,t}$. Our framework is the following:

$$\text{SSR}_{i,t} = \sum_{k=1}^{\text{age}_{i,t}-1} w_{i,t}(k, \mathbf{d}_{i,t}, \mathbf{b}_{i,t}) \times \text{SR}_{t-k} \tag{7}$$

$$\text{where} \quad \mathbf{w}_{i,t} = g(\text{age}_{i,t}, \mathbf{d}_{i,t}, \mathbf{b}_{i,t}) \tag{8}$$

To ensure that each individual's experience weights vector $\mathbf{w}_{i,t}$ captures the economic mechanism of experience effects, we require that

$$\sum_{k=1}^{\text{age}_{i,t}-1} w_{i,t}(k, \mathbf{d}_{i,t}, \mathbf{b}_{i,t}) = 1 \quad \text{and} \quad w_{i,t}(k, \mathbf{d}_{i,t}, \mathbf{b}_{i,t}) \geq 0 \ \ \forall \ \ k \in [1, \text{age}_{i,t} - 1]. \tag{9}$$

In our specification above, estimating each individual's experience weights vector $\mathbf{w}_{i,t}$ reduces to estimating the function $g(\cdot)$, which we achieve using a flexible neural network that naturally embeds nonlinear effects and interactions between different demographics and beliefs. We rely on recent advances from the computer science literature and implement a Transformer neural network to estimate the above specification for all households.

Before that, we first note that beliefs about macroeconomic conditions play a clear role in Equation (7): each household will attend to potentially different past experiences depending on its perception of the short-run evolution of the state of the macroeconomy, which is encoded in the household's macroeconomic beliefs vector $\mathbf{b}_{i,t}$. Prior findings in the literature support our assumption that macroeconomic beliefs are an input into individuals' return expectation formation process: Boons, Ottonello, and Valkanov (2024) find that professional forecasters'

stock market beliefs causally respond to macro shocks. Even clearer evidence is provided by Von Gaudecker and Wogrolly (2022, Figure 1), who survey households on what they most thought about when forming subjective stock market return probabilities, and find the "state of the economy" was the dominant response.

3.3 Estimation Procedure

The Transformer was introduced by Vaswani et al. (2017) and has shown success in the tasks of natural language processing, image recognition and multi-modality learning.[7] The central innovation by Vaswani et al. (2017) is the so-called attention mechanism, of which we use a variant that compares two sets of information, known as *cross-attention*. Consider that we have a database of information extracted from past Sharpe Ratios that realized over the lifetime of respondent i, $\mathbf{SR}_{(t-\text{age}_{i,t}-1):t}$. The database is characterized by keys $\mathbf{k}$ and values $\mathbf{v}$. $\mathbf{k}$ is a nonlinear transformation of $\mathbf{SR}_{(t-\text{age}_{i,t}-1):t}$, which is directly learned from the data:[8]

$$\mathbf{k}_{i,t} = \mathbf{W}_2^{\mathbf{k}} \times \text{GELU}\left(\mathbf{W}_1^{\mathbf{k}} \times \mathbf{SR}_{(t-\text{age}_{i,t}-1):t} + \mathbf{a}_1^{\mathbf{k}}\right) + \mathbf{a}_2^{\mathbf{k}}. \tag{10}$$

GELU (Hendrycks and Gimpel, 2016) is a smoother variation of the commonly-used ReLU activation function, which typically produces better results in the family of Transformers. Let $\mathbf{v} = \mathbf{SR}_{(t-\text{age}_{i,t}-1)}$. Once we set up the database on past Sharpe Ratios, we may retrieve information from it using queries $\mathbf{q}$. Given that we instrument the variation in experience effects through demographics and various macroeconomic beliefs, queries $\mathbf{q}$ are obtained from demographics $\mathbf{d}_{i,t}$ and macroeconomic beliefs $\mathbf{b}_{i,t}$. In a first step, we extract interaction effects across and nonlinearities within different demographic characteristics and various macroeconomic beliefs. The details of this step are given in Appendix G. After this interaction step, we obtain queries $\mathbf{q}$ through a linear transformation:

$$\mathbf{q}_{i,t} = \mathbf{W}^{\mathbf{q}} \times [\widehat{\mathbf{d}}_{i,t}, \widehat{\mathbf{b}}_{i,t}], \tag{11}$$

Attention is related to relational database lookups (e.g. SQL databases), which retrieve the values underlying a query if and only if the query is found in the database's keys. Instead, attention acts as a probabilistic lookup, wherein we receive a weighted average of the values, for which the corresponding keys are most closely related to the query. This property makes it

[7] For example, OpenAI's ChatGPT, a now well-known machine learning product, is itself a Transformer adapted to language tasks.

[8] For the sake of simplicity we discuss the attention mechanism for keys, values and queries for a single entity i. In the implementation, we make use of matrix multiplications which speed up the computation considerably.

appealing for the extraction of experience effects:

$$\mathbf{Attn}_{i,t} = \frac{\mathbf{q}_{i,t} \times \mathbf{k}'_{i,t}}{\sqrt{d^{\text{Attn}}}} \tag{12}$$

We compute the dot-product similarity between the queries, obtained from a respondent's demographics and other beliefs, and keys to the internal "database," which come from experienced past Sharpe Ratios. We normalize this dot-product similarity by the square-root of the embedding size, which stabilizes predictions.

We then transform the raw attention scores to experience weights through SoftMax normalization, which informs us about which past Sharpe Ratios (indexed by k) are most informative for explaining the subjective Sharpe Ratios of respondents who share a given combination of demographics and other beliefs:[9]

$$w_{i,t}(k, \mathbf{d}_{i,t}, \mathbf{b}_{i,t}) = \frac{e^{\text{Attn}_{i,t}(k)}}{\sum_{k=1}^{\text{age}_{i,t}-1} e^{\text{Attn}_{i,t}(k)}}, \quad \forall\ k \in [1, \text{age}_{i,t} - 1] \tag{13}$$

In words, we assign a larger weight to past stock market Sharpe Ratios that are more informative about the expectation formation process of respondents with a certain set of demographics and beliefs about macroeconomic state variables.

We estimate the experience weights such that predicted $\widehat{\text{SSR}}_{i,t}$ values best explain the variation in households' observed SSRs. To do this, we minimize the mean squared error of model-predicted Subjective Sharpe Ratios across subjects i and months t:

$$\text{MSE} = \sum_{i}^{I} \sum_{t}^{T} \left(\widehat{\text{SSR}}_{i,t} - \text{SSR}_{i,t} \right)^2 \tag{14}$$

As is standard for neural networks, we estimate the model parameters using stochastic gradient descent. Specifically, we opt for the Adam optimizer with proper weight decay (Loshchilov and Hutter, 2017), which combines accurate and stable model parameter estimation with fast convergence. We also follow the paper by Reddi, Kale, and Kumar (2019) which discusses better convergence for Adam.

We estimate the model using data from 2002 to the end of 2019, holding out the 2020-2022 period for out-of-sample verification.

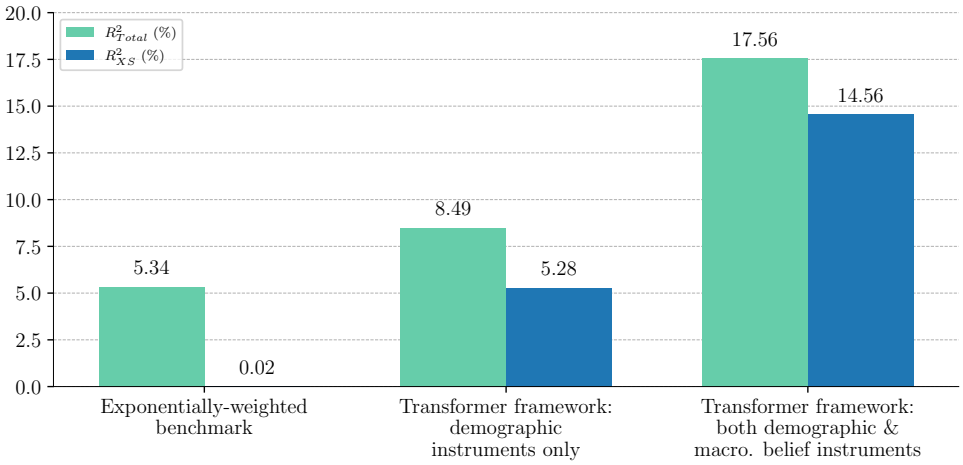


Figure 3: Explanatory Power Using Different Models Of Experience Effects

This figure compares the proportion of variance in respondents' subjective Sharpe Ratios that is explained by three empirical models of experience effects. These models incorporate past experiences in different ways. In the first model, individuals adaptively weight their past experiences according to Malmendier and Nagel (2011). The second two experience effects models are estimated by our Transformer framework, with different sets of instruments used for experience weights. Blue columns report the overall fraction (%) of variance explained R^2_{Total} (15), while green columns focus on the fraction (%) of cross-sectional variance explained R^2_{XS} (16).

3.4 Explanatory Power of Model-Implied SSRs

As a first result, we investigate how well our models of experience effects perform in explaining respondents' Subjective Sharpe Ratios, $\text{SSR}_{i,t}$. Figure 3 displays the performance of Transformer-based experience effects models (Equation (7)), alongside a classic exponentially-weighting baseline.[10] Our two Transformer-estimated models differ solely in which variables are used as instruments for experience weights in Equation (8): the first uses only demographic variables $\mathbf{d}_{i,t}$ (including age), and the second augments these with households' own beliefs about macroeconomic state variables $\mathbf{b}_{i,t}$.

We begin by reporting the classic fraction of total variance R^2_{Total} explained by our models. We then also assess how much of the *cross-sectional* variation of $\text{SSR}_{i,t}$ can be explained in comparison to a benchmark of the cross-sectional mean in month t, using an R^2_{XS} metric that is

[9]Note that the SoftMax normalization does *not* impose an exponential decay of experience weights, but is merely a convenient way to normalize weights and abide by the restrictions in Equation (9).

[10]The baseline model is described in Equation (5). We set the optimal value of $\widehat{\lambda} = 10.39$ for this exponentially weighted baseline using a grid search, see Appendix E.

equivalent to the "within-R^2" of a panel regression with time effects:[11]

$$R^2_{Total} = 1 - \frac{\sum_{i=1}^{N}(\text{SSR}_{i,t} - \widehat{\text{SSR}}_{i,t})^2}{\sum_{i=1}^{N}(\text{SSR}_{i,t} - \overline{\text{SSR}})^2} \tag{15}$$

$$R^2_{XS} = 1 - \frac{\sum_{i=1}^{N}\left((\text{SSR}_{i,t} - \overline{\text{SSR}_t}) - (\widehat{\text{SSR}}_{i,t} - \overline{\widehat{\text{SSR}}_t})\right)^2}{\sum_{i=1}^{N}(\text{SSR}_{i,t} - \overline{\text{SSR}_t})^2} \tag{16}$$

The R^2 results of Figure 3 show that our Transformer-estimated models of experience effects explain a large fraction of the variation of survey respondents' subjective expectations. Our approach succeeds in explaining 17.6% of the overall and 14.6% of the cross-sectional variation when both macroeconomic belief and demographic variables are used as instruments for experience weights. Although not intended as an exhaustive comparison between approaches, we note that a baseline that exponentially weights past experiences does not capture cross-sectional variation, with an R^2_{XS} value of 0.02%. We emphasize that demographic variables are not used as direct predictors by any of the models in Figure 3 – they merely instrument the weights nonlinearly in the Transformer estimation framework.

When excluding macroeconomic beliefs from the set of instruments on experience weights, our Transformer-based model continues to outperform the adaptive weighting baseline, with an R^2_{Total} of 8.5% and R^2_{XS} of 5.3%. However, the dropoff in explanatory power suggests that the excluded variables are useful in capturing cross-sectional heterogeneity across households. In other words, incorporating macroeconomic beliefs into households' mental models of the stock market is important to explain the risk-adjusted return expectations that they hold.

The strong performance of our approach indicates that pure experience effects can explain a large fraction of the variance in respondents' subjective expectations, both in total and within each cross-section – so long as the weights on past experiences are allowed to be driven by households' own jointly-held macroeconomic beliefs, as well as the role played by their demographic characteristics (in addition to age).

Out-of-Sample Test In addition to checking the explanatory power in-sample, we also verify that the performance decay out-of-sample is small. To do this, we assess explanatory power during the unprecedented Covid-related market conditions that occurred from 2020 onwards. Figure 4 shows the performance remains comparable during the two periods.

[11] $\overline{\text{SSR}}$ denotes the benchmark overall mean, $\overline{\text{SSR}}_t = \frac{1}{N_t}\sum_{i=1}^{N}\text{SSR}_{i,t}$ is the cross-sectional mean of benchmark values (stated by respondents), and $\overline{\widehat{\text{SSR}}}_t = \frac{1}{N_t}\sum_{i=1}^{N}\widehat{\text{SSR}}_{i,t}$ is the cross-sectional mean of model-predicted values, where N_t is the number of samples at each time t.

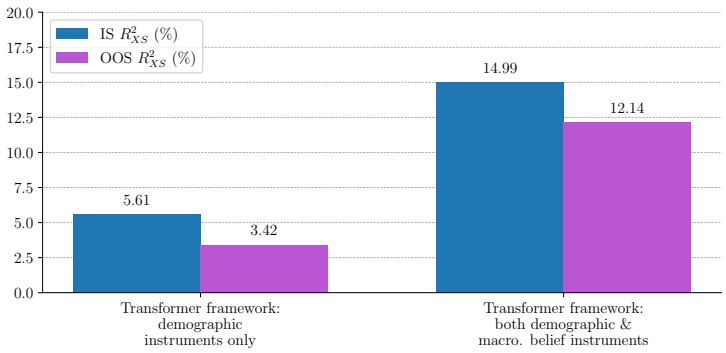


Figure 4: Out-of-Sample Explanatory Power

This figure compares the proportion of cross-sectional variance in respondents' subjective Sharpe Ratios that is explained during the in-sample (blue) and out-of-sample (purple) periods, for the two Transformer-based experience effects models.

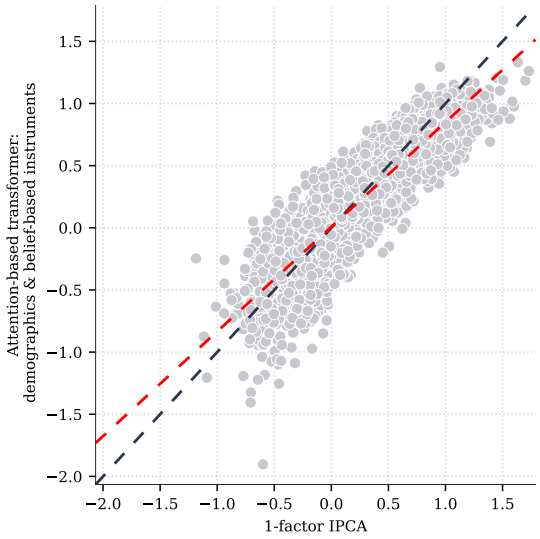


Figure 5: Predictions of the Latent 1-Factor Model vs. Transformer-based Experience Effects

The figure shows the predictions of the 1-factor IPCA, specified in Equation (11) with $\alpha_{i,t}$ restricted to zero, on the x-axis and the predictions of the experience effects estimated using demographics and macroeconomic beliefs as instruments in our machine learning model. We overlay the 45-degree line in black and the best linear fit in red.

3.5 Comparison to an IPCA benchmark

As a final validation check, we compare our main model's explained SSRs with the portion of SSRs that are explained by the state-of-the-art Instrumented Principal Components Analysis (IPCA) model, due to Kelly, Pruitt, and Su (2019). We describe the procedure for estimating the IPCA model and picking the optimal number of latent dynamic factors in Internet Appendix F. We also note that the IPCA model does not admit an economic interpretation in this particular setting, and should be considered a statistical reduced-form alternative. Yet its strong reduced-

form performance mean it can serve as a viable benchmark for our own approach.

Figure 5 shows a scatterplot of the predictions of the 1-factor IPCA model on the x-axis and the predictions of the transformer-based experience model on the y-axis. The agreement of the models is striking. We overlay the 45-degree line, which represents a perfect agreement between the two models' predictions in black and the best linear fit in red. The slope coefficient of this linear fit is 0.84. The intercept is close to zero and statistically insignificant. This close correspondence suggests that experience effects models can match the explanatory power of time-varying latent factor models – with the crucial advantage of capturing an interpretable economic mechanism.

4 Interpreting the Empirical Model

4.1 Heterogeneity

After estimating the Transformer framework with the richest specification (i.e., both macroeconomic belief and demographic instruments), we obtain a panel of experience weights for each household i surveyed on date t stretching back with various historical lags k. Figure 6 visualizes these experience weights, and highlights the substantial heterogeneity in households' models of the stock market.

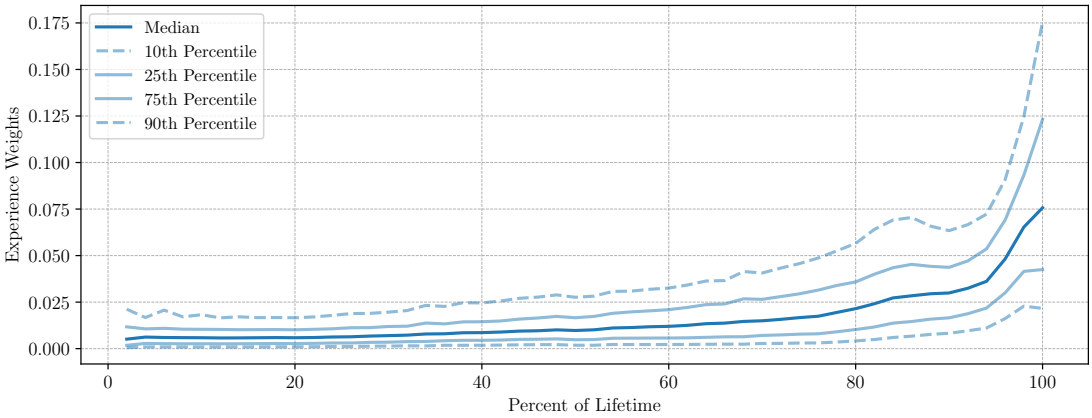


Figure 6: Heterogeneity in Estimated Experience Weights

This figure plots statistics of estimated experience weights generated by the Transformer framework with the richest specification. For the purposes of plotting, ages (which vary across households) are normalized to a common scale (x-axis), and experience weights are interpolated accordingly.

Recency Remarkably, the estimated experience weights in Figure 6 exhibit an exponential-like decay, although no explicit functional form was specified by our Transformer framework.

Furthermore, we note the most recent period of experience is consistently assigned the highest weight in the expectation formation process. Our estimates thus recover the well-known recency bias in return extrapolation (e.g., Greenwood and Shleifer (2014) and Cassella and Gulen (2018)).

Variation in Experience Weights We next conduct an exploration of the sources of heterogeneity in estimated experience weights. First, we analyze to what extent experience weights are time-invariant (while allowing weights to vary with age profiles, as in Malmendier and Nagel (2011), and as shown in Figure 6). To do so, we project our estimated experience weights onto interacted Lag $k \times \text{Age}_{i,t}$ fixed effects, and find that only 36.32% of the total variation in experience weights can be explained. Next, we allow for the possibility that experience weights vary with time (as shown by Cassella and Gulen (2018)), and thus project the estimated experience weights onto doubly-interacted Survey Date $t \times$ Lag $k \times \text{Age}_{i,t}$ fixed effects. These doubly-interacted fixed effects now explain 77.54% of the total variation, showing that our framework identifies time-variation in experience effects as well. We note that the remaining variation must be attributable to cross-sectional variation in experience effects. This analysis of the explanatory power of our framework highlights its ability to successfully identify multiple dimensions of heterogeneity.

4.2 Business Cycle

Given our focus on the perceived link between the stock market and macroeconomic conditions espoused by households, we turn immediately to interpreting cross-sectional and time-series variation in experience weights with an eye on business cycle variation. Our approach for doing so is simple: a panel regression framework that allows us to condition on and control for various household attributes (including a normalization for the length of experience; i.e., the age), and test for statistically significant differences in weights.

Influence of Past Recessions Table 3 shows that households place an outsize weight on past experienced recessions. The business cycle therefore plays a key role in the extrapolation process. Current conditions also matter: there is an incrementally significant additional weight placed on these past experienced recessions when the Michigan survey was run during a current recession period. Including time effects in the panel regression reveals the effect is cross-sectional and independent of the objective state of the macroeconomy.

We emphasize throughout this section that the Transformer framework is unaware of actual macroeconomic conditions (past, current, or future); only respondents' stated beliefs about

future macroeconomic conditions. The estimation merely attempts to best fit the data on SSRs, lifetime experienced realized Sharpe Ratios, and answers to Michigan survey questions about demographic characteristics and macroeconomic beliefs. The outsized influence of past realized recessions emerges independently as an effect.

Beliefs About the Macroeconomy Table 4 shows that the differential influence of recession periods is due to beliefs about the evolution of the macroeconomy. We calculate an index for the macroeconomic pessimism of each household, and include it as an interaction term:

$$\text{Macro Pessimism}_{i,t} = \frac{1}{3}(\text{Unemployment Direction}_{i,t} - \text{Business Conditions Direction}_{i,t} - \text{Interest Rates Direction}_{i,t}). \tag{17}$$

The positive loading on this interaction suggests a linkage between past realized Sharpe Returns and perceptions about the evolution of macroeconomic conditions. Including time effects in the panel regression reveals the effect is cross-sectional and independent of the objective state of the macroeconomy.

Similarity to Past Macroeconomic Conditions Next, we compute cosine similarity measures between a vector of each household's macroeconomic beliefs and a vector of each past year's change in equivalent observed macroeconomic variables, using multiple measures. We then repeat the exercise for the cosine similarity between the future conditions that are yet to realize and the same past realizations. Including both in a panel regression, we find in Table 5 that households are more likely to attend to and be influenced by returns that occurred when past macroeconomic conditions were most similar to what the household believes they will be. (Actual realizations do not have a positive predictive coefficient.) In a sense, households are attempting to predict future risk-adjusted returns by looking up past risk-adjusted returns that occurred in periods similar to those the household belies will be likely to ensure. Including time effects in the panel regression reveals the effect is cross-sectional and independent of the objective state of the macroeconomy.

Onset of a Recession Finally, we establish a link between our estimates and an argument – advanced by Gómez-Cram (2022) – that investors are slow to react to business cycle turning points. We conduct an event study on the 2008 recession to understand whether past returns that occurred in recession periods are overweighted to past returns that did not occur in recession periods. Our panel regression includes interacted lag-age fixed effects to control for mechanical weighting distortions, as well as baseline estimates for past returns that did not overlap with a recession period. Event study coefficient estimates (with double-clustered standard error-derived confidence intervals) are displayed in Figure 7. Interpreting the chart,

we observe a lagged response of experience effects, which is consistent with a lagged response of return expectations during a business cycle turning point.

Table 3: Experience Weights on Past Recession Periods

	Experience Weight$_{i,t,k}$ (%)			
	(1)	(2)	(3)	(4)
$\mathbb{1}\{\text{Experience}_{i,t,k}$ Overlaps NBER Recession$\}$	0.6670***	0.7540***	1.159***	1.155***
	(0.0519)	(0.0422)	(0.0456)	(0.0246)
$\mathbb{1}\{\text{Experience}_{i,t,k}$ Overlaps NBER Recession$\}$	0.9328***	0.9100***	0.6314***	0.6366***
$\times\ \mathbb{1}\{\text{Survey}_t$ During NBER Recession$\}$	(0.1950)	(0.0565)	(0.0380)	(0.0669)
Fixed-effects				
Survey date t	✓	✓	✓	✓
$\text{Age}_{i,t}$		✓	✓	
Lag k			✓	
Lag $k \times \text{Age}_{i,t}$				✓
N	3,804,951	3,804,951	3,804,951	3,804,951
R^2	0.01362	0.06144	0.38291	0.39283

Table 4: Experience Weights and Pessimism About the Macroeconomy

Dependent Variable:	Experience Weight$_{i,t,k}$ (%)			
	(1)	(2)	(3)	(4)
$\mathbb{1}${Experience$_{i,t,k}$ Overlaps NBER Recession}	0.9450***	1.029***	1.409***	1.407***
	(0.0567)	(0.0515)	(0.0529)	(0.0300)
Macro Pessimism$_{i,t}$	-0.2871***	-0.2999***	-0.2956***	-0.2964***
	(0.0142)	(0.0084)	(0.0072)	(0.0079)
$\mathbb{1}${Experience$_{i,t,k}$ Overlaps NBER Recession}	1.175***	1.168***	1.152***	1.155***
× Macro Pessimism$_{i,t}$	(0.0521)	(0.0440)	(0.0396)	(0.0240)
Fixed-effects				
Survey date t	✓	✓	✓	✓
Age$_{i,t}$		✓	✓	
Lag k			✓	
Lag k × Age$_{i,t}$				✓
N	3,689,522	3,689,522	3,689,522	3,689,522
R^2	0.01855	0.06654	0.39017	0.40007

Table 5: Experience Weights and the Similarity of Macroeconomic Conditions

	Experience Weight$_{i,t,k}$ (%)			
	(1)	(2)	(3)	(4)
Similarity of Beliefs to Past$_{i,t,k}$	0.2754***	0.2930***	0.3632***	0.3622***
	(0.0069)	(0.0075)	(0.0065)	(0.0070)
Similarity of Truth to Past$_{t,k}$	-0.3854***	-0.4586***	-0.1803***	-0.2161***
	(0.0129)	(0.0148)	(0.0109)	(0.0117)
Fixed-effects				
Lag k × Age$_{i,t}$	✓	✓	✓	✓
Survey Date t	✓	✓	✓	✓
N	2,472,733	2,242,704	2,557,088	2,353,802
R^2	0.34742	0.33999	0.35441	0.34774

Standard errors are clustered at the level of the fixed-effects shown.

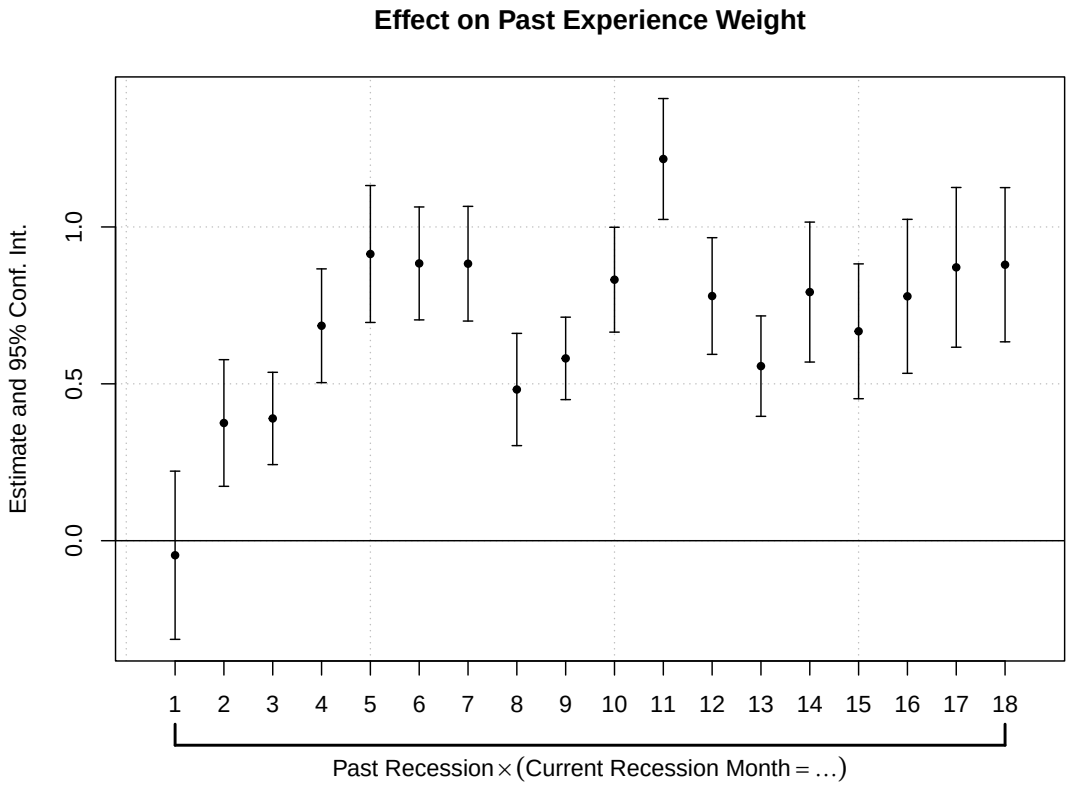


Figure 7: Weighting of Past Recessions During a Current Recession

This figure plots event study estimates for (significant) incremental experience weight assignments to past recession periods, evolving during the 2008 crisis recession period.

4.3 Local Thinking

Since households appear to be trying to predict the future return by looking up past returns, we finally examine to what extent such behavior is rational. Our test follows Meeuwis, Parker, Schoar, and Simester (2022) in exploiting the surprise Presidential election result in 2016, which influenced households' expectations. Table 6 Panel A shows that Democrat-affiliated respondents place a greater weight on past recession periods following the Republican candidate's victory, while Republican-affiliated respondents do the opposite. Interestingly, the effect is not driven out by the respondents' own stated macroeconomic beliefs, which suggests a role for heuristic, local thinking in this expectation formation process (Gennaioli and Shleifer, 2010).

As a further check, we repeat the test for the 2008 Presidential election in Table 6 Panel B, and find a similar result, albeit with less statistical significance. This is most likely due to the fact a much smaller subset of households were surveyed on their beliefs about the macroeconomy around those dates, leading to a loss in statistical power. We note also that economic and financial conditions were particularly volatile in 2008.

Table 6: Experience Weights and Heuristic Thinking

	Experience Weight$_{i,t,k}$ (%)			
Sub-sample:	Republicans		Democrats	
	(1)	(2)	(3)	(4)
Panel A: 2016 victory of President Trump				
$\mathbb{1}${Survey$_t$ POST Election}	-1.035**	-0.7000**	0.3024*	0.2843*
× $\mathbb{1}${Experience$_{t,k}$ Overlaps NBER Recession}	(0.0236)	(0.0110)	(0.0370)	(0.0257)
N	8,080	3,878	7,588	3,881
R^2	0.76215	0.83095	0.86477	0.90109
Panel B: 2008 victory of President Obama				
$\mathbb{1}${Survey$_t$ POST Election}	0.2478**	1.082	-0.2176*	-0.4072
× $\mathbb{1}${Experience$_{t,k}$ Overlaps NBER Recession}	(0.0131)	(0.2059)	(0.0299)	(0.1811)
N	7,926	2,227	7,224	1,866
R^2	0.83314	0.92525	0.88963	0.95920
Controls				
Baseline Effect	✓	✓	✓	✓
Demographic Characteristics	✓	✓	✓	✓
Macroeconomy Beliefs		✓		✓
Fixed-effects				
Lag k × Age$_{i,t}$	✓	✓	✓	✓
Survey Date t	✓	✓	✓	✓

Standard errors are clustered at the level of the fixed-effects shown.

5 Conclusion

This paper introduces a new source of subjective expectations data about the conditional Sharpe Ratio of the market portfolio. We contribute a set of stylized facts about investors' Subjective Sharpe Ratios, building upon recent work documenting the properties of other subjective expectations.

We then generalize the experience effects framework of Malmendier and Nagel (2011) to incorporate a link to perceived macroeconomic conditions. We find that past recessions exert an outsize influence on return expectations, that households appear to try to forecast returns based on their heterogeneous macro beliefs, but furthermore that other factors unrelated to the macroeconomy may play a role.

Our findings can inform the development of macro-finance asset pricing models.

References

Amromin, G., and Sharpe, S. A. (2014). From the horse's mouth: Economic conditions and investor expectations of risk and return. *Management Science*, *60*(4), 845–866.

Andre, P., Pizzinelli, C., Roth, C., and Wohlfart, J. (2022). Subjective models of the macroeconomy: Evidence from experts and representative samples. *Review of Economic Studies*, *89*(6), 2958–2991.

Andre, P., Schirmer, P., and Wohlfart, J. (2024). Mental models of the stock market. *Available at SSRN*, *4589777*.

Barberis, N., Greenwood, R., Jin, L., and Shleifer, A. (2015). X-capm: An extrapolative capital asset pricing model. *Journal of Financial Economics*, *115*(1), 1–24.

Beckmeyer, H., and Wiedemann, T. (2022). Recovering missing firm characteristics with attention-based machine learning. *Available at SSRN*, *4003455*.

Ben-David, I., Graham, J. R., and Harvey, C. R. (2013). Managerial miscalibration. *Quarterly Journal of Economics*, *128*(4), 1547–1584.

Bianchi, F., Ludvigson, S. C., and Ma, S. (2022). Belief distortions and macroeconomic fluctuations. *American Economic Review*, *112*(7), 2269–2315.

Boons, M., Ottonello, G., and Valkanov, R. I. (2024). What drives the volatility of professional stock return forecasts? causal evidence from macro shocks. *Available at SSRN*, *4537181*.

Bose, D., Cordes, H., Nolte, S., Schneider, J. C., and Camerer, C. F. (2022). Decision weights for experimental asset prices based on visual salience. *Review of Financial Studies*, *35*(11), 5094–5126.

Bybee, J. L. (2023). The ghost in the machine: Generating beliefs with large language models. *arXiv preprint arXiv:2305.02823*.

Cao, S., Jiang, W., Wang, J., and Yang, B. (2024). From man vs. machine to man + machine: The art and AI of stock analyses. *Journal of Financial Economics*, *160*, 103910.

Cassella, S., and Gulen, H. (2018). Extrapolation bias and the predictability of stock returns by price-scaled variables. *Review of Financial Studies*, *31*(11), 4345–4397.

Chen, Z., and Huang, J. (2023). Memory and beliefs in financial markets: A machine learning approach.

Cochrane, J. H. (2017). Macro-finance. *Review of Finance*, *21*(3), 945–985.

Cookson, J. A., and Niessner, M. (2020). Why don't we agree? Evidence from a social network of investors. *Journal of Finance*, *75*(1), 173–228.

Das, S., Kuhnen, C. M., and Nagel, S. (2020). Socioeconomic status and macroeconomic expectations. *Review of Financial Studies*, *33*(1), 395–432.

Dominitz, J., and Manski, C. F. (2007). Expected equity returns and portfolio choice: Evidence from the Health and Retirement Study. *Journal of the European Economic Association*, *5*(2-3), 369–379.

Dominitz, J., and Manski, C. F. (2011). Measuring and interpreting expectations of equity returns. *Journal of Applied Econometrics*, *26*(3), 352–370.

Fedyk, A., Kakhbod, A., Li, P., and Malmendier, U. (2024). ChatGPT and perception biases in investments: An experimental study. *Available at SSRN*, *4787249*.

Gennaioli, N., and Shleifer, A. (2010). What comes to mind. *Quarterly Journal of Economics*, *125*(4), 1399–1433.

Giglio, S., Maggiori, M., Stroebel, J., and Utkus, S. (2021). Five facts about beliefs and portfolios. *American Economic Review*, *111*(5), 1481–1522.

Goetzmann, W. N., Watanabe, A., and Watanabe, M. (2022). Cohort effects on expected comovement. *NBER Working Paper*, *29949*.

Gomes, F., Haliassos, M., and Ramadorai, T. (2021). Household finance. *Journal of Economic Literature*, *59*(3), 919–1000.

Gómez-Cram, R. (2022). Late to recessions: Stocks and the business cycle. *Journal of Finance*, *77*(2), 923–966.

Greenwood, R., and Shleifer, A. (2014). Expectations of returns and expected returns. *Review of Financial Studies*, *27*(3), 714–746.

Grossman, S. J., and Shiller, R. J. (1981). The determinants of the variability of stock market prices. *American Economic Review*, *71*(2), 222.

Hendrycks, D., and Gimpel, K. (2016). Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*.

Kelly, B. T., Pruitt, S., and Su, Y. (2019). Characteristics are covariances: A unified model of risk and return. *Journal of Financial Economics*, *134*(3), 501–524.

Lochstoer, L. A., and Muir, T. (2022). Volatility expectations and returns. *Journal of Finance*, *77*(2), 1055–1096.

Loshchilov, I., and Hutter, F. (2017). Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.

Malmendier, U., and Nagel, S. (2011). Depression babies: Do macroeconomic experiences affect risk taking? *Quarterly Journal of Economics*, *126*(1), 373–416.

Meeuwis, M., Parker, J. A., Schoar, A., and Simester, D. (2022). Belief disagreement and portfolio choice. *Journal of Finance*, *77*(6), 3191–3247.

Reddi, S. J., Kale, S., and Kumar, S. (2019). On the convergence of adam and beyond. *arXiv preprint arXiv:1904.09237*.

Sheikh, A. M. (1993). The behavior of volatility expectations and their effects on expected returns. *Journal of Business*, 93–116.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, *30*.

Von Gaudecker, H.-M., and Wogrolly, A. (2022). Heterogeneity in households' stock market beliefs: Levels, dynamics, and epistemic uncertainty. *Journal of Econometrics*, *231*(1), 232–247.

Appendix

In this appendix we derive Proposition 1. A more detailed discussion and verification of the assumptions, together with a characterization of the functional form, follows in in Internet Appendix A.

Preliminaries Recall that the error function is defined as

$$\operatorname{erf}(z) = \frac{2}{\sqrt{\pi}} \int_0^z e^{-t^2} dt, \tag{1}$$

and that the probit function is defined as the inverse of the cumulative distribution function $\Phi(z)$ of the standard (zero mean and unit variance) normal distribution:

$$\operatorname{probit}(p) = \Phi^{-1}(p) \quad \text{for} \quad p \in (0,1). \tag{2}$$

Proposition 1 (Subjective Sharpe Ratio). *Assume that an individual i perceives the market return to be conditionally normally distributed, $r_{t+1} \sim \mathcal{N}(\mu_{it}, \sigma_{it})$, and the riskless rate to be zero. Then*

$$SSR_{it} = \frac{\mu_{it}}{\sigma_{it}} \;=\; \operatorname{probit}\big(\mathbb{P}_{it}(r_{t+1} > 0)\big) \tag{2}$$

is her subjective perception of the market's conditional Sharpe Ratio over the next period.

Proof. To simplify the notation, we omit all subscripts i. We write the cumulative distribution function of the normal distribution

$$\mathbb{P}_t(r_{t+1} \le x) = \Phi\left(\frac{x-\mu_t}{\sigma_t}\right) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\frac{x-\mu_t}{\sigma_t}} e^{\frac{-t^2}{2}} dt \tag{3}$$

in terms of the error function erf as

$$\mathbb{P}_t(r_{t+1} \le x) = \frac{1}{2}\left[1 + \operatorname{erf}\left(\frac{x-\mu_t}{\sigma_t\sqrt{2}}\right)\right]. \tag{4}$$

For the value $x = 0$, the complement is therefore

$$\mathbb{P}_t(r_{t+1} > 0) = \frac{1}{2}\left[1 - \operatorname{erf}\left(\frac{-\mu_t}{\sigma_t\sqrt{2}}\right)\right]. \tag{5}$$

Exploit the property $\text{erf}(-z) = -\text{erf}(z)$ to write

$$\mathbb{P}_t(r_{t+1} > 0) = \frac{1}{2}\left[1 + \text{erf}\left(\frac{\mu_t}{\sigma_t\sqrt{2}}\right)\right] \tag{6}$$

and then rearrange as

$$\frac{\mu_t}{\sigma_t} = \sqrt{2}\,\text{erf}^{-1}\left(2\,\mathbb{P}_t(r_{t+1} > 0) - 1\right). \tag{7}$$

Finally, apply the property $\text{probit}(p) = \sqrt{2}\,\text{erf}^{-1}(2p-1)$ for probabilities p to obtain

$$\frac{\mu_t}{\sigma_t} = \text{probit}\left(\mathbb{P}_t(r_{t+1} > 0)\right). \tag{8}$$

Because the investor perceives the riskless rate to be zero, $\frac{\mu_t}{\sigma_t}$ is the Sharpe Ratio of the market portfolio, according to her subjective perception. □

Internet Appendix

A Mapping Survey Responses to Subjective Sharpe Ratios . . . 2
A.1 Probit Transformation . . . 2
A.2 Assumption of a Gaussian Return . . . 2
A.3 Assumption of a Zero Riskless Rate . . . 3
B Timeseries Variation in Subjective Sharpe Ratios . . . 6
B.1 Subjective Perceptions vs. Realizations . . . 6
B.2 A Rational Benchmark for the Conditional Sharpe Ratio . . . 8
C Cross-Sectional Variation in Subjective Sharpe Ratios . . . 11
D Subjective Sharpe Ratios Predict Household Demand . . . 13
D.1 Direct and Indirect Equity Holdings by Households . . . 13
D.2 Fund Flows . . . 13
E Estimating the Malmendier and Nagel (2011) Benchmark . . . 15
F Estimating the IPCA Benchmark . . . 17
G Overview of the Transformer Framework . . . 20
References in Internet Appendix . . . 22

A Mapping Survey Responses to Subjective Sharpe Ratios

A.1 Probit Transformation

Following directly from Proposition 1, one alternative interpretation of our Subjective Sharpe Ratio (SSR) measure is simply as a probit-transformed subjective probability of a stock market gain.

Figure IA.1 illustrates the probit transformation. While it is nonlinear, it is well-approximated by a linear relationship over most of the input domain.

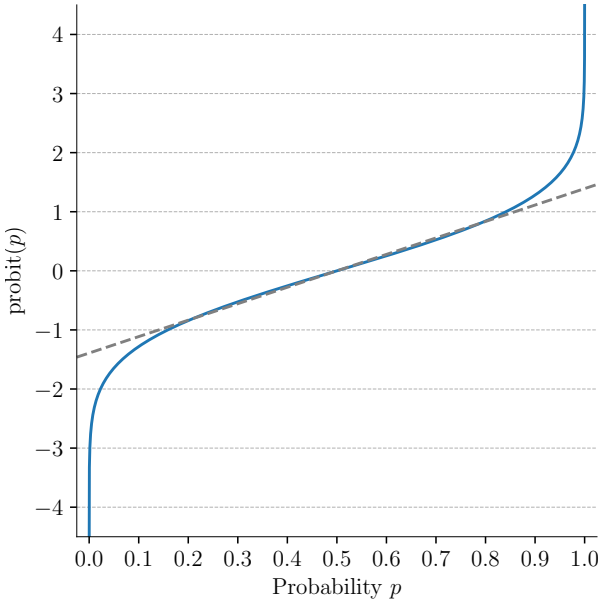


Figure IA.1: Probit transformation of example survey responses (x-axis) to subjective Sharpe Ratios (y-axis), following Proposition 1.

Another alternative interpretation of our results stems from the observation that the probit function is well-approximated by a linear mapping for the majority of survey responses, so for most responses the transformation will merely be a near-constant scale adjustment. This suggests our paper's main results can alternatively be interpreted as linking together experience effects and investors' subjective perceptions of the probabilities of future losses and gains in the stock market. Recent empirical evidence by Zeisberger (2022) suggests that the perceived loss probability is itself an important quantity.

A.2 Assumption of a Gaussian Return

The probit functional form follows from our assumption that investors perceive the market return to be conditionally Gaussian.

We build upon the existing subjective expectations literature by using the parametric assumption of normality to link survey responses to our economic quantity of interest. Notably, Dominitz and Manski (2007, pp. 376) and Lochstoer and Muir (2022, pp. 1060) make similar assumptions to measure the first and second moments (respectively) of investors' perceived distributions of the market return. Ben-David, Graham, and Harvey (2013, pp. 1553) also rely on this assumption.

We now directly test whether annual returns can be modeled as Gaussian. We collect annual returns for the S&P 500 index since 1925, and conduct three tests for Normality. Under each test, the Null hypothesis is that the returns are drawn from a Normal distribution, and the results are presented in Table IA.1. For each test, the Gaussian Null cannot be rejected.

Table IA.1: Hypothesis Test Results

Test	Statistic	p-value
Shapiro–Wilk	0.9809	0.1650
D'Agostino–Pearson	3.5625	0.1684
Kolmogorov–Smirnov	0.0618	0.8256

A.3 Assumption of a Zero Riskless Rate

The assumption of a zero riskless rate being available to (or perceived by) the investor is realistic in our setting of individual investors; for example, cash holdings in a retail current account facility insured by the FDIC are likely to provide a negligible return.

Table IA.2 shows that any measurement error introduced by this assumption will be small if the respondents perceive the riskless rate to be low – as interest rates were for much of the survey sample period. Recall that the Sharpe Ratio for a single asset is defined as

$$\frac{\mu_t}{\sigma_t} - \frac{r_f}{\sigma_t}, \tag{9}$$

where μ_t and σ_t are the perceived first and second moments of the asset's return distribution, and r_f is the perceived riskless rate. Our assumption in deriving Proposition 1 implies that the second term r_f/σ_t is zero. To understand the consequences of an agent who perceives a non-zero riskless rate, Table IA.2 shows possible values of the error term for a grid of r_f & σ_t values ranging up to extreme historical values during recent times. It is immediately apparent that low perceived riskless rate values result in a low magnitude of the possible error, irrespective of the (unobserved) perceived level of volatility. Given the low-interest rate regime that took

place throughout our sample period, we believe the assumption of a perceived zero riskless rate is a reasonable one.

Figure IA.2 plots cross-sectional means of our recovered SSR values, and additionally includes an explicit adjustment for the wedge $\frac{r_f}{\sigma_t}$, which we assume is constant for all respondents in a given survey year-month; we estimate this adjustment using 10-year Treasury yields and realized S&P 500 return volatility over the past year. Therefore, even *if* households were able to access the true riskless rate, the magnitude of any wedge with our estimates of households' SSRs would be minimal for most of our sample period.

Table IA.2: Relaxing the Zero Riskless Rate Assumption

Possible values of r_f/σ_t for a grid of r_f columns and σ_t rows.

		r_f (%)										
		0	0.5	1	1.5	2	2.5	3	3.5	4	4.5	5
σ_t	10	0.00	0.05	0.10	0.15	0.20	0.25	0.30	0.35	0.40	0.45	0.50
(%)	20	0.00	0.03	0.05	0.08	0.10	0.13	0.15	0.18	0.20	0.23	0.25
	30	0.00	0.02	0.03	0.05	0.07	0.08	0.10	0.12	0.13	0.15	0.17
	40	0.00	0.01	0.03	0.04	0.05	0.06	0.08	0.09	0.10	0.11	0.13
	50	0.00	0.01	0.02	0.03	0.04	0.05	0.06	0.07	0.08	0.09	0.10
	60	0.00	0.01	0.02	0.03	0.03	0.04	0.05	0.06	0.07	0.08	0.08
	70	0.00	0.01	0.01	0.02	0.03	0.04	0.04	0.05	0.06	0.06	0.07
	80	0.00	0.01	0.01	0.02	0.03	0.03	0.04	0.04	0.05	0.06	0.06
	90	0.00	0.01	0.01	0.02	0.02	0.03	0.03	0.04	0.04	0.05	0.06
	100	0.00	0.01	0.01	0.02	0.02	0.03	0.03	0.04	0.04	0.05	0.05

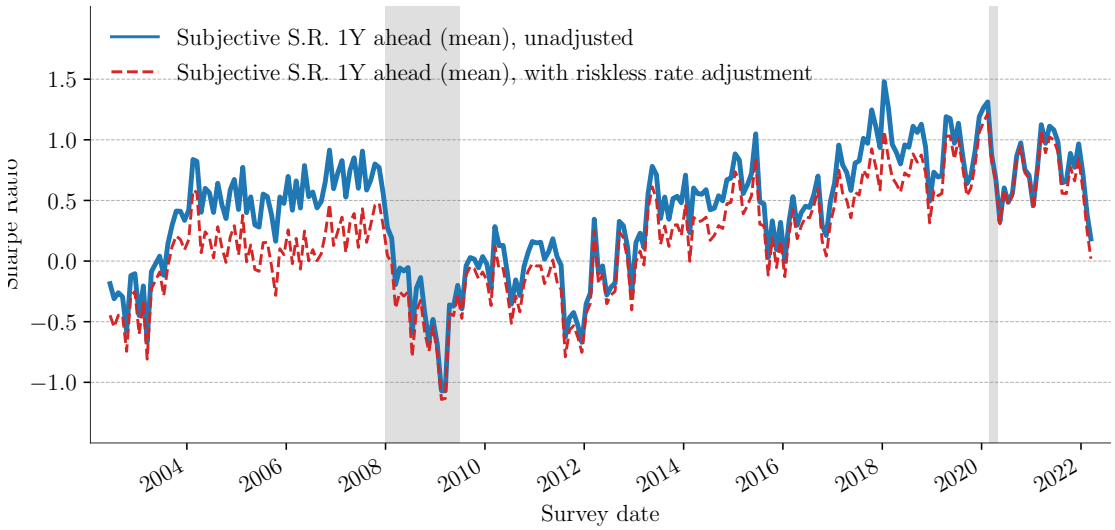


Figure IA.2: Mean Survey Subjective Sharpe Ratio With Riskless Rate Adjustment

Primary series (solid blue) shows the cross-sectional mean of individual $\frac{\mu_t}{\sigma_t}$ values, which we recovered from Michigan Survey responses following Proposition 1. Secondary series (dashed red) is our augmented estimate for $\frac{\mu_t}{\sigma_t} - \frac{r_f}{\sigma_t}$ using realizations. Shaded areas are NBER recession periods.

B Timeseries Variation in Subjective Sharpe Ratios

B.1 Subjective Perceptions vs. Realizations

How do Subjective Sharpe Ratios (SSRs) relate to realized Sharpe Ratios? We now compare the timeseries of cross-sectional mean SSRs with the timeseries of S&P 500 Sharpe Ratio realizations.

Figure IA.3 shows that mean SSRs diverge from the future realizations, over the same 1-year horizon. In this sense, SSRs are poor predictors, on average, of the future realization.

Figure IA.4 points to one reason why. Overlaying the same series of mean SSRs with *past* 1-year realizations of the S&P 500 Sharpe Ratio, the cyclicalities now match. That is, expectations about the future depend (positively) on recent past realizations; the magnitude of the correlation is 0.11, which is significant at the 1% level.

On the one hand, this positive correlation is consistent with the survey evidence of Greenwood and Shleifer (2014) that expectations (of returns) are extrapolative from the recent past; on the other hand, a large part of the variation remains unexplained, suggesting a need for incorporating investors' more distant past experiences into their expectation formation function.

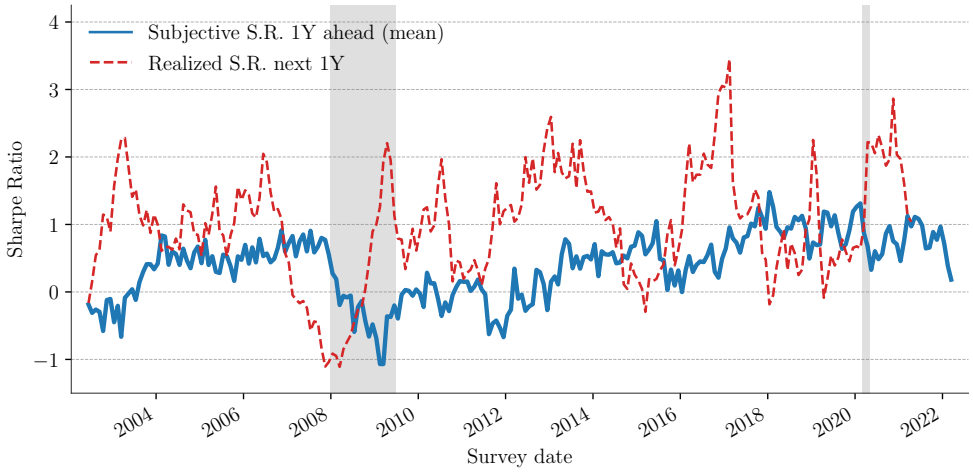


Figure IA.3: Mean Survey Subjective Sharpe Ratio for Next Year vs. Actual Future Realization

Primary series (solid blue) is the cross-sectional mean of probit-transformed responses on the probability of an up-move in the stock market over the next twelve months; this measures the mean subjective Sharpe Ratio perceived by survey respondents, according to Proposition 1. Secondary series (dashed red) is the actual, realized future twelve-month ahead Sharpe Ratio at each survey date. Shaded areas are NBER recession periods.

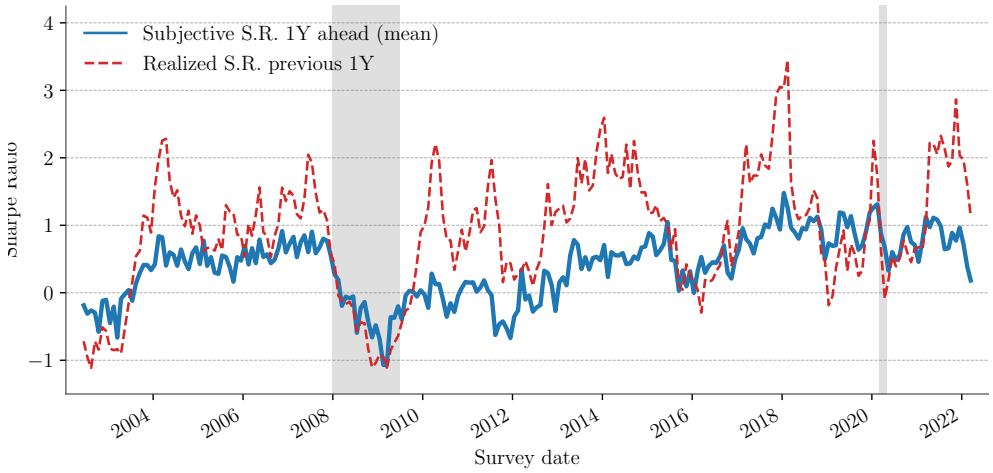


Figure IA.4: Mean Survey Subjective Sharpe Ratios for Year Ahead vs. Previous Year Realizations

Primary series (solid blue) is the cross-sectional mean of probit-transformed responses on the probability of an up-move in the stock market over the next twelve months; this measures the mean subjective Sharpe Ratio perceived by survey respondents, according to Proposition 1. Secondary series (dashed red) is the realized rolling twelve-month Sharpe Ratio at each survey date. Shaded areas are NBER recession periods.

B.2 A Rational Benchmark for the Conditional Sharpe Ratio

This appendix highlights the contrast between the procyclical nature of elicited subjective Sharpe Ratios and the countercyclical nature of a rational benchmark. We first explain the construction of the rational benchmark, which is based on a classic consumption-based asset pricing model, and then present our comparison.

Campbell and Cochrane (1999) Theoretical Benchmark We justify our rational benchmark using the habit formation model of Campbell and Cochrane (1999). A time-varying Sharpe Ratio naturally arises as part of their main mechanism because the conditional mean and standard deviation of the (market) asset's return are both functions of the dynamic model's state variable, the surplus consumption ratio. Furthermore, this conditional Sharpe Ratio is countercyclical: *"In our model, recessions are times of low consumption relative to habit, low prices, somewhat higher standard deviations of returns, very much higher expected returns, and correspondingly high Sharpe Ratios."* (Campbell and Cochrane, 1999, pp. 224)

The asset's P/D ratio is also a function of the model's state variable. Therefore, a key advantage of the Campbell and Cochrane (1999) model for our purposes is that the P/D ratio and the conditional Sharpe Ratio are tightly linked, via the model's state variable. To illustrate this point, we first numerically solve the model using their suggested parameters, and we then characterize the relationship that the solved model generates between the P/D ratio and the Sharpe Ratio of the dividend claim. Figure IA.5 plots the relationship between the P/D ratio and the Sharpe Ratio:[12] according to Campbell and Cochrane (1999)'s model, the Sharpe Ratio is negatively correlated with the P/D ratio, and the relationship is close to linear.

Empirical Benchmark Based on P/D Predictability Since the rational Campbell and Cochrane (1999) model implies a negative and (almost) linear relationship between the P/D ratio and the Sharpe Ratio, we linearly project the S&P 500 P/D ratio onto the year-ahead realized Sharpe Ratio, using annual data from 1970 onwards. Empirically, the two timeseries are indeed negatively correlated.

This projected Sharpe Ratio is our empirical benchmark for the conditional Sharpe Ratio. Figure IA.6 plots this rational benchmark Sharpe Ratio over time, and shows it is countercyclical: the benchmark Sharpe Ratio increases during recessions, and falls during good times.

Comparison to the Benchmark Figure IA.7 displays investors' subjective Sharpe Ratios, which we recover using Proposition 1 and then plot (in blue) as cross-sectional means. This

[12] These results are presented separately in Figures 3 and 6 of the Campbell and Cochrane (1999) paper. Our Figure IA.5 combines the two.

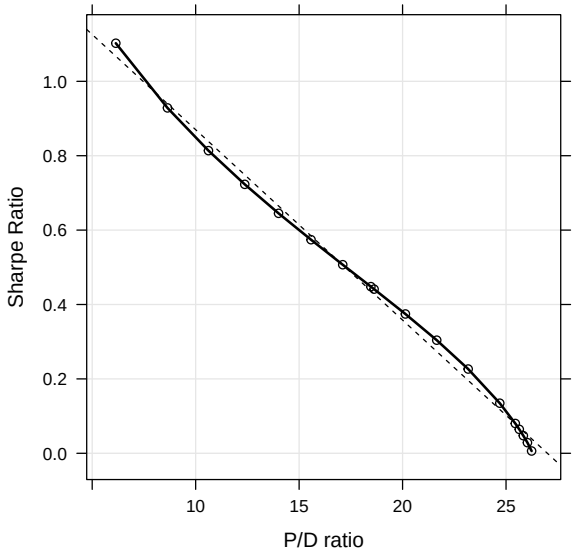


Figure IA.5: Rational model-implied relationship between the Sharpe Ratio and P/D ratio

This figure depicts the relationship between the Sharpe Ratio and P/D ratio of the dividend claim in Campbell and Cochrane (1999)'s numerically solved asset pricing model. The dashed gray line is a linear approximation.

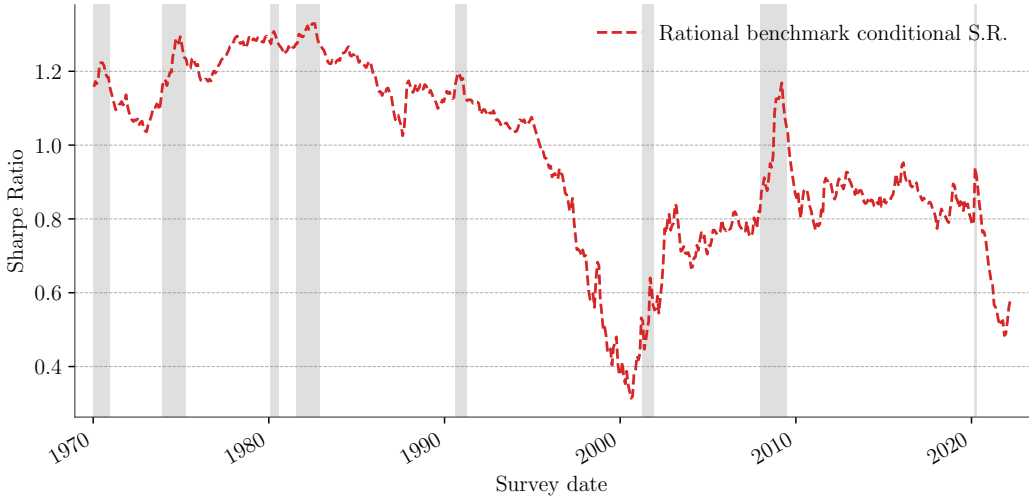


Figure IA.6: Empirical benchmark for the conditional Sharpe Ratio

This empirical benchmark is implied by a rational model in which the P/D ratio negatively predicts the Sharpe Ratio, such as the habit formation model of Campbell and Cochrane (1999). Shaded areas are NBER recession periods.

aggregate time series is procyclical, with a fall in the average subjective Sharpe Ratio during the financial crisis in 2008 and the Covid crash of 2020; the procyclical nature of this expectations series is consistent with the survey evidence of Amromin and Sharpe (2014) and Greenwood and Shleifer (2014), who document similar procyclical patterns in subjective expectations collected

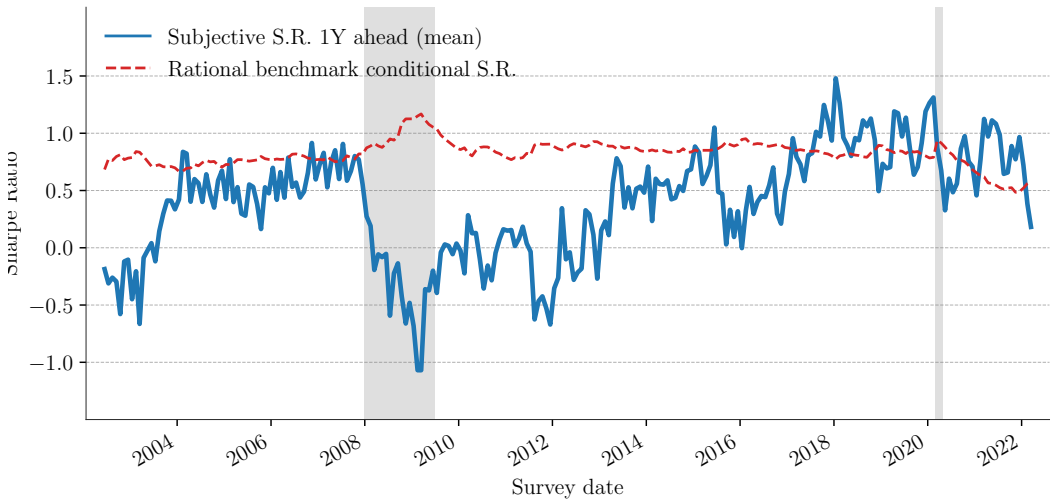


Figure IA.7: Mean Survey Subjective Sharpe Ratios for Year Ahead vs. a Rational Benchmark

Primary series (solid blue) is the cross-sectional mean of probit-transformed responses on the probability of an up-move in the stock market over the next twelve months; this measures the mean subjective Sharpe Ratio perceived by survey respondents, according to Proposition 1. Secondary series (dashed red) is a rational benchmark for the conditional Sharpe Ratio: its projection onto the market P/D ratio, as justified by the asset pricing model of Campbell and Cochrane (1999). Shaded areas are NBER recession periods.

from other surveys.

In comparison, the time-varying Sharpe Ratios produced by our benchmark rational asset pricing model (Campbell and Cochrane, 1999) is countercyclical, in order to match the equity premium. Figure IA.7 also overlays (in red) the conditional Sharpe Ratios implied by Campbell and Cochrane (1999)'s benchmark asset pricing model, which by contrast increase in bad times and fall in good times.

C Cross-Sectional Variation in Subjective Sharpe Ratios

The prior literature finds that demographic attributes have explanatory power for the cross-section of subjective return expectations (Dominitz and Manski, 2007; Das, Kuhnen, and Nagel, 2020). In this appendix, we document that the same is true in our setting.

Figure IA.8 plots various timeseries of the cross-sectional means of respondents' subjective Sharpe Ratios. Younger investors, male investors, and investors who support the current President's party all expect higher subjective Sharpe Ratios for the market in the year ahead, when compared to older investors, female investors, and those who oppose the current President's party.[13] Moreover, their expectations align with their investments in the stock market: investors with above-median investments in stocks expect a higher Sharpe Ratio than those with below-median investments in stocks.

The empirical framework that we introduce in Section 3 explicitly allows a role for demographic characteristics to shape households' response to past experience.

[13]The alignment of subjective beliefs to political affiliation is consistent with the findings of Kamdar and Ray (2022) for other expectations series.

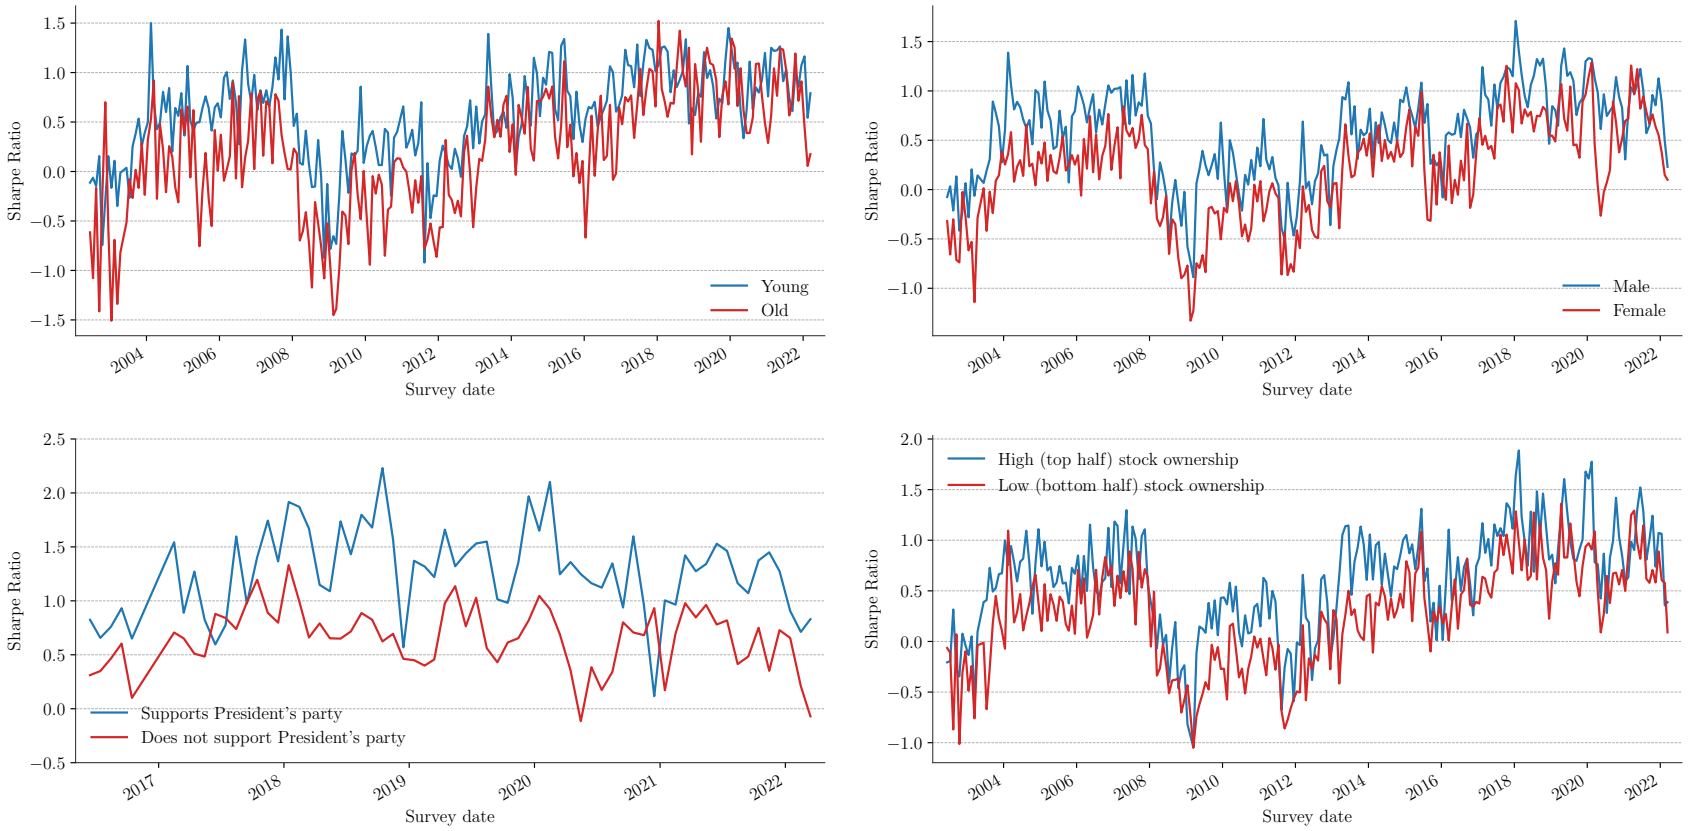


Figure IA.8: Mean Survey Subjective Sharpe Ratio Conditional on Various Attributes of Respondents

Each chart shows mean subjective Sharpe Ratios (SSR) conditional on one attribute of the survey respondents, who are split by that attribute (blue vs. red lines). The upper-left chart conditions on age: younger respondents ($\leq$ 40 years old) on average report a higher SSR than older respondents ($>$ 60 years old). The upper-right chart conditions on gender: male respondents on average report a higher SSR than female respondents. The bottom-left chart conditions on political affiliation: respondents who support the current President's party on average report a higher SSR than those who do not. The bottom-right chart conditions on the amount invested in the stock market: respondents in the top half of investors by size (calculated cross-sectionally) on average report a higher SSR than respondents in the bottom half.

D Subjective Sharpe Ratios Predict Household Demand

This section shows that investors' subjective Sharpe Ratios determine aggregate holdings in equities.

D.1 Direct and Indirect Equity Holdings by Households

An investor who chooses between a risky market portfolio and a riskless asset in a classic portfolio choice model (such as mean-variance optimization) will in theory hold the market portfolio in direct proportion to its perceived Sharpe Ratio, as shown in Equation (4). We now confirm that this relationship holds empirically at the aggregate level.

We are able to measure wealth shares held in equities – the fraction x in Equation (4) – at an aggregate level, thanks to US national accounts data.[14] Figure IA.9 compares the timeseries variation in the aggregate household's equity share with the timeseries variation in the subjective Sharpe Ratio perceived by this representative investor. The two series covary strongly: the correlation coefficient of the two series is 0.66 (and is significant at the 1%-level). Furthermore, the effect is not driven by a mechanical increase in the market's value value over time: in a regression framework (unreported), the positive association survives controlling for the level of the S&P 500, and is significant even when standard errors are clustered by year-month.

D.2 Fund Flows

Another consequence of Equation (4), which relates holdings to expectations, is that *changes* to the Subjective Sharpe Ratio should be positively correlated with changes to this holding; i.e. to an investor's net *flow* into the stock market. Under the assumption that households' flows exclusively target funds, we confirm that this is indeed the case at the level of a representative investor and the aggregate market portfolio.

We measure the representative investor's Subjective Sharpe Ratio by the mean of individual investors' values at each point in time. Aggregate flows into the stock market are calculated

[14] The results shown in this section make use of the FRED series "Households and Nonprofit Organizations; Directly and Indirectly Held Corporate Equities as a Percentage of Financial Assets." The results are not sensitive to using the series "Households and Nonprofit Organizations; Directly and Indirectly Held Corporate Equities as a Percentage of Total Assets" instead. Indeed, the magnitude of the correlation between those two series is 0.99.

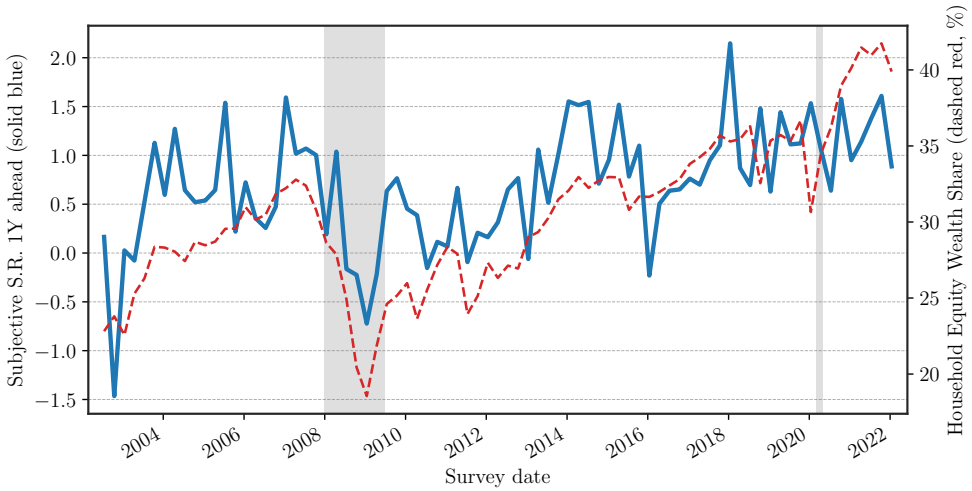


Figure IA.9: Mean Survey Subjective Sharpe Ratios for Year Ahead vs. Households' Share of Wealth in Equities

Primary series (solid blue) is the cross-sectional mean of probit-transformed responses on the probability of an up-move in the stock market over the next twelve months; this measures the mean subjective Sharpe Ratio perceived by survey respondents, according to Proposition 1. This aggregate mean is weighted by individuals' dollar amounts invested amount in the stock market. Secondary series (dashed red) is the (quarterly) share of households' wealth held in equities, both directly and indirectly, as published by FRED. Shaded areas are NBER recession periods.

using Investment Company Institute data on fund flows:[15]

$$\text{Flows}_{t\to t+1} = \frac{\$\text{ Domestic Equity Fund Flows}_{t\to t+1}}{\$\text{ Domestic Equity Market Capitalization}_{t}} \times 100\%. \tag{10}$$

As expected, we find that monthly flows ($\text{Flows}_{t\to t+1}$) are positively correlated with month-to-month differences in the aggregate mean Subjective Sharpe Ratio. The magnitude of the correlation is 0.26 and it is significantly non-zero at the 1%-level. This finding suggests that the aggregated subjective Sharpe Ratio of individual investors has explanatory power for aggregate flows, and therefore contributes to Gabaix and Koijen (2021)'s call for understanding the determinants of flows into the stock market.

[15]Monthly aggregate fund flows are calculated by the Investment Company Institute, and are publicly available from 2007 onward as their "Long-Term Mutual Fund and Exchange-Traded Fund (ETF) Flows" dataset. We use the Domestic Equity category. The total market capitalization is calculated from the CRSP database at a monthly frequency.

E Estimating the Malmendier and Nagel (2011) Benchmark

This appendix provides further details on how the adaptive weighting model in Section 3.1 was estimated.

We estimate the adaptive expectation formation model given by Equations (5) and (6) on our dataset of Subjective Sharpe Ratios, choosing λ so that subjects' past experiences best explain their current expectations. Subjects' experienced Sharpe Ratios SR_{t-k} are specified as the lagged annual S&P 500 Sharpe Ratios that were realized during their lifetimes.

The goal of the optimization is to pick the value of the decay parameter λ in Eqn. (6). We specify an objective of minimizing the Mean Squared Error between the values predicted by the exponential weighting procedure, in Eqn. (5), as compared to the subjective Sharpe Ratios recovered from subjects' responses that are recorded in our Michigan dataset, as in Eqn. (2).

We run a grid search, similar to Malmendier and Nagel (2011). Over the range $\lambda \in [-2, +20]$, we find the optimal value to be $\widehat{\lambda} = 10.39$. Figure IA.10 visualizes the objective function minimized over the grid search procedure.

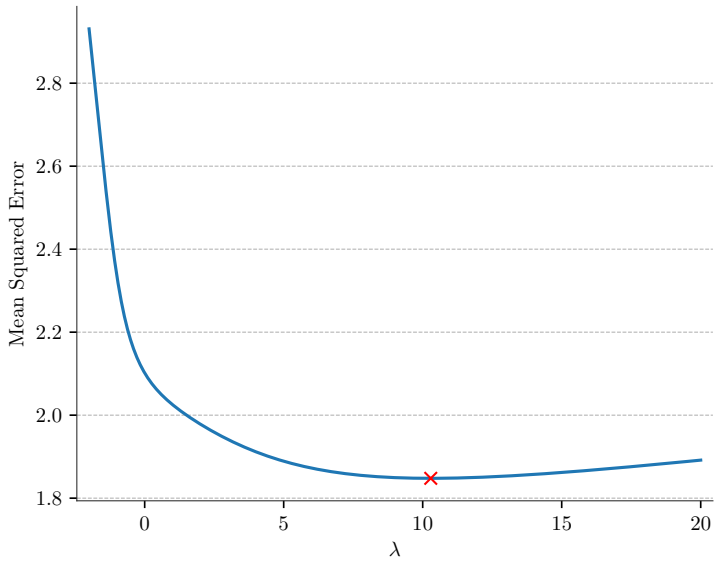


Figure IA.10: Objective function during the grid search procedure for estimating the λ parameter of the Malmendier and Nagel (2011) exponential weighting function.

We estimate $\widehat{\lambda} = 10.39$, suggesting that subjects place a relatively high weight on their recent past experiences (compared to their early lifetime experiences) when forming expectations about the future Sharpe Ratio of the stock market. The experience weights that correspond to this value of λ are illustrated in Figure IA.11.

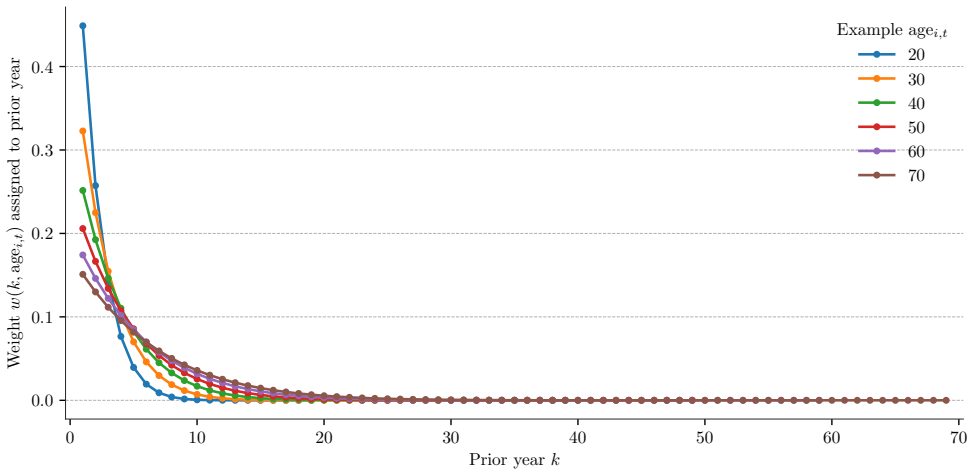


Figure IA.11: Weights $w(k, \text{age}_{i,t})$ assigned by the Malmendier and Nagel (2011) experience effects model to prior years k for various example ages of respondents. The optimal value $\widehat{\lambda} = 10.39$ is used.

F Estimating the IPCA Benchmark

We follow recent advances from the asset pricing literature and apply the *instrumented principal component analysis (IPCA)* technique of Kelly, Pruitt, and Su (2020) to subjective Sharpe Ratios.

IPCA instruments each household i's time-varying exposure $\beta_{i,t}$ to common factors f_{t+1} using the household's characteristics, i.e., their observable demographic attributes and beliefs about macroeconomic state variables:

$$\text{SSR}_{i,t} = \alpha_{i,t} + \beta_{i,t} f_{t+1} + \varepsilon_{i,t+1}, \qquad \beta_{i,t} = [\mathbf{d}, \mathbf{b}]'_{i,t} \mathbf{\Gamma}^{AC}_{\beta}; \quad \alpha_{i,t} = [\mathbf{d}, \mathbf{b}]'_{i,t} \mathbf{\Gamma}^{AC}_{\alpha} \tag{11}$$

Since individual households are identified by their characteristics, which we can track through time, we can consequently estimate a common factor structure. Besides the common factors in f_{t+1} we also include a time-varying and household-specific intercept ($\alpha_{i,t}$) in Equation (11), which too is also instrumented through household characteristics. This way, household characteristics can influence subjective Sharpe Ratios above and beyond their impact through f_{t+1}. A secondary benefit of this setup is that we can test for the significance of $\alpha_{i,t}$ and select the appropriate number of common factors to drive out the explanatory power of household characteristics that is not yet picked up by these factors. This significance test is done through a bootstrap analysis, the details of which are provided in Kelly, Pruitt, and Su (2019).

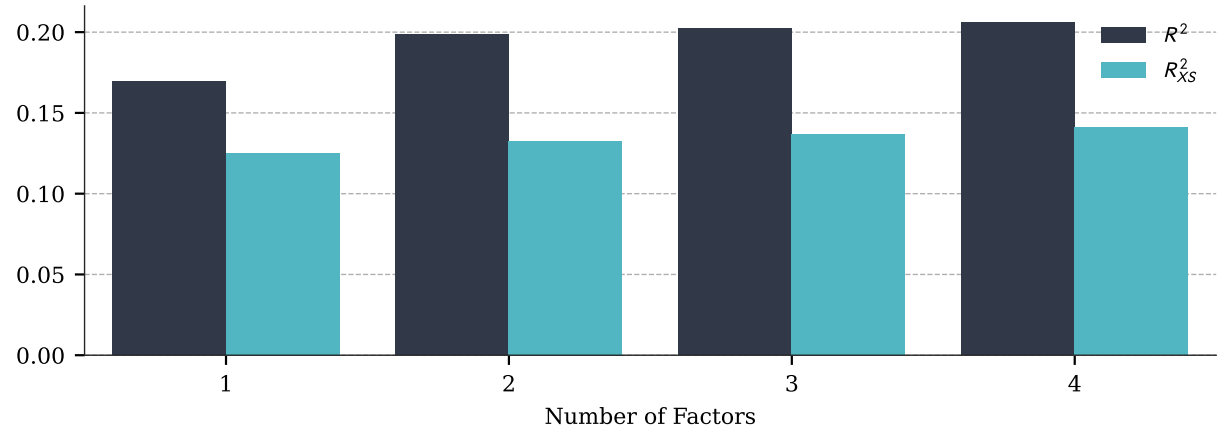


Figure IA.12: R^2 of IPCA with 1 to 4 Factors

The figure shows the R^2 and R^2_{XS} of IPCA models with 1 to 4 factors, following the specification in Equation (11) with $\alpha_{i,t} = 0$.

When estimating Equation (11), we rank household characteristics cross-sectionally, and standardize them to the range $[-0.5, 0.5]$. For the education code variable we assign larger ranks for increasing education levels. We report the classic fraction of total variance R^2 explained by a $K \in [1, 4]$-factor IPCA model, together with the cross-sectional R^2_{XS} measure of

Table IA.3: $\mathbf{\Gamma}_\beta$ Estimates for the IPCA Factor Model

This table shows the $\mathbf{\Gamma}_\beta$ estimates for a one-factor IPCA model following Equation (11) with $\alpha_{i,t} = 0$. The input characteristics are standardized to the range $[-0.5, 0.5]$, making the coefficients directly comparable.

	Coefficient
Panel A: Demographics	
(Constant)	0.170
Age	−0.115
Education level	0.187
Is female	−0.147
Stock holdings value	0.310
Home value	−0.032
Household income	0.018
Supports current President's party	0.009
Panel B: Beliefs About the Macroeconomy	
Forecast probability of income increase	0.742
Forecast direction of business conditions	0.366
Forecast direction of interest rates next year	0.051
Forecast direction of unemployment	−0.312
Forecast 1Y inflation rate	−0.110
Forecast 5Y inflation rate	−0.067
Forecast house prices change	−0.045

Eqn. (16). Figure IA.12 shows that a single factor model explains roughly 17% of the variation in subjective Sharpe Ratios across households. This number increases to 20% for a two-factor model and afterwards saturates at around 20-21% for three to four factors. The cross-sectional R^2_{XS} shows a similar picture: a single-factor model explains roughly 12.5% of the cross-sectional variation in subjective beliefs, which increases slightly as we increase the number of factors. When we test for the significance of the household-specific intercept $\alpha_{i,t}$, we find it to already be insignificant for a single factor. A one-factor model is sufficient to drive out the explanatory power of household demographics and macroeconomic beliefs that is not already captured by a household's exposure to the latent systematic factor.

Following this insight, Table IA.3 presents estimates for the $\mathbf{\Gamma}_\beta$ vector of Equation (11) for the one-factor model. Since all input characteristics are standardized between −0.5 and 0.5, the coefficients are directly comparable. Overall, we find that households' other beliefs have the largest influence on a household's factor exposure. We find the largest coefficients for the stated probability of an increase in personal income over the next year of 0.742. The next-largest coefficients are for the household's forecast of business conditions (0.366) and the direction of the unemployment rate (−0.312). Demographics are also important drivers of a household's exposure: a household's exposure is increasing in its stock holdings (0.310) and education level (0.187), decreasing in age (−0.115) and is lower for women (−0.147).

Interestingly, despite the support for the current president's party driving a marked wedge between household expectations of future Sharpe Ratios (Figure IA.8), it has little influence on the household's factor exposure.

The described IPCA analysis shows that subjective Sharpe Ratios follow a low-dimensional factor structure across households. Both the demographics and other beliefs of a household are important in instrumenting its exposure to this factor structure. As with any latent specification, we lack a clear understanding of what these factors represent. In the main paper, we compare the predictive power of this 1-factor IPCA model with our economically interpretable experience effects model.

G Overview of the Transformer Framework

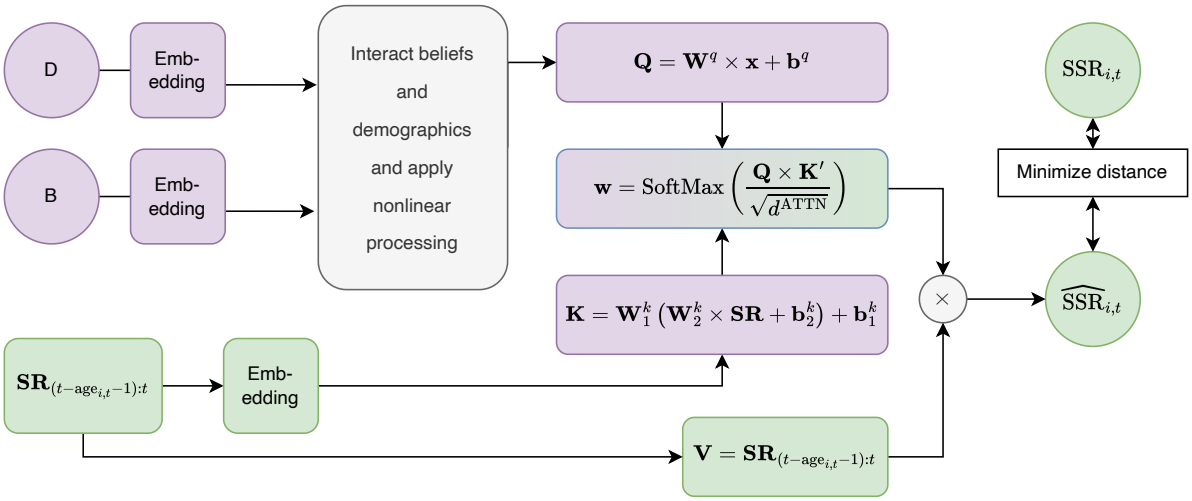


Figure IA.13: Model Setup

The figure shows how our model processes information and instruments variation in experience effects across households and over time using their beliefs and demographics. Purple-shaded parts use information about a respondent's beliefs and demographics, whereas green-shaded parts relate to past Sharpe Ratios.

Embedding As the first modeling step, we feed demographics, beliefs and past Sharpe Ratios through individual embeddings. For continuous variables, embeddings simply amount to a linear transformation:

$$\mathbf{x}^{\text{emb}} = \mathbf{x} \times \mathbf{W}, \tag{12}$$

where $\mathbf{W}^x \in \mathbb{R}^{1 \times d^{\text{ATTN}}}$. Here, d^{ATTN} is called the "embedding dimension", which governs the model's internal capacity. The logic behind expressing categorical variables is similar. For example, a respondent's age will be represented by a d^{ATTN}-dimensional numeric vector, which is learned in the estimation step. That way, the model has the flexibility of applying similar transformations to similar age groups. The embeddings are global and do not depend on month t or respondent i.

Interacting Macroeconomic Beliefs and Demographics Next, we extract interaction effects between and within demographic variables and beliefs about macroeconomic state variables, and apply a layer of nonlinear processing. Specifically, we first use the attention mechanism described in Equation (11), wherein keys, queries and values are all functions of the same inputs $\mathbf{d}_{i,t}$ and $\mathbf{b}_{i,t}$, giving rise to the name "self-attention." Let $\mathbf{d}_{i,t}$ be a N^{dem}-dimensional and $\mathbf{b}_{i,t}$ a N^{bel}-dimensional vector. The attention mechanism retrieves N^{total} weighted averages

of the stacked inputs $\mathbf{d}$ and $\mathbf{b}$, where $N^{\text{total}} = N^{\text{dem}} + N^{\text{bel}}$. These weighted averages are *context-specific*, in that they may differ across households and over time. To also incorporate possible nonlinearities, we next feed these weighted averages through a linear layer, followed by a nonlinear activation function and another linear layer. This procedure is standard in the Transformer literature. We use Gaussian Error Linear Units (GELU) as the nonlinear activation function, which generally provides better results in Transformers (Hendrycks and Gimpel, 2016). We found that repeating this self-attention step for demographics and beliefs a few times yields a more stable fit and weights that are easier to interpret. We therefore repeat this step three times. The resulting interaction effects are denoted by $\widehat{\mathbf{d}}_{i,t}$ and $\widehat{\mathbf{b}}_{i,t}$.

Extracting Experience Weights and Predict Stock Market Expectations With the extracted information about a household's demographics and beliefs about macroeconomic state variables to hand, we extract experience weights from experienced returns. The details of this are discussed in Section 3. In short, we compute keys $\mathbf{K}$ as a nonlinear transformation of embedded past Sharpe Ratios, queries $\mathbf{Q}$ as a linear transformation of the interacted macroeconomic beliefs and demographics, and compare the two matrices using their dot-product similarity to extract experience weights $\mathbf{w}$, which we multiply with the raw, realized Sharpe Ratios, to come up with the model's prediction for respondent i's subjective Sharpe Ratio in month t. We fit the model by minimizing the mean squared prediction error using stochastic gradient descent. The model takes between 100–150 epochs to converge and is robust to alterations of the hyperparameters.

References in Internet Appendix

Amromin, G., and Sharpe, S. A. (2014). From the horse's mouth: Economic conditions and investor expectations of risk and return. *Management Science*, *60*(4), 845–866.

Ben-David, I., Graham, J. R., and Harvey, C. R. (2013). Managerial miscalibration. *Quarterly Journal of Economics*, *128*(4), 1547–1584.

Campbell, J. Y., and Cochrane, J. H. (1999). By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy*, *107*(2), 205–251.

Das, S., Kuhnen, C. M., and Nagel, S. (2020). Socioeconomic status and macroeconomic expectations. *Review of Financial Studies*, *33*(1), 395–432.

Dominitz, J., and Manski, C. F. (2007). Expected equity returns and portfolio choice: Evidence from the Health and Retirement Study. *Journal of the European Economic Association*, *5*(2-3), 369–379.

Gabaix, X., and Koijen, R. S. (2021). In search of the origins of financial fluctuations: The inelastic markets hypothesis. *NBER Working Paper*, *28967*.

Greenwood, R., and Shleifer, A. (2014). Expectations of returns and expected returns. *Review of Financial Studies*, *27*(3), 714–746.

Hendrycks, D., and Gimpel, K. (2016). Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*.

Kamdar, R., and Ray, W. (2022). Polarized expectations, polarized consumption. *Available at SSRN 4251955*.

Kelly, B. T., Pruitt, S., and Su, Y. (2019). Characteristics are covariances: A unified model of risk and return. *Journal of Financial Economics*, *134*(3), 501–524.

Kelly, B. T., Pruitt, S., and Su, Y. (2020). Instrumented principal component analysis. *Available at SSRN 2983919*.

Lochstoer, L. A., and Muir, T. (2022). Volatility expectations and returns. *Journal of Finance*, *77*(2), 1055–1096.

Malmendier, U., and Nagel, S. (2011). Depression babies: Do macroeconomic experiences affect risk taking? *Quarterly Journal of Economics*, *126*(1), 373–416.

Zeisberger, S. (2022). Do people care about loss probabilities? *Journal of Risk and Uncertainty*, *65*(2), 185–213.